ELECTRONICS WORLD

APRIL 2006 £4.50

SPECIAL REPORT:

Software Development

Unveiling of the New Multi-Threaded MIPS Processor

Automatic Conversion Between Analogue and Digital Transfer Functions

Technology Review of the Bus&Board Conference in California

9 770959 833103 04>

New B² Spice V5
Our hottest Spice ever

New B² Spice Version 5 has all the power and functions you expect from a professional Spice package, but without the high cost:

- Real design flexibility with over 30,000 models, unlimited circuit size and a huge range of new virtual instruments
- New Circuit Wizard saves time by auto-generating many designs for you
- Sweep all parameters for any component and simulation type with the powerful new Scenario Editor
- Live Circuit feature allows values to be adjusted while simulations are running, displaying the results in real time

Professional standard Spice simulation for just £229 + VAT. Plus educational and multi-user licence discounts available and FREE comprehensive telephone technical support. **Try the full version completely free for 30 days.**

www.spice-software.com
Tel: 01603 872331

Research House, Norwich Road, Eastgate
Norwich. NR10 4HA. Fax: 01603 879010
Email info@looking.co.uk

RD Research

Contents

April 2006 Volume 112 Number 1840

Win a Microchip rfPIC Development Tool

See page 50

Front cover image: MIPS Technologies

WHAT BETTER CHOICE ?

FOR QUALITY, PERFORMANCE & AFTER SALES SERVICE

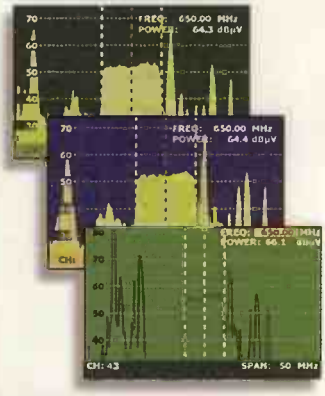

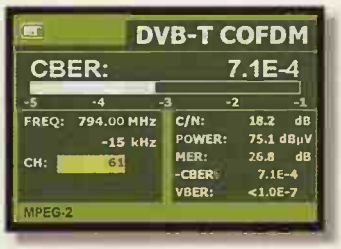

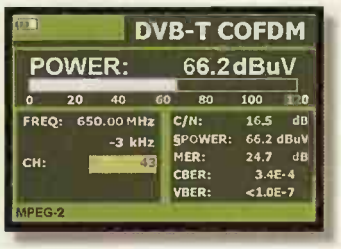

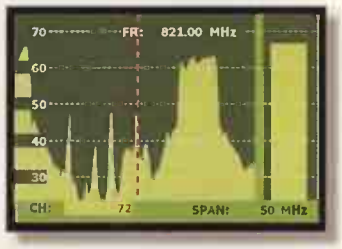

PROLINK 4C PREMIUM

PROLINK 3C PREMIUM

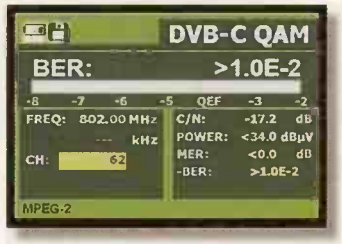

PRODIG 1+

PRODIG 2

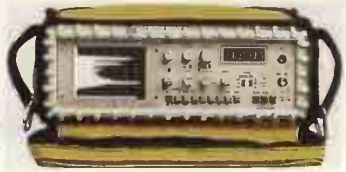

MC 577

Alban

ALBAN ELECTRONICS LIMITED

THE PROMAX SERVICE CENTRE

6 Caxton Centre, Porters Wood, St. Albans, Herts, AL3 6XT.

TEL : 01727 832266 FAX : 01727 810546

www.albanelectronics.co.uk info@albanelectronics.co.uk

SALES + SERVICE + CALIBRATION

Savings or just a hassle?

It's about time that high-tech and IT companies abandoned India and other far-flung places as their offshore bases.

For over ten years now, many firms, including services and utilities – and even financial institutions – have been flocking to set up offshore operations in India, the Phillipines and the like. As they stood in line to select the best location for their operations, they sang the praises of the local workforce: inventive, hard-working, skilled, and above all, cheap.

It appears that this is no longer the case, however.

The low-cost base has now turned into a higher-cost, unstable platform, which makes many western firms re-think their strategies. Hidden costs have come from the very places they praised – the workforce. They spend time and money on training, send western managers abroad to carry that out, and after all this, the workforce then leaves in droves after only several months of employment, for marginally better paid jobs elsewhere.

"Training takes time and it costs," said one high-tech executive. "Just about when you make them productive, they leave."

Then, there are the issues revolving around intellectual property (IP) – and in the case of utilities and financial institutions, confidential customer information – that gets "sold on" by local employees after they leave, for small amounts of money. The British daily press report whole databases being sold for as little as £1.

On the other side of the coin, we have the customers who are increasingly becoming dissatisfied with the quality of products and services they are receiving from western companies that have placed operations there. Twenty-four hour service is a good thing if done properly, but not if it leaves the customer frustrated, dissatisfied and angry to the point of giving up on the firm altogether.

Equally, when producing high-tech products, many have complained that not everything is done to spec; in some offshore operations, they are happy to approximate as long as they get the job done. But approximation does not come even close enough. It has to be perfect. This means, more costs in getting products right for the second, third, or indeed umpteenth time.

The lookout for cheaper locations continues, however. Some companies are planning to go even further afield, in places like Vietnam, for example; some are looking closer to home and hoping to get cheap operations in Eastern Europe.

But eventually, all of these benefits will disappear as cheaper-labour countries and their local workforce become more acquainted with the ways of the capitalist western world and how high-tech companies tend to do business.

Indeed, the stream of firms announcing their departure to warmer climates continues.

But, the trend is definitely slowing and, hopefully, at some point even come to a halt. Several high-tech companies have been decreasing their involvement there operationally and hoping to leave those areas for good.

Let's hope these firms are followed by the western utility companies and service organisations – it's time we got our western quality back.

Svetlana Josifovska
Editor

EDITOR: Svetlana Josifovska **E-mail:** svetlana.josifovska@nexusmedia.com **EDITORIAL E-mail:** EWeditor@nexusmedia.com
EDITORIAL ADMINISTRATION: +44 (0) 1322 611274 **E-mail:** EWadmin@nexusmedia.com
PRODUCTION EDITOR/DESIGNER: Jane Massey
SUBSCRIPTIONS: Customer Interface Ltd, Cary Court, Somerton, TA11 7BR **Telephone:** 0870 4287950, **Fax:** 01458 271146
SUBSCRIPTION RATES: 1 year: £45 (UK); €115 (Europe); $150 US & worldwide
DISPLAY SALES EXECUTIVE: Reuben Gurunlian +44 (0) 1322 611261 **E-mail:** reuben.gurunlian@nexusmedia.com
PRODUCTION EXECUTIVE: Dean Turner +44 (0) 1322 611206
PRINTER: William Gibbons Ltd • **ORIGINATION:** Impress Repro A1 Parkway, Southgate Way, Orton Southgate, Peterborough, PE2 6YN
NEWSTRADE: Distributed by Seymour Distribution Ltd, 86 Newman St, London W1T 3EX. • **PUBLISHING DIRECTOR:** Tony Greville
If you are experiencing problems getting copies through your newsagent, please call Debbie Jenner on +44 (0) 1322 611210

ISSN 0959-8332

ABC

NEXUS MEDIA COMMUNICATIONS

Electronics World is published monthly by **Nexus Media Communications, Media House, Azalea Drive, Swanley, Kent, BR8 8HU** Nexus Media Communications is a trading name of Nexus Holdings Limited. Registered in England. Registered Number 5346404. Registered Office: Hanover House, 14 Hanover Square, London W1S 1HP

Class AB and Class D drivers appear in a single codec

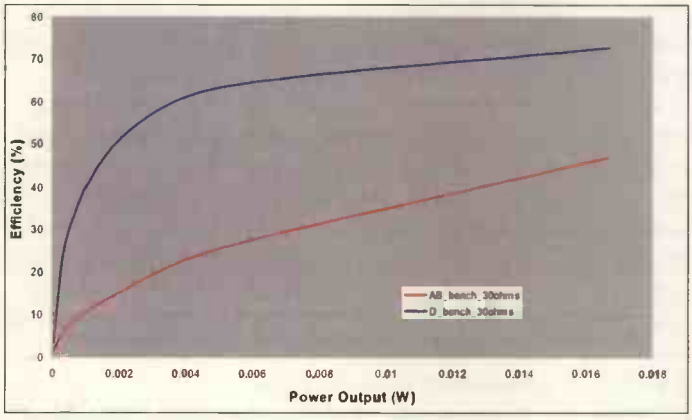

Headphone driver power consumption reduced by up to 50%

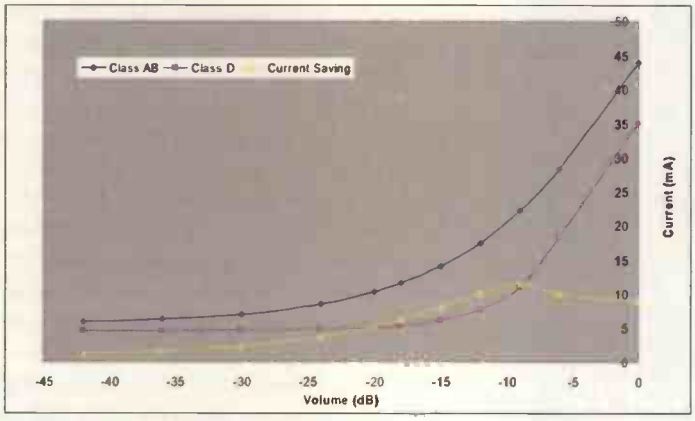

WM8985 Power Consumption

As mobile phone companies continue to pack in various applications into a single handset, such as MP3, FM, even, mobile television capabilities – the power drainage on the battery becomes a key requirement for designers to work on. Edinburgh-based Wolfson Microelectronics has come up with a novel – and some would say unconventional – way in solving the power issue in portable devices, by combining Class AB and Class D amplifiers in a single, switchable stereo codec.

"MP3 decode is consuming lower power than it used to, for example. The same goes for headphone drivers also, but to crank up the [sound] volume, you'll need higher current and that's more power. Class AB headphone driver is good for quality and low distortion but it is not so efficient when it comes to power – a bit like an LDO. The solution is to use a more efficient driver technology and the most efficient is Class D switching, non-saturating technology used for high-power audio systems," said Mark Jacob, strategic marketing manager at Wolfson Electronics.

Wolfson's solution – the WM8985 – uses a closed loop Class D architecture that has a power supply rejection ratio (PSRR) of -70dB at a 3.3V supply voltage. The device then becomes a switchable Class AB/D headphone driver that switches between modes seamlessly, depending on the application the phone is in. There are a few milliseconds when the two modes could be used together, for example when the phone is used in audio mode but it still needs a burst of power to communicate with the base station.

"With this design, you can switch between Class AB and Class D at the same times and you will not notice the difference in [sound] volume [when listening to audio]," added Jacob.

The device includes a built-in audio-enhancement DSP for wind noise filter, notch filter, 5 Band EQ and 3D audio. It supports up to two differential microphone and analogue stereo line inputs, enabling direct mixing from an FM radio, and has two pairs of headphone drivers enabling two users on one audio player. An integrated PLL supports input clock between 9 and 27MHz.

Jacob says that the WM8985 is a unique offering to mobile phone designers and a turning point for the industry. "There's been a reluctance [so far] to use Class D in mobile phones for audio [applications] such as headphones amplifiers, but that's changing," said Jacob.

Odyssé shuns black and white

French firm Aures Technology has launched a colourful series of POS terminals for the retail and hospitality markets. There are no less but eight differently coloured fascias that simply clip onto the terminal itself.

In addition to its colourful appearance, Odyssé is technically advanced too. At its heart is the Intel M (mobile) processor and a hard disc drive with a capacity of between 40 and 80GBytes, situated in an accessible, removable drawer. The system has two USB retail ports (12V and 24V) and 14 communication ports, enabling the connection of a complete set of POS additional equipment and peripherals. Odyssé is Wi-Fi compatible via the mini PCI slot.

Aures plans to launch a range of complementary products in matching colours, including printers and portable payment systems.

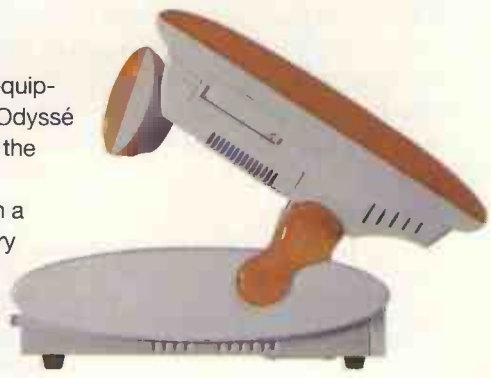

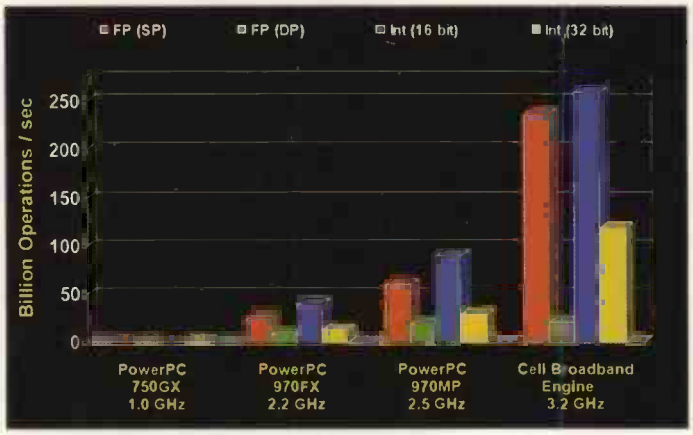

Performance comparisons of the Cell processor with others

Cell processor is now aimed at embedded applications

Mercury Computer Systems, one of the leading board and system makers in the US, is looking to push IBM's Cell processor into embedded applications, such as medical imaging and video inspection and surveillance systems, in both the US and Europe. The Cell processor already powers the Sony PlayStation 3. Being produced in high volume will help bring its cost down for embedded systems.

This processor uses a new variant of the PowerPC core with the AltiVec DSP extensions and eight slave cores for "number crunching", particularly for video applications such as rendering data from a scanner in close to real time. IBM also has a demonstration of the Cell processor rendering satellite data into a 3D video 'fly-through' in real time.

To help system developers, Mercury has developed an evaluation system for embedded applications running YellowDog Linux and Eclipse-based open source compilers, and supporting maths libraries and its own middleware software with own APIs.

IBM is still keeping the power consumption of the processor confidential, which makes it difficult to assess its suitability for embedded systems, but a system with the Cell processor, two PowerPCs and PMC mezzanine cards consumes around 400W, says Mercury. This puts the power at around 90W, with around 50Gflops of processing at 3.2GHz, compared to a single core Opteron with 60W at 2.2GHz.

However, some equipment makers are skeptical about the roadmap of the Cell processor and its support over a period of 10 to 15 years needed for industrial systems. There are also no bridge chips available to link it to other parts of the system. Toshiba has a design, but only for internal use, and IBM offers to develop a dedicated ASIC for it through its services business.

However, support for this processor is mounting from other parts of the industry. French equipment maker Thales Computer is working on a design, which is still at the R&D stage.

"It's a bit early; it's a mid-term project," said Joe Eicher, US engineering manager at Tales Computer.

"There's a lot of software to be ported."

• Nanowatch • Nanowatch • Nanowatch • Nanowatch • Nanowatch • Nanowatch • Nanowatch • Nanowatch • Nanowatch • Nanowatch • Nanowatch •

Microscience, Europe's largest microscopy and imaging event, this year will also focus on nanomaterials and nanotechnology. In the keynote session, chaired by Peter Dobson, American and European scientists will discuss how new nano-scaled materials and nano-scale structures can be exploited in new products.

Research sponsored by the National Science Foundation, NASA Vehicle Systems Program, the US Department of Defence and the Chinese Academy of Sciences has found a previously unknown zinc oxide nanostructure that resembles the helical configuration of DNA could provide engineers with a new building block for creating nanometre-scale sensors, transducers, resonators and other devices that rely on electromechanical coupling. The "nanohelix" structure is part of a family of so-called nanobelts, which have semiconductor and piezoelectric properties.

Limited and Fujitsu Labs have developed a carbon nanotube-based heatsinks for semiconductor chips. It has been found that the use of carbon nanotubes as heatsinks for high-frequency high-power amplifiers successfully achieves heat dissipation and high amplification simultaneously. Carbon nanotubes exhibit superior thermal conductivity and are suitable for use in bumps in the flip-chip structure. The technology uses an iron catalyst coating to grow carbon nanotubes to a vertical length of at least 15 micrometers on the wafer substrate.

NEC Electronics, Sony and Toshiba will jointly develop system LSI process technologies for the 45nm generation. As applications in the digital consumer, mobile and communications areas evolve, there are greater requirements for advanced semiconductors to achieve higher performance and functionality, such as high-speed data processing, as well as lower power consumption and smaller chip dimensions.

The three firms' joint work will take place at the Advanced Microelectronics Center.

Fujitsu is planning a new fab for producing logic chips in 65nm process technology and 300mm wafers. The fab will be constructed at Fujitsu's Mie semiconductor plant in central Japan.

This is Fujitsu's second 300mm fab with capacity for 2007, projected at 10,000 wafers per month (wpm) and maximum capacity reaching 25,000wpm. Operation starts in April 2007.

Researchers at Toshiba Research Europe (TREL) and the University of Cambridge have discovered that a simple semiconductor device can generate light possessing quantum entanglement. This could lead to long-distance, highly secure optical networks, more sensitive medical diagnosis, more powerful computer chips and scalable quantum computing.
Unlike normal light in which the photons (the 'particles', or quanta of light) can be regarded as distinct, the new source emits a stream of photons in pairs at regulated times with 'entangled', or interrelated, properties. The new entangled photon source is similar to an ordinary semiconductor light source, but contains a tiny, nanometer-sized quantum dot that emits the coupled photons.

Ω

The Centre for Integrated Photonics (CIP), Bookham, Epichem, Loughborough Surface Analysis, the University of Sheffield and the University of Surrey have been jointly awarded £1.7m by the UK's DTI to develop new technologies for uncooled operation of advanced photonic components. The 2.5-year project called ETOE (Extended Temperature Optoelectronics) will focus on the development of advanced InP-based photonic materials and devices in two strains. The first is the development of active devices containing aluminium, to enable the high temperature operation of a range of advanced devices, including fixed frequency and widely tuneable lasers, semiconductor optical amplifiers, superluminescent diodes and avalanche photo diodes. The second is the development of improved processes for the MOVPE growth of semi-insulating current-blocking layers using ruthenium doping to enable higher speed modulation of devices.

Ω

ARM and Handshake Solutions (part of Philips Electronics), unveiled a new clockless IC processor. The ARM996HS is based on a self-timed circuitry, which has been used in hundreds of millions of chips for smartcards, advanced pagers, in-vehicle network transceivers and cordless handsets. By removing the clock and associated architecture of standard ICs, clockless designs offer significant reductions in power consumption and EMI. As such, the processor will be pitched at automotive, medical and embedded control applications.

Not just a pretty face

Tektronix's latest launch of a family of oscilloscopes – the DPO 4000 series for the low end of the market – holds some great new innovations. Apart from being a slick, thin (only 13.7cm in depth), lightweight (5kg) portable system with a large viewing screen, it also has additional features that will help engineers in pursuit of elusive glitches in highly complex signals pinpoint them quickly and with ease.

A feature called Wave Inspector provides the search for certain points/data within the record memory. The controls for it are on the instrumentation panel itself, unlike with other units where the user needs to go into a set of menus to identify and zoom in on certain events in the memory.

"Interestingly, the Wave Inspector turning knob on the front panel was an idea of our DPO7000 series scopes [development] team. The two teams [DPO4000 and DPO7000] have been discussing ideas and the DPO7000 was so close to launch that the idea could not be used there, but then the DPO4000 team took it and run with it. So, here it is on this scope and a completely unique and new idea on the market," said Pete Derby, applications engineer at Tektronix.

Slick, light and portable – the new DPO4000 series scopes are packed with features

Wave Inspector further allows the marking of events of interest. A search menu, covering pulse width, runt, bus, rise/fall time etc, allows the user to specify the parameter for each and the instrument automatically finds all events that meet that search criterion. Events could easily then be compared and zoomed into or out of.

The DPO4000 series also has capture and decode of a long stream of serial data for three buses: CAN, I2C and SPI.

With a little help from National Instruments, the scope comes with a signal Express version that allows users to view and handle the scope from their PC.

High-tech firms downscale development in India

Costs of being based in India have soared to the point that western firms are now reducing their operations there or pulling out altogether. Among them are EDA firm Mentor Graphics and AMI Semiconductor.

"Costs of being in India have risen by a factor of five in the last ten years," said Hanns Windele, vice president for Europe at Mentor Graphics. "It's not worth being there any more."

One main reason for the rising costs is the high employee attrition rate. India is notorious for having employees leave jobs for the smallest increment in salary elsewhere. "We recruit them, train them and in six months, just as they become productive, they leave for a job round the corner that is only slightly better paid," said Windele. "We are not investing or growing in India any more."

Similarly, European firm AMI Semiconductor inherited operations in India from Alcatel through an acquisition some ten years ago. Alcatel used to develop ADSL products there. "Our centre in India is also becoming smaller," said Sven Ake Shonborg, VP of mixed signal products for Europe and Asia. "Alcatel had 100 engineers there at one point, now there are only about 30."

Instead, Mentor Graphics and other high-tech firms have turned their eyes to Eastern Europe, where they find a highly educated workforce that is unlikely to change jobs as quickly. "The level of education [in Eastern Europe] is very high and they don't leave you. They are also closer to our customers and markets," said Windele.

Conglomerat Wipro's main role is labour arbitrage in India, where it initially helped western firms with labour outsourcing, followed by outsourcing of call centres and then the outsourcing of processes. Kees Ten Nijenhuis, VP of sales at Wipro Technologies, confirmed that Indian workforce changes jobs frequently but that this is a problem that Nasscom, a trade body and "voice" of the IT software and service industry in India, is trying to address.

Business will lead the way in demand for electronics

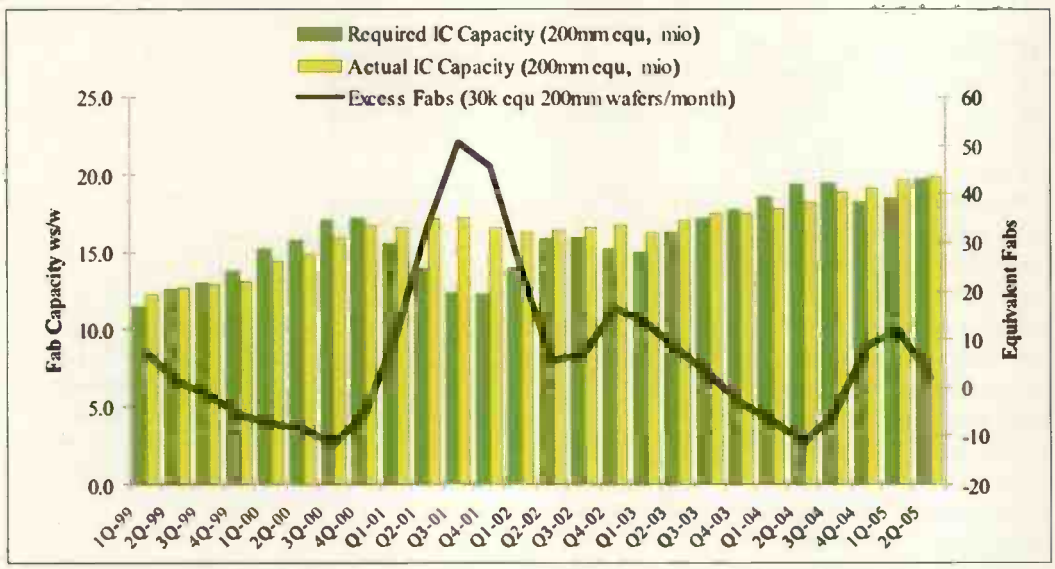

Is semiconductor capacity shortage on the horizon?

The electronics industry will change from being consumer-driven to business-driven as we enter next year. This is the forecast of Malcolm Penn, CEO and key analyst at research house Future Horizons. "Business demand crashed in 2001 and it hadn't improved so much until now. There haven't been any serious investments [made by businesses]. But, profits are growing, businesses are getting stronger and more confident and cheque books are getting opened," he said. "We will see investments in everything for solid manufacturing; in equipment, infrastucture, networks, employing people – everything."

"Consumer [oriented spending] will not disappear though – it will just slow and it will be overtaken by business spending," he said.

In the semiconductor business, supply is too close to demand for comfort, warned Penn. "Capacity is tight. If demand increases, there'll be supply problems," he said. "If supply is constrained, demand begins to have a distortive effect [on the industry]. So far, investment in capacity is not happening. I don't know what it takes to tell people [in this business] that they should start investing in capacity."

ARC launches a video platform

Customisable processor core developer ARC International has taken its 700 core and added specific instructions to it for handling video, all fitting in a 260k-270k gates.

"We've added 105 certain instructions and the SIMD [single instruction multiple data] macro we've added to it [the core] accelerates those instructions," said Derek Meyer, marketing VP at ARC.

The 700 core also has the memory management unit (MMU), which a designer can easily do away with if required, which saves some 10k gates.

"We are seeing customers saying 'keep power low even if the gates go up, then we can drive down the clocks', [the number of gates affect silicon estate and power]. Power is a lot more important than space."

Depending on the application, ARC's video sub-system for a 750 core consumes 44mW when decoding D1 H.264, which is considered low in the process-intensive video applications.

ARC Intentional already has an equivalent audio sub-system, which it launched last year. Its roadmap includes further advances made for that subsystem as well as for the newly-launched video sub-system, which will eventually focus on high-definition (HD) applications.

The use of active RFID is growing rapidly, says a new research from IDTechEx. According to its latest report, the value of sales of active systems, including the tags, will grow from half-a-billion US dollars in 2006 to nearly $7bn in 2016. The three primary technologies fuelling this growth will be real-time location systems, disposable RFID sensor systems, including ones in the form of Smart Active Labels (SALs), and sophisticated multi-functional devices.

Ω

European patent attorneys are currently in talks with the Chinese government officials over intellectual property (IP) matters. The Chinese delegation wants to learn more about the patent systems in Europe and, in particular, the way in which patent attorneys are regulated. As such, they are visiting the UK, France and Germany, where they meet with the local patent attorney associations.

Ω

It is expected that in 2006, US businesses will spend more on wireless voice services than on wireline, says research house In-Stat. Expenditures by enterprise firms (1000 or more employees) on wireless data will grow an average of 18% per year through 2009, reports the firm. According to its researchers, the transformation of telecom spending continues unabated since "the relentless march toward IP-based networks, the unequivocal adoption of mobile solutions and the pervasiveness of broadband have changed the face of business networks".

Ω

China is preparing to roll out its third-generation digital phone service. Hisense Communication, a major handset manufacturer in China, is developing the HT28, a trial TD-SCDMA handset. TD-SCDMA is the Chinese home-grown 3G standard among the three international 3G standards. The on-going TD-SCDMA pre-commercial tests will have a direct impact on the development and commercial operation of 3G in China, so the tests are being watched closely by many leading handsets vendors both in China and abroad.

RF test instruments use software-defined radio architecture

Keithley Instruments launched its first ever vector signal generator, part of a new line of RF test instruments. The 2910 RF vector signal generator uses software-defined radio architecture and off-the-shelf components, such as readily available DSPs and FPGAs.

The software-defined radio architecture replaces functions that used to be done in analogue. Keithley's executives say that this adds speed and flexibility to testing but without the cost associated with customised hardware.

As such, the 2910 settles in 1.5ms and the sync-out source-settled indicator is only set when the source is settled. This eliminates the need to slow down tests by adding wait states to ensure the signal generator has settled. Also, its 64 mega-sample Arbitrary Waveform Generator supports simultaneous loading of multiple signal waveforms and enables switching between such signals in less than 5ms.

Fast switching among waveforms is a critical attribute for reducing test time in devices operating within multiple wireless standards, just as faster execution times increases throughput.

The new series of instruments is aimed at all types of RF applications, including establishing protocols for wireless connectivity through WLAN, Bluetooth and WiMax, as well as RFID. "You'll find that the people employing these new types of RF technologies are new to them; this is not the old military crowd [dealing with RF] any more. That's why the instruments have to be flexible, easy-to-use, low-cost and all the time offer high-quality measurements," Walter Strickler, marketing director at Keithley Instruments.

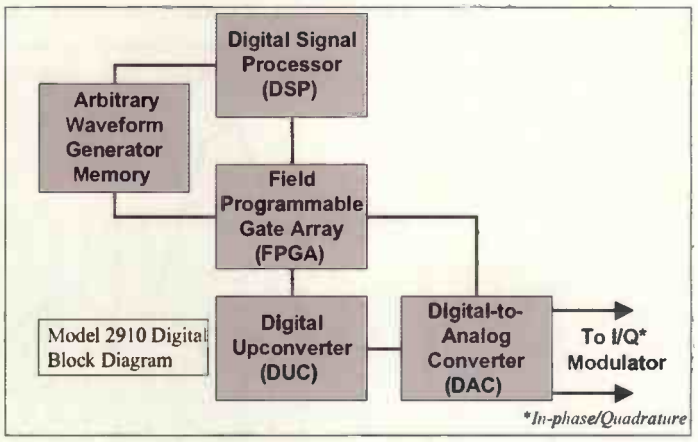

Block diagram of the 2910 RF vector signal generator from Keithley Instruments

In the summer, Keithley plans to launch two complementary products – the 2810 RF vector signal analyser and the 3500 portable RF power meter.

The price for the 2910 model starts from around £9300.

QinetiQ pushes FPGA interconnect in VME

The embedded board business of UK defence group QinetiQ has teamed up with its US partner TEKMicro to sponsor a new standard for linking FPGAs in VME systems.

The VITA55 standard will use the Aurora protocol that is used in Xilinx FPGAs as a standard way of connecting devices, particularly creating a 'mesh' of processors in a VME rack.

QinetiQ has been working with US partner TEKmicro on the development of systems using Aurora, such as TEKmicro's new Jazz system, and has developed a prototype mesh system with backplane maker Elma Bustronic.

Aurora tends to be a point-to-point protocol rather than a bus or switch fabric technology, but being the native protocol for Xilinx FPGAs means there is no need for interface or translation chips and has a good latency. It uses a simple FIFO interface to the DSP cores in the FPGA and supports multiple streams on a single link. It would use around 1000 logic slices for 1X interface, plus a block RAM for each FIFO stream, or around 2000 slices for a 4X interface with two block RAMs per FIFO stream.

The VITA55 group is sugesting that Aurora can also be used in other VITA standards such as VXS and VPX as a low overhead point-to-point interconnect protocol for DSP systems, as many DSP cards use Xilinx FPGAs anyway.

QinetiQ applied the technology in an adaptive beam-forming antenna system, using the Aurora links to feed back the weights for each of the radar segments. These were all connected up on a VXS (VITA41) standard backplane.

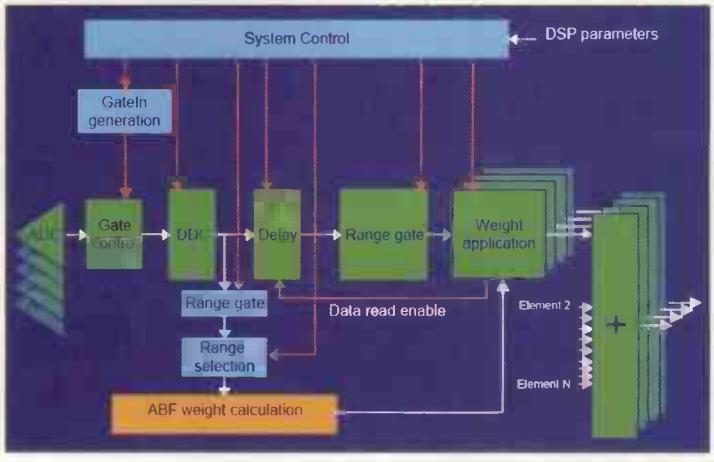

Using Aurora for an adaptive beam forming antenna

Next destination for GPS: mobile phones

Swiss fabless firm NemeriX is readying itself for the era of the GPS-enabled mobile phone with the launch of a low power, small size, inexpensive receiver. The NJ1836A GPS IC, implemented in a 0.13µm TSMC CMOS process, is only 7 x 9mm in size. Although NemeriX's CEO Vincent Mouret would not state its price, he said that most solutions on the market today currently sell between $8 and $9, which, he added, is a price a lot lower than before.

"GPS was not consumer oriented [until now]. It was aimed at the automotive sector mainly. Only now there're solutions that can be consumerised," he added. "The market will start seeing [GPS receivers] at below $5 in 2007."

NemeriX uses its own power management techniques to lower the power in the receiver to 25mW per fix. The engine runs on the Ceva DSP and it typically shares the memory with the CPU.

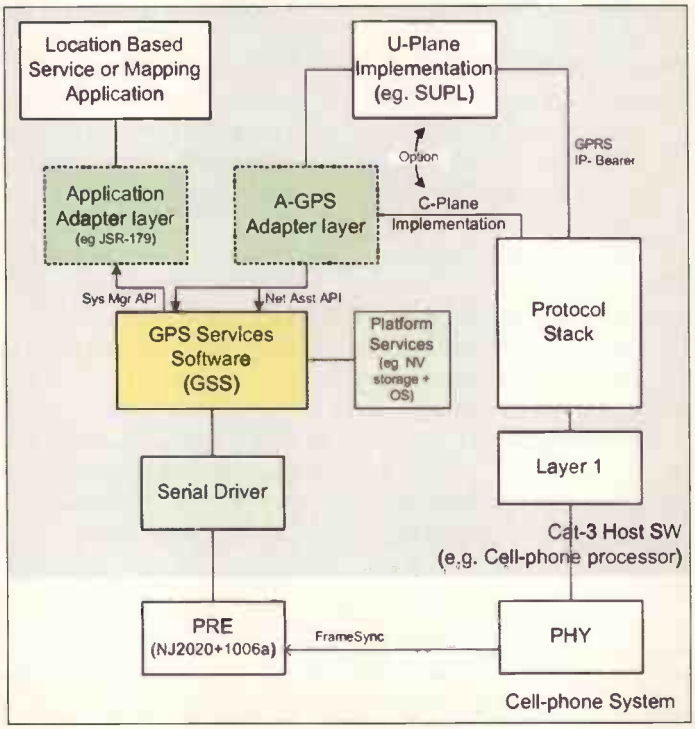

Block diagram of NemeriX's GPS engine

Accuracy is 3-5m outdoors and 20m indoors. The next solution from NemeriX will be a single chip GPS receiver, with a fully integrated RF and baseband, in a package of 6 x 6mm. It is likely to be implemented in a 90nm CMOS process.

"People are already working on mobile phones with integrated GPS," said Mouret. "GPS will be in most handsets by 2009-2010."

That estimate is even more optimistic from Chris Ryan, principal analyst at research house Future Horizons. "I think it is only two years away [from now]. It is not the technology that's the issue [to have these phones on the market sooner] but how the operators will extract money for such services out of users. You'll have to have the business model for such services in place first."

Board makers split over future VME standards

Board makers are split over which of the coming switched VME standards to back with products.

One, Curtiss Wright Embedded Computing (CCEW), is committing fully to one standard, VPX, for its high performance VME boards for military and aerospace applications, while others are offering a range based on the previous VXS standard.

The VPX standard (VITA46) uses all the pins on the backplane for a switching fabric – such as Rapid I/O, Infiniband or Ethernet – rather than a bus structure. This means that it cannot be used in the same backplane with existing VME64 boards. But alongside traditional 6U systems it is also well suited for a 3U form factor to challenge CompactPCI systems, particularly in vehicle and aircraft equipment.

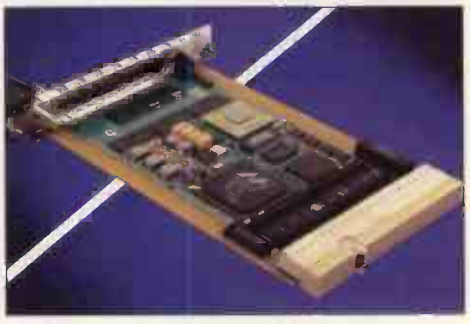

Curtiss-Wright's SCP/DCP-124

On the other hand, the VXS standard (VITA41), launched last year, uses the existing P0 connector in the middle of a VME board for switch connections, so it can be used with existing VME boards. But the P0 connector is limited in the number of connections it can offer and many equipment makers use this for their own protocols. One of the problems has been the slow launch of VXS cards, particularly switch cards, but these are now coming to the market from companies such as Mercury Computer Systems, which has launched a family of VXS products.

While CCEW will continue supporting its one existing VXS board, it will not be launching any new ones, said Stuart Dewar, product manager, and all its focus will be on VPX instead. It has announced a 6U VPX board, the VPX6-185 that will ship later in the year. This is a nominal backplane bandwidth of 8GB/s via four Advanced Switching Interconnect (ASI) ports, two PCI Express VITA 42 XMC/PMC sites and Freescale's 8641 single/dual-core PowerPC processor.

However, some companies are seeing a hybrid approach. Backplane maker Elma Bustronic has developed a custom system that has slots for VME/VXS and VPX on the same backplane, allowing both standards to be used in a single, custom chassis.

Selecting User Interfaces

- Try to understand the users of the system and how they are likely to use the technology. For example, Europeans prefer direct contact using mobile phones rather than leaving voicemail messages. The opposite is true for North Americans.
- Learn as much about the environment in which the technology will go into. For example, in a large warehouse, if pallets were not barcoded, the employees would need to type in all the details every time a pallet is moved.
- Learn as much about the likes and dislikes of the company where the technology will be used. For example, the purchasing department may always buy a particular make of a PC, screen, keyboard and mouse.
- Establish how easy to learn the device needs to be. Direct pointing devices, like a stylus, tend to be intuitive to use as touching and pointing come naturally in comparison to using an indirect pointing device, like a mouse for example.
- Establish how accurate the device needs to be. Touch screens are generally inaccurate, in which case you may have to increase the size of the screen elements.
- Establish how much time the user will spend on the system. Graphics tablets are less tiring than touch screens, and joysticks require a wrist rest if used for lengthy periods.
- Establish how much space is available. The trackball, joystick and touch tablet require very little room, but you need to consider what the users might prefer.
- Establish how robust the device needs to be. Pen systems are not as good in public-access systems as touch screens are; the stylus could be broken or stolen.
- Establish how manually dexterous the user is. Children, for example, may find it difficult to use devices that require a high degree of accuracy.
- If you have to use alternative approaches to outputting information, such as head-up displays, head-mounted displays and stereoscopic displays, don't forget to train those who will need to use them, such as firefighters for example.

This month's Top Ten Tips were extracted from the book '*User Interface Design and Evaluation*' by D Stone, C Jarrett, M Woodroffe and S Minocha and published by Morgan Kauffmann, an imprint of Elsevier.

To evaluate this book, please contact the editor at

EWadmin@nexusmedia.com

WiMax Has Landed

Gilles Karolkowski outlines that, despite years of speculation, WiMax has finally emerged with many commercial deployments already up and running across the globe

The benefits of fixed WiMax (802.16-2004), especially in emerging markets and rural areas, are as an altemative to deploying wired broadband infrastructure. It's now commonly accepted that WiMax will help spread broadband to more users more quickly, complementing existing traditional technologies. For consumers and businesses, it will give people further choices about how and whether they can access broadband Internet. In addition, the introduction of a competitive 'third broadband access pipe' to DSL and cable will lead to prices being driven down.

But where are we at today? The hype surrounding WiMax has built steadily, but are there any deployments, how is the technology being used and is it bringing benefits as expected?

Carriers are now rolling out full commercial deployments of fixed WiMax solutions in cities as well as suburban and rural communities, allowing broadband wireless networks to reach locations where previously they were either impossible or too costly to pursue.

These deployments support a range of uses, from basic high-speed access for homes to Internet telephony, business connectivity and support for schools and government offices. For example, VoIP services are being offered for the first time in the Dominican Republic and home users in Austria. Residential and small business users now have previously unavailable access to high-speed Internet, from the Philippines to Spain. Schools and government offices in Dublin now benefit from cost-effective, high-speed Internet access, and in Poland, WiMax services are on offer to public safety, local government offices and schools.

At the recent annual Tour de France cycle race, hundreds of international journalists and other accredited people enjoyed quick and easy Intemet access via a WiFi hotspot linked to a WiMax backhaul. This provided low-cost, broadband Internet connections to the heart of the pressroom, from which news were sent out to editorial offices across the world. Broadband symmetric connections enabled quicker and easier uploads compared to DSL, allowing the press to send multimedia information quickly and efficiently.

For each new stage of the race, staff were able to install all of the technical equipment overnight at each of the new venues, because WiMax brings portability and fast installation, typically within one hour. Without disruptive cable laying, it means that technical staff can easily set up WiMax capabilities for similar mobile events.

In Spain, trials carried out in a network of 300 base stations and 20,000 subscribers have recently been completed, kicking off migration to Intel technology and ensuring that WiMax technology is fully optimised for commercial public WiMax services in that region. Initially, Iberbanda, a Spanish broadband telecommunications operator, will deploy fixed WiMax technology, updating the WiMax-ready trial networks that have already been deployed across 25% of Spain. The WiMax technology provided telephony services and broadband Internet access across areas of over 30km and at speeds of up to 10Mbps per second.

Another example can be found in Sweden with recent trials in the municipality of Skellefteå, in co-operation with Intel and MobileCity. The project's success has paved the way for other regional governments to accelerate the deployment of wireless broadband access to several municipalities outside of Skellefteå.

In the UK, at the Science Museum in Wroughton, where many of the museum's large object collections are housed in 11-old aeroplane hangers across a 545-acre site, WiMax is allowing the curators at the various sites to be able to communicate electronically with each other for the first time ever and also keep updated records of the various artefacts within the remote buildings. As a result, the Science Museum is now looking to move forward on new plans that will allow its collections to be more widely accessible to the public and provide greater benefit to the museum staff.

In Indonesia, Intel has installed a large wireless broadband 'umbrella' that lets humanitarian and disaster relief groups in the tsunami-hit Banda Aceh communicate with each other and the outside world. In June, local engineers switched on an Intel-supported pre-WiMax network that today covers some 1,500km^2 of Aceh province, where the tsunami wreaked the greatest destruction.

The network consists of three pre-WiMax base stations providing high-speed Internet data connection at speeds up to 6Mbps within the coverage area and 28Mbps backhaul connections between base stations and connections to multiple VSATs.

The WiMax network is making it possible for the many organisations offering disaster relief to reach the local population with the help they need.

So after much speculation and analysis, it's clear that the fixed WiMax version has begun the transition from an exciting emerging technology to one that starts bringing real benefits to businesses and consumers in its first commercial deployments.

Gilles Karolkowski is Marketing Manager for WiMax at Intel

Israeli Fabless Semiconductor Start-Ups Find Their Niche

By Joel Bainerman

"Cutting edge" has always been the defining term to describe what was unique about Israeli R&D. Creative and cost-effective R&D has always been considered the core and crux of Israel's relative advantage in international high-tech. Since entering the high-tech market in the late 1970s, Israeli companies and Israeli-based R&D centres of foreign multinationals have produced some of the leading advances in microelectronics and chip design.

One of the reasons why Israel has done so well in these fields is due to the training all Israelis receive as part of their military service in the elite unites of the Israel Defence Forces (IDF). The training methods employed by the IDF emphasise flexibility and the ability to adapt to changing circumstances is encouraged – a critical component when working on R&D for the next generation of microprocessor. IDF officers are trained to analyse and comprehend a situation in the broad context and to identify solutions to solve specific problems and challenges. IDF training also encourages its recruits to be goal- and project-focused.

All of these attributes, when taken into the civilian R&D arena, contribute to the success of Israeli R&D.

Israelis are also very good at blending together various disciplines – once again, due to their army training. When serving in the army, the IDF engineers are taught to blend together well from different technological disciplines, such as hardware and software. Thus, when the electronics software engineer works with the software engineer, both know something about the other's field of expertise, which provides for a smooth working relationship between the two disciplines.

In the past it was Motorola, National Semiconductor, Intel and other multinational electronics firms that benefited from Israeli R&D. Today a growing number of local start-ups, many of them who once worked for the multinationals, which are spotting the landscape of Israel's Silicon Valley – a 100km stretch on Israel's Mediterranean Coast between Tel Aviv and Haifa. While many of them have tapped into US public equity markets, such as Zoran Corp, Audiocodes, Metalink, Orckit Communications and Saifun Semiconductors, a whole new crop of start-ups are making their presence felt in electronics and communications markets worldwide.

The up-and-rising stars of fabless semiconductor start-ups developing products for the communications market include: Color Chip, CopperGate, Provigent, Passave Technologies, EZchip Technologies, TeraChip and Wintegra among others.

ColorChip's core technology is based on a proprietary Ion Exchange (IE) process in glass substrate, which forms planar structures of circular shaped waveguides. These ultra-low-loss fibre optic components and modules provide system vendors and integrators with advantages in cost, performance and footprint, insertion loss and polarisation sensitivity. Light guiding is achieved by producing a high refractive index core in the glass substrate. The company claims almost zero coupling loss to standard SM optical fibre.

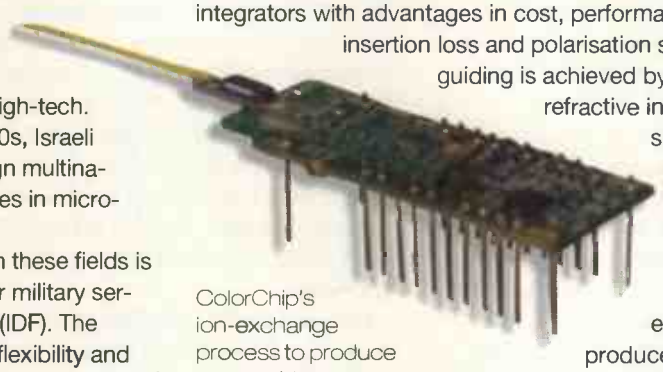

ColorChip's ion-exchange process to produce waveguides embedded in a glass substrate is essentially a system-on-glass

"ColorChip's ion-exchange process to produce waveguides embedded in a glass substrate is essentially a system-on-glass, which is particularly well suited for fibre-to-the-home applications, where transceivers like this would be at every end-user's location," says CEO Moshe Price. "The critical element at the home is going to be the optical transceiver, the element that converts the incoming data format to a format usable in the home."

Price says the cost advantages of the ColorChip approach make it an attractive technology for many elements in optical telecommunications networks because each end-user needs a transceiver.

A developer of fibre to the home (FTTH) semiconductors, Passave Technologies is preparing for large-scale commercial installations. It develops and manufactures systems-on-a-chip and integrated solutions for broadband fibre optic access. Passave CEO, Victor Vaislieb, says the FTTH market is waking up. "After several years during which there was only talk about the potential of this market, we're now seeing it wake up for real. Leading Asian communications operators are rushing to adopt the new EPON standard. FTTH broadband service is now available at speeds 1000 times faster than ADSL and cable."

EZchip is a fabless semiconductor company providing high-speed network processors. The company has joined forces with another Israeli start-up, Dune Networks, a provider of switching fabric and traffic management solutions. EZchip's breakthrough TOPcore technology enables integration of packet processing, classification search engines and traffic management at 10-Gigabit throughput on a single chip. EZchip's single-chip solutions are used for building networking equipment with dramatic savings in chip count, power and cost.

The company has developed a new breed of network processors that integrates multiple TOPs (Task Optimised Processors) rather than generic RISC (Reduced Instruction Set Computer)

processors, to provide the performance required for next generation networking products. With its Task Optimised Processing technology – TOPcore – different types of TOPs are tailored to perform a specific networking task. These fast and efficient processors are then integrated into a super-scalar architecture, designed and optimised for packet processing.

The technology integrates both packet processing and classification into a single chip with no need for external classifiers, CAMs or SRAMs. Systems based on EZchip's network processors require only a few low-power, low-cost DRAM for the lookup tables.

Together, the two Israeli companies will offer a combined chipset that provides a flexible and scalable solution for implementing line cards and switch blades. The chipset is currently in design by several customers for deployment in carrier-class switches and routers.

A typical line card configuration of the joint solution consists of EZchip's NP-2 network processor and Dune Networks's SAND-FAP10/20V traffic manager and fabric interface device. The devices are then interconnected via an industry-standard SPI4.2 interface.

> *"After several years during which there was only talk about the potential of this market, we're now seeing it wake up for real"*
> **Victor Vaislieb, CEO, Passave Technologies, speaking about the FTTH market**

Michal Kahan, Dune Networks' director of marketing, said: "NP-2's processing capabilities together with the FAP10V's traffic management capabilities provide a solution, addressing carrier's requirements in Metro, Core and Edge markets. This enables mutual customers to enjoy the best of breed of NP, TM and fabric technologies."

Amir Eyal, EZchip's VP for business development, added: "Networking system vendors will gain the processing flexibility to add new features and support evolving standards, as well as scale their system designs to Terabit switching capacity. To enable the emerging converged network, carrier-class equipment needs to be flexible and able to guarantee bandwidth not only per subscriber but also per each subscriber-application."

EZchip has also hooked-up with TeraChip, another Israeli fabless semiconductor company providing high-performance switch fabrics. The solution, which integrates the TeraChip 160Gbps TCF16X10 switch fabric with EZchip's NP-2 10Gbps network processor (NPU) with integrated traffic manager, will enable efficient designs for carrier-class switches and routers.

The two Israeli start-ups believe that network system providers are looking for true multiservice architectures that can meet end-user performance requirements and service-level agreements. Supporting converged triple-play services on a single infrastructure requires not only high-bandwidth switching, but also advanced traffic management to deliver the required QoS to applications.

Using TeraChip's high-capacity shared-memory switch fabric and EZchip's NP-2, which combines packet processing and traffic management, system designs require fewer devices on both switch fabric boards and line cards, resulting in lower power consumption and reduced form factor. The combination of NP-2's advanced flow control with TeraChip's highly scalable multicast capabilities enables solutions especially well-suited for the delivery of IP video services.

CopperGate, on the other hand offers something different. Its CopperStream technology addresses the emerging IPTV market with its home networking distribution technology. The company is the sole provider of the HPNA V3 standard chipsets, which utilise existing coax cables and telephony wiring in the home.

The core technology exploits the capacity of telephony and coax wiring by delivering broadband services at data rates of up to 128Mbps to multiple network nodes. It includes a robust frequency-diverse modulation scheme, which achieves high data rates on a variety of wiring topologies, overcomes various impairments and coexists with other services on the same wires, such as POTS, dial-up modems, ISDN and ADSL. The embedded synchronous Media Access Control (MAC) protocol layer allows distribution of multiple HDTV streams, digital audio and toll-quality voice streams within the home.

Also in telecom domain is ProviBand from Provigent. Provigent develops broadband wireless transmission technologies that allow two data streams on two polarisations of a single channel to be transmitted simultaneously, thus doubling net channel capacity. The chips enable transmission speeds of 311Mbps over 28MHz channels or up to 622Mbps over 50-56MHz channels.

In Provigent's PVG310 single-chip modem, the implementation of differential services is enabled by the on-the-fly adaptive code and modulation (ACM) mechanism, which allows a dynamic link capacity as modulation and coding change on a frame-by-frame basis according to the link conditions, thus allowing more efficient use of the spectrum.

The company believes that with the proliferation of 2.5G and 3G networks, cellular operators are migrating from traditional exclusive voice services to a rich mix of time division multiplexing (TDM) and IP traffic for voice, data and video services. Thus, there is a steadily growing demand for differential services over cellular backhaul links, as new services are added.

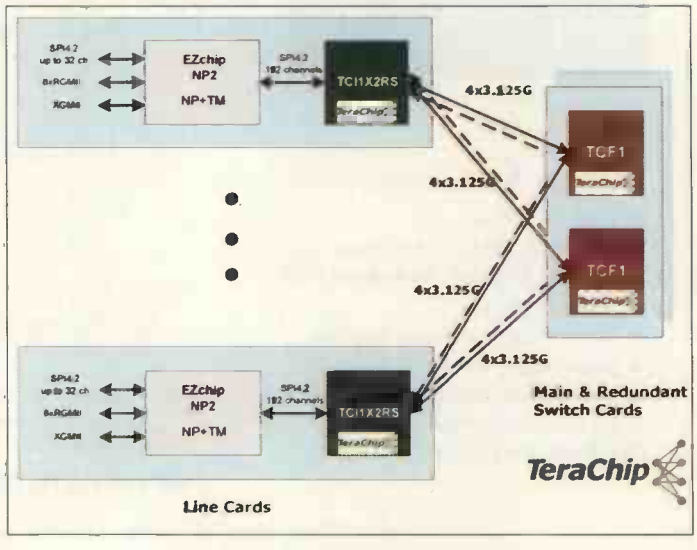

Diagram of the EZchip NP-2 network processor

Provigent's CEO Dan Charash said: "Using our PVG310 with ProviBand, vendors can design adaptive rate systems that meet the requirements of mixed (TDM and IP) high-capacity traffic in next-generation cellular backhaul networks. The ProviBand technology enables cellular operators to support transmission of differential services over multiple classes of availability (CoA); each CoA defines the priority level and capacity per service."

And last but not least of our Israeli crop of high-tech start-ups is Wintegra. The firm develops products to enable the next generation of access networks with its family of access processors. Wintegra recognised early that in access networks the need to integrate new services with legacy systems and multiple existing protocols is critical. It was the first semiconductor company to solely focus on access protocol handling.

The Company's WinPath access processors allow equipment carriers to migrate from legacy networks to IP without jeopardising current revenue or services, while also reducing cost and increasing flexibility. The WinPath chips also enable communications infrastructure equipment vendors to adapt their production lines to the next generation of access networks with a single chip solution.

Wintegra recently announced a joint venture with UK-based picoChip for the joint development of reference designs for a new generation of WiMax base stations. The solutions combine picoChip's picoArray family of massively parallel digital signal processors (DSPs) with Wintegra's WinMax processor and software for a complete MAC solution.

A unique aspect of the joint design is that it is software-defined: both companies offer fully tested software and vendor-independent APIs that permit rapid integration of MAC, PHY and RF components, as well as additional system software. To date, most WiMax reference designs have focused solely on "closed" board-level solutions, providing benefits for those seeking to directly manufacture the design, but provide little flexibility for those seeking to add their own features.

In another partnership, together with Nasdaq-traded Ikanos Communications, a US developer of Fibre Fast broadband solutions, Wintegra has developed a VDSL2 line card reference platform that will speed time to market and reduce costs for access equipment designers.

Intelligent line cards utilise functionality partitioned between Wintegra network processors and Ikanos VDSL2 chipsets, enabling on-board QoS management, traffic shaping, ADSL termination and bonding, without expensive custom engineering. In addition, the tight integration between access processor and VDSL chipset functions offers the potential for higher performance than designs that employ larger numbers of components.

SERIAL COMMUNICATIONS SPECIALISTS
Test and Measurement Solutions

EasySYNC™

featured products

£81.50 £61.00

CAN-USB & CAN-232
affordable CAN Bus to USB and RS232 Interface solutions

=== Bronze Prize Winner ===
NASA Tech Briefs 2004
Products of the Year

£125.00

DS1M12 USB Scope / Logger
2 x 1MS/s Input Channels + waveform generator output. EasyScope & EasyLogger s/w included.

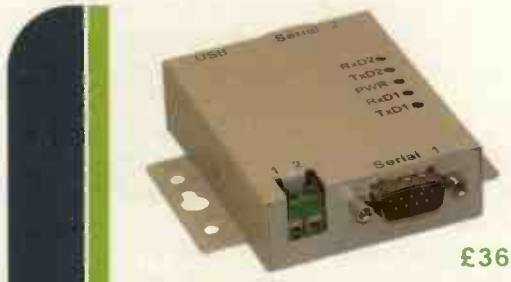

£36.00

USB-2COM-M
2 Port Industrial USB RS232 Serial with wall mount bracket and DC auxiliary output

Low Profile Version also available

£65.00 £10.00

UPCI-400HS
4 Port UPCI RS232 Serial Card with Spider Cable or COMBOX IO (extra)

15

Affordable CAN Bus Solutions
CANUSB is a very small dongle that plugs into any PC USB Port and gives an instant CAN connectivity. This means it can be treated by software as a standard COM Port (serial RS232 port) which eliminates the need for any extra drivers or by installing a direct driver DLL for faster communications and higher CAN bus loads. CAN232 is a very small dongle that plugs into any PC COM Port, or any other RS232 port in an embedded system and gives an instant CAN connectivity. This means it can be treated by software as a standard COM Port (serial RS232 port) which eliminates the need for any extra drivers. Sending and receiving can be done in standard ASCII format.

priced from £61.00 (CAN-232)

USB Instruments - PC Oscilloscopes & Logic Analyzers
Our range of PC Instruments may be budget priced but have a wealth of features normally only found in more expensive instrumentation. Our DS1M12 and PS40M10 oscilloscopes have sophisticated digital triggering including delayed timebase and come with our EasyScope oscilloscope / spectrum analyzer / voltage and frequency display application software and our EasyLogger data logging software. We also provide Windows DLLs and code examples for 3rd party software interfacing to our scopes.
Our ANT8 and ANT16 Logic Analyzers feature 8/16 capture channels of data at a blazing 500MS/S sample rate in a compact enclosure.

priced from £125.00 (DS1M12 & ANT8)

1 to 16 port USB to Serial Adapters
With over 16 different models available, we probably stock the widest range of USB Serial Adapters available anywhere. We offer converter cables, multi-port enclosure style models in metal and plastic, also rack mount units with integral PSU such as the USB-16COM-RM. Serial interfaces supported include RS232, RS422 and RS485. We also supply opto-isolated RS422 and RS485 versions for relaible long distance communications. All our USB Serial products are based on the premium chipsets and drivers from FTDI Chip for superior compatibility, performance and technical support across Windows, MAC-OS, CE and Linux platforms.

priced from £20.00 (US232B/LC)

UPCI Serial Cards
Discover our great value for money range of multi-port UPCI serial cards. Supporting from two to eight ports, the range includes RS232, RS422, RS485 and opto-isolated versions. Our 4 port and 8 port models can connect through external cables or the innovative wall mounting COMBOX.

priced from £21.00 (UPCI - 200L)

EasySync Ltd
373 Scotland Street
Glasgow G5 8QB U.K.
Tel: +44 (141) 418-0181 Fax: +44 (141) 418-0110
Web : http://www.easysync.co.uk
E-Mail: sales@easysync.co.uk

* Prices shown exclude carriage and VAT where applicable

Multi-Threading:

Powering The Next-Generation Of Embedded SoCs

Vivek Sardana of MIPS Technologies describes the workings of the latest microprocessor architecture

Designers everywhere face ever-increasing constraints on system cost and power consumption, while at the same time are required to add more performance and functionality to their designs.

Some previous approaches have been to ramp up the clock speed of a processor, but this usually results in increased power consumption. Additionally, memory performance has not kept pace with processor technology (**Figure 1**), and this mismatch limits any significant gains in system performance. Consequently, the higher frequency approach is leading to diminishing returns.

The multi-core system is a possible solution, but this suffers from a larger die area and higher cost. Any increase in performance comes at a fairly substantial cost in silicon and system power consumption.

Multiple issue processors with two or more execution units offer another option, but they struggle to make best use of hardware resources and, also, have an area penalty. Additionally, the software has to be revised in many cases to make best use of the multiple pipelines.

A multi-threading solution offers a more interesting approach. By running multiple threads within the same execution pipeline, we are able to mask the effect of memory latency by increasing processor utilisation. As one thread stalls, additional threads are instantly fed into the pipeline and executed, resulting in a significant gain in application throughput. This effectively mitigates the inherent inefficiency of a single pipeline when suffering from thread misses caused by cache stalls. Such inefficiencies can account for the loss of over 50% of the processor cycles in a typical single-threaded processor running at high clock speeds.

Figure 1: Processor-memory mismatch bottlenecks system performance [Source:http://www.sun.com/aboutsun/investor/financials/Yen03_03_03.pdf]

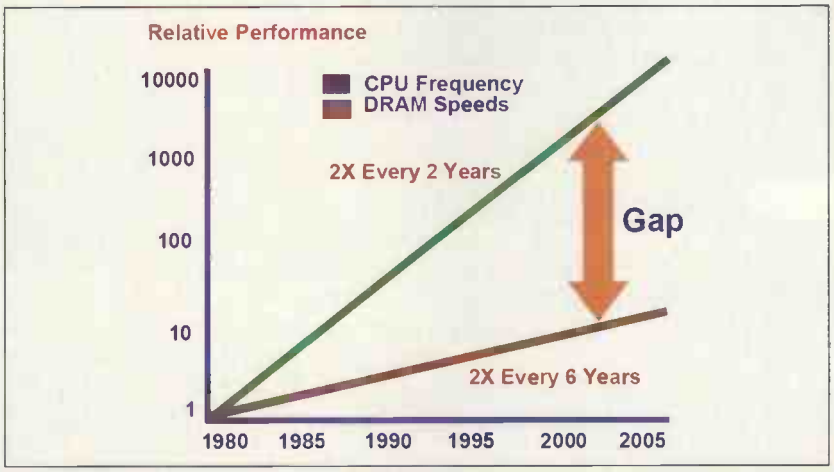

Multi-threading on a single core

Implementing multi-threading capability on a single core, where the design exploits system-level parallelism, is done by processing multiple threads of software simultaneously. This is a novel approach that results in higher system performance, lower system cost and lower power consumption. Such SoCs are targeted at highly concurrent applications such as a set-top-box, VoIP, multifunction printers or digital TV.

This higher efficiency also enables designers to eliminate other hardware blocks in the system, such as a dedicated digital signal processor (DSP). Overall power consumption is decreased because the higher application throughput ensures that the required performance can be achieved at a lower frequency than that of a traditional processor delivering similar performance.

This additional performance is a result of the multi-threaded design, which makes the most out of the full computing power in each clock cycle. The processor resources in the core, therefore, are used more efficiently than in traditional designs, even multi-core designs. The problem is that the execution pipeline of a traditional single-threaded processor will stall for many reasons including cache misses, branch mis-predicts and other pipeline interlocking events (**Figure 2**).

The key to squeezing the maximum performance out of any processor core is controlling the way the threads are executed in the pipeline. **Figure 3** illustrates how the various threads are scheduled into

Figure 2: Pipelines can stall easily, slowing down applications

Figure 3: Improving pipeline efficiency with multiple threads

the execution unit of MIPS's 34K core family for maximum efficiency. When one thread stalls waiting for memory, the other thread is introduced to ensure that the common hardware is as busy as possible.

The 34K technology

Derived from the proven 24KE microarchitecture, the 34K core family uses a nine stage execution pipeline coupled with a small amount of hardware to handle the virtual processing elements (VPEs), the thread contexts (TCs) and the quality of service (QoS) prioritisation.

As illustrated in **Figure 4**, each thread has its own dedicated hardware, called the thread context (TC). This allows each thread to have its own instruction buffer with pre-fetching so that the core can switch between threads on a clock-by-clock basis to keep the pipeline as full as possible. All this avoids the costly overheads of context switching.

Each TC has its own set of general-purpose registers, a PC (program counter) that allows a TC to run a thread from a complex operating system such as Linux. A TC also shares resources with other TCs, particularly the CP0 registers used by the privileged code in an OS kernel.

The set of shared CP0 registers and the TCs affiliated with them make up a VPE (Virtual Processing Element). A VPE running one thread (i.e. with one TC) looks exactly like an independent MIPS CPU, and is fully compliant with the MIPS32 architecture specification.

All threads (in either VPE) share the same caches, so cache coherency is not an issue. This eliminates the problem in multi-core and multi-processor systems, where many cycles and additional logic are used to manage the different processors and ensure cache coherency.

Depending on the application requirements, the 34K core can be configured for up to five TCs that are supported across a maximum of two VPEs. It is this combination of the VPEs with the TCs that provides the most area-efficient and flexible solution.

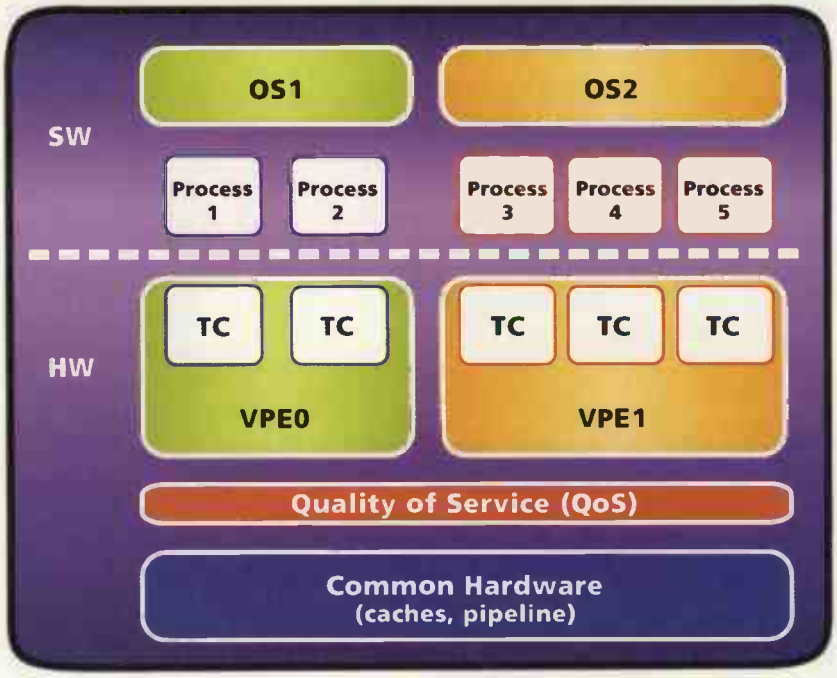

Figure 4: MIPS32 34K processor design

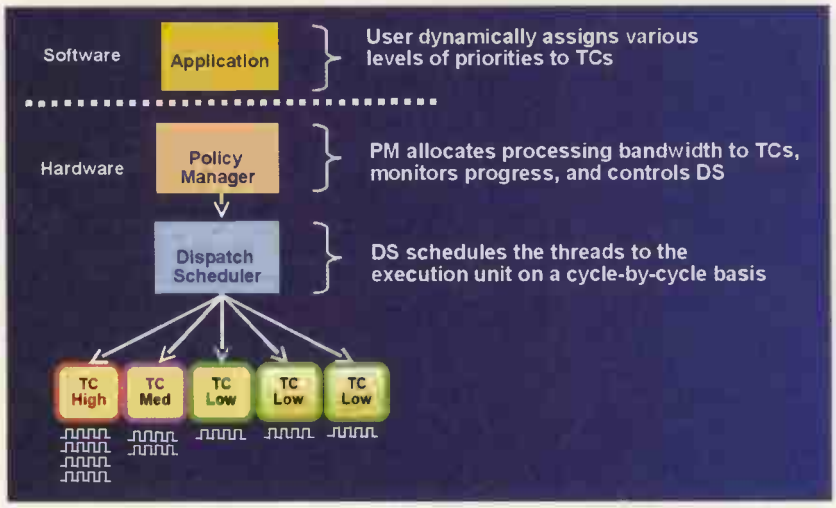

Figure 5: QoS hierarchy

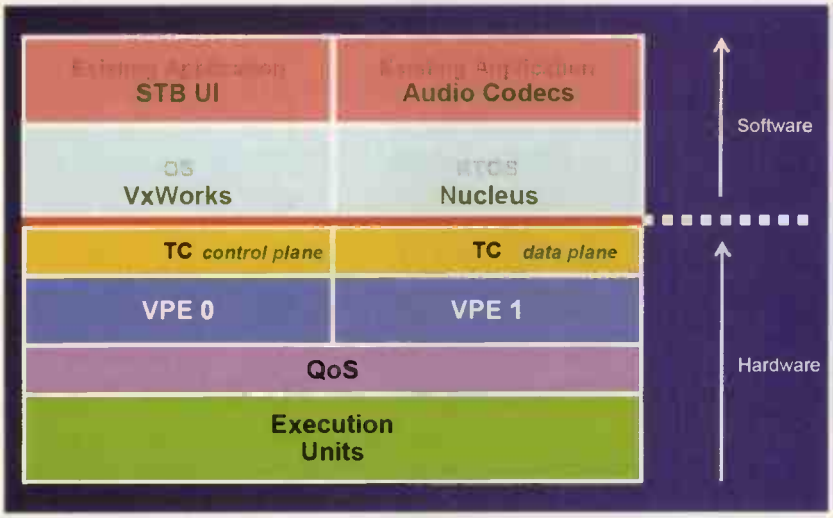

Figure 6: An AMVP (Asymmetric Multiprocessing on VPEs) system running two operating systems side by side

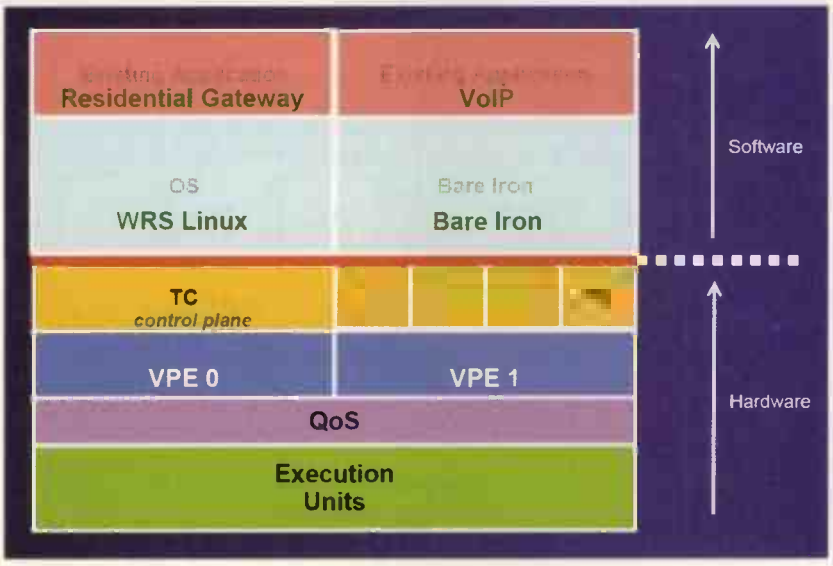

Figure 7: A Residential Gateway supporting multiple VoIP channels

Quality of Service (QoS)

The QoS engine picks instructions from runnable threads in round-robin, interleaving instructions on a cycle-by-cycle basis for maximum overall application throughput. In this manner, the processing bandwidth is shared finely between the different threads, thus using each other's processing "gaps". Alternatively, it can also achieve QoS for real-time tasks such as communications, and video and audio processing, by allocating dedicated processing bandwidth to specific thread(s).

The QoS is handled with a hierarchical approach (**Figure 5**) where the user can program various levels of processing bandwidth to the available TCs. Based on this allocated bandwidth, the integrated Policy Manager assigns priorities to the individual TCs, constantly monitors the progress of the threads, and provides critical "hints" to the Dispatch Scheduler as needed. The Dispatch Scheduler in turn, schedules the threads to the execution unit on a cycle-by-cycle basis, ensuring that the QoS requirements are met.

Software usage models

Another key point is that the majority of these configurations will run existing software without modification, which is a fundamental advantage to the system developer in getting products to market in time.

In its simplest multi-threaded form, the AMVP (Asymmetric Multiprocessing on VPEs) system allows two operating systems to run side by side, each one running its own TC on its own VPE (**Figure 6**). In most cases, these two operating systems would most likely be two legacy operating systems (e.g. Linux for a control plane application and an RTOS for a data plane application) that were previously run on separate MIPS cores.

The 34K is highly configurable core and this is highlighted in **Figure 7**. Customers can significantly lower system cost by running the control plane functions of a Residential Gateway on one VPE with four TCs supporting multiple VoIP channels on the second VPE.

Facing impovements

Although the embedded environment is highly concurrent, the current class of single-threaded processors have not been designed to exploit this concurrency with any reasonable efficiency. The ever-increasing clock frequency does offer incremental benefit in performance, but the memory latencies seriously limit any significant gains in system throughput. The multiple-issue superscalar processors also provide diminishing returns, given the additional cost in silicon and design complexity. Multi-core designs, another alternative, have enjoyed some success in delivering system performance but quite often the associated increases in silicon cost and power consumption have prevented them from being designed-in for numerous embedded applications.

Processing multiple software threads in parallel, masks the effect of memory latency to deliver significant gains in system performance, with a very modest increase in die size. MIPS Technologies's internal benchmarks indicated that the performance improvement of 34K was 60% for an increase in die size of only 14% compared to the equivalent single-threaded processor, the 24KE.

PCB-POOL®

SERVICING YOUR COMPLETE PROTOTYPE NEEDS

1 EUROCARD

(160 x 100 mm)
+ Tooling
+ Photoplots
+ VAT

€49

Price example
Any size and contour possible!

Over 15.000 Customers

Optional:

- Soldermask
- Fast-turnaround
- Silkscreen
- 4-Layer Multilayer
- 6-Layer Multilayer

DOWNLOAD OUR FREE LAYOUT SOFTWARE!

Industry Quality LEAD FREE

Freephone 0800-3898560

Beta LAYOUT
Tel.: +353 (0)61 701170
Fax: +353 (0)61 701165
pcb-pool@beta-layout.com

p-cad TARGET Protel EDWIN orcad a Cadence product family GraphiCode PROTEUS Electronics WORKBENCH Easy-PC Sprint Layout

Simply send your files and order ONLINE:

PCB-POOL.COM

Embedded C++ Key to Smaller And Faster Code

Mike Haden, Engineering Manager at Green Hills Software, here presents an overview of the capabilities of Embedded C++ (EC++)

Embedded system software developers currently using C are faced with an important decision when beginning a new project – should they change their programming language to the increasingly popular C++, Embedded C++ (EC++) or remain with ANSI C?

Embedded system software developers today may desire to adopt the C++ programming language to benefit from features such as classes, templates, exception handling and class inheritance, which have proven invaluable for native application development on desktop computers. However, most embedded system applications do not have tolerance to deal with the overhead and complexity of using C++. Some C++ features can dramatically increase the size of the application object code requiring additional target resources and impacting execution speed.

The object-oriented features of C++ can generally simplify the source code and, therefore, the development process, both by allowing code reuse and by placing the onerous housekeeping functions such as range checking and memory allocation in the class definitions and separate from the main application. Typically, C++ is more readable than standard C but, once compiled, the C++ code size may be a factor of five larger than a C implementation.

Embedded C++ is a true subset of C++

An industry group, led largely by major Japanese CPU manufacturers, developed a true subset of ISO/ANSI C++, called Embedded C++ (EC++). This group, known as the EC++ Technical Committee, set out to retain the object oriented concepts of C++ whilst eliminating those most responsible for boosting memory requirements and reducing efficiency.

EC++ omits several C++ features such as multiple inheritance, virtual base classes, templates, exceptions, runtime type identification, virtual function tables and mutable specifiers. Whilst each of these features is useful in its own right, none is compelling for a sufficiently broad range of embedded applications. The support for some of these features will bloat the generated code, whether or not the features are actually used in an application. For example, exception handling is one of the worst offenders and can adversely affect the deterministic response to external events required in real-time systems. So, eliminating these particular features yields substantial reductions in the size of the compiler code and corresponding improvements in runtime efficiency.

Differentiating C, EC++ and C++

Through a series of examples, from a baseline EC++ implementation through to a full C++ implementation, the features that differentiate C, EC++ and C++ can be illustrated. The code fragment in **Example 1** illustrates some of the key advantages that EC++ offers over the C language. The concept of classes is probably the single most important concept originally in C++ and also supported in EC++.

Classes build on the data structures found in the standard C language.

In addition to allocating memory for a number of variables of mixed types, classes can be used to initialise variables, dynamically allocate additional memory for variables and arrays, perform range checking and many other useful functions. In C programs, these tasks have typically been distributed throughout the main code.

Classes and object definition

In Example 1 is an illustration in the use of classes for an array operation. An embedded application such as a data acquisition system might use such a class to create arrays for storage of data samples. The array class named Array includes two integer members. The first is a pointer named "elements" to members of an array and the second is used to track the array size.

As the main() portion of the code fragment implies, the declaration, Array a(6), can be used to create an array object with the name "a" that contains six elements. The class definition includes several important features such as the constructor code, necessary to create the array and to ensure that the size specified is greater than zero. The constructor is located in the public section of the class definition, which allows code located outside the class definition a window through which it can access the elements and element_cnt class members. Each time an array of type Array is created, the compiler automatically calls the constructor function Array (int n). This function first assigns the value passed in the array declaration to element_cnt, then checks for a valid size and, finally, allocates space in main memory for the array by calling "new".

In this simple example, a bad array size such as zero or a negative number causes the constructor to call a simple "die" function that outputs an error message. An embedded system would typically use a more elaborate scheme to handle runtime errors.

The Array class also demonstrates two other key features of classes in EC++ or C++, function definitions within a class and overloaded operators. First consider the "size()" function which

illustrates the simpler of the two concepts. Because element_cnt is a private member of the class, code outside the class can't directly access the counter. The size() function, however, allows the two "for" loops located in the main() section to indirectly access element_cnt for use as an upper limit of the loop.

Overloaded operators

The class definition also includes an example of overloaded operators. Operator overloading allows the programmer to develop new definitions of standard C/C++ operators such as "=", ">", or "+" that are customised for the type of object defined in a class. For example, a class could be developed that defined an object such as a circle or sphere. A size comparison could be made of two objects that were created using the class definition – objects A and B. The expression A>B or A=B have well understood meanings when A and B are integers, but the compiler could have trouble evaluating the expressions when A and B are spheres of given size or composition. To eliminate ambiguity, EC++ and C++ permit the programmer to include a new definition for such operators that works in a way that is advantageous to the specific application and object type.

The sample code in Example 1 overloads the subscripting operator "[ ]" used to store the index for an array. In this case, the overloaded function gets called each time an indexed array reference occurs – for example, "a[i]=i". Instead of changing the effective meaning of the subscripting operator, the example uses the overloaded operator to automatically detect for out-of-range array indices. Note that should the sample program be executed, the output statement used in the second loop would generate an error on the sixth pass through the loop because a[7] would exceed the valid index test.

It can be seen that classes significantly streamline the mainline code in an EC++ or C++ program. For example, a C program would require explicit data-structure definitions for every array declared while EC++ or C++ handles creation of all similar objects with a single class. Moreover, C programs would require memory allocation, error checking and element-count code in the main part of the program or in dedicated C functions. The compiled code overhead of EC++ relative to standard C code is minimal as well. Adding classes and overloaded operators only marginally increases the generated code size.

An important note here is that these code savings cannot be realised simply by not using the memory-hungry C++ features in an application and then compiling the code with a standard C++ compiler. See the highlighted section "Why Use An EC++ Compiler?" details the consequences of attempting this approach.

Using C++ functions omitted from EC++

Having looked at the advantages offered by EC++, some omitted C++ functions prove to be extremely desirable for given applications. To this end, compiler vendors can provide programmers with some flexibility as to what C++ functions are available for use in each application. For example, a programmer can use a compiler directive with the Green Hills C++/EC++ compiler to strictly limit the source code to the EC++ subset and fully realise the savings in code size and the boost in efficiency. Additionally, a programmer can use compiler switches to add support for one or a few specific C++ functions that were left out of EC++. The granular support for optional C++ features allows the programmer to trade off compiled code size with ease of development and maintainability. Furthermore, the Green Hills C++/EC++ compiler allows the programmer to choose libraries appropriate to their application thereby eliminating a significant amount of redundant library code.

Adding template support

Templates provide an excellent example of a C++ feature not included in EC++ that provides significant advantages in development with only a modest increase in memory requirements when used carefully. **Example 2** illustrates the benefit of templates. In Example 1, the Array class was defined in such a way that all members of the array had to be integers. With EC++, additional classes would have to be defined to handle arrays for short, long, char or floating-point data types. Templates allow a single class definition to support creation of arrays for any valid C++ data type.

The only real addition to the Array class definition with a template is the template label that precedes the class definition and the use of the "T" specifier, each time the code addresses an element of the array. Consider the main() code, however, and see that Array works equally well to instantiate array "a1" to store integer data types and array "a2" to store short data types. It would be just as easy to define more arrays to store other data types. An embedded system performing data acquisition, for example, might well require floating-point arrays.

The use of templates as illustrated would result in little or no increase in compiled code size, so programmers that don't need to strictly meet the EC++ specification can leverage a valuable tool. Caution: the code size realised will depend specifically on the embedded application. Much of the code bloat found in C++ code comes not from using a feature such as a template but from referencing templates that are found in large C++ libraries. Reference one of these standard templates and the resultant code will include many aspects of the library that are not required.

Experienced C++ programmers will enable a number of C++ features in an EC++ environment and not decrease efficiency. Such an environment might be called extended EC++. The EC++ development effort eliminated features such as templates, name spaces, mutable specifiers and new-style casts more so due to the complexity of using the features properly than due to inherent inefficiency. C programmers may find it difficult to use the features correctly, but careful, experienced C++ programmers can leverage the benefits of these features without penalty. Moreover, compilers such as the Green Hills C++/EC++ compiler makes extending EC++ as simple as using compiler switches.

Full C++ with exception handling

What other features become available in moving towards full C++? Exception handling proves to be among the most valuable of features to embedded system designers, yet is also among the leading in causing compiled code bloat. Exception handling

provides a systematic approach to trapping errors caused by operator input or even out-of-range errors in a data-acquisition environment.

The code fragment in **Example 3** illustrates C++ exception handling. This example is more typical of the kind of error handling required in complex embedded systems than was the simple die() function used in the first two examples.

C++ defines the keywords "try", "throw" and "catch" for use in exception handling. Typically, programmers organise code within blocks called try blocks that are enclosed within braces, as in { }.

A second block of code called the catch block is dedicated as a centralised runtime exception/error dispatch service. Anywhere within the try block, a throw directive can originate an exception condition and transfer control to the catch block based on evaluation of a C++ "if" statement. The throw statements can be in-line within the try block or located in functions within the class definition.

Example 3 uses the throw mechanism at two different places in the Array class definition. The first throws to the catch block when an Array instantiation has zero or a negative number of elements. This throw is located in the constructor function. The second throw is used to handle array index errors detected by the overloaded array subscripting operator.

C++ offers significant flexibility in how exceptions are handled and in all cases allows separation of the exception handling code from the mainline application. As in the example, it is possible to dedicate a catch block to each try block. Alternatively, define a single catch block to service an entire main() program. The code in a catch block only executes when a throw evaluation fails. For example, should the code within the first try block in Example 3 execute with no exception the following catch block will be skipped and the second try block will execute.

Finding the right mix for a specific application

EC++ promises to provide embedded-system programmers a valuable path to leverage the most significant aspects of an object-oriented language. With formal approval of the EC++ standard imminent, programmers should demand C++ compilers that include EC++ support. To minimise development time and simplify code maintenance, however, programmers should not dismiss all of the C++ features that were eliminated in EC++. By carefully choosing which features to use on an application-by-application basis, programmers can both simplify development and maintain reasonable runtime efficiency. Moreover, complex embedded systems easily benefit from the robust libraries and features of C++.

Why Use An EC++ Compiler?

Since EC++ is a proper subset of ISO/ANSI C++, one might assume that the efficiencies promised by EC++ could be delivered by simply avoiding certain C++ features. EC++ code can certainly be compiled on a C++ compiler but the resulting code will still require five or six times more memory than the same code compiled with an EC++ compiler.

There are three main problems that arise from attempting to increase efficiency using EC++ code and a standard C++ compiler.

First problem:

The compiler will still use C++ libraries and link a multitude of code into the finished product that is surplus to requirements for your application. An EC++ compiler, conversely, uses libraries that are optimised for the new dialect and thereby generates smaller more efficient code.

Second problem:

The compiler is not as efficient in code optimisation. EC++ compilers can achieve superior optimisation relative to C++ compilers. EC++ compilers can optimise code without presuming the possibility that complex features such as exception handling may be used at some point.

Third problem:

Standard C++ compilers have no mechanism to enforce EC++ compliance within a programming team. With a standard C++ compiler, one out of 10 or 20 programmers on a team can use an offending feature and destroy the efficiency of the entire code base.

Green Hills developed a sample EC++ program to demonstrate memory efficiency. The program solves a form of the classical "travelling salesman" (Travelling Salesman Program – http://www.ghs.com/wp/citlesdemo.html) problem often used to teach programming. When compiled on the Green Hills C++/EC++ compiler using EC++ mode the total code size is 57kB. When compiled using the identical EC++ source file but in C++ mode with no exception handling library, the code size is 322kB.

Adding an exception handling library brings the code size to 378kB.

Example 1

```
//
// Embedded C++ example of a simple array class that does range checking
// on creation of an array, and on array subscripting operations.
//

#include <iostream>

extern "C" void exit(int);

// Deal with a runtime error.  A real embedded application would
// probably choose a different error handling strategy.

void die(const char *msg, int n)
{
   cout << msg << n << endl;
   exit(1);
}

// The integer array class

class Array {
private:
   int *elements;                    // array elements
   int element_cnt;                        // array size
public:
   Array(int n) : element_cnt(n) {    // construct a new array
         if (n > 0)
            elements = new int[element_cnt];
         else
            die("Bad Array size ", element_cnt);
   }
   int &operator [](int indx) const { // overloaded subscripting operator
         if (indx < 0 || indx >= element_cnt)
           die("Bad Array index ", indx);
         return elements[indx];
   }
   int size() { return element_cnt; } // return the size of the array
};

main()
{
   Array a(6);
   for (int i=0; i<a.size(); i++)
         a[i] = i;
   for (int i=0; i<a.size()+1; i++)      // error on a[7]
         cout << i << ". " << a[i] << endl;
}
```

Example 2

```
//
// Extended embedded C++ example of a simple array class that does range
// checking on creation on an array, and on array subscripting operations.
//
// This time we use a template class for the array class.
//

#include <iostream>
```

```
extern void die(const char *, int);
extern "C" void exit(int);

// Deal with a runtime error.  A real embedded application would
// probably choose a different error handling strategy.

void die(const char *msg, int n)
{
   cout << msg << n << endl;
   exit(1);
}

// The array class using templates

template <class T>
class Array {
private:
   T *elements;                    // array elements
   int element_cnt;                        // array size
public:
   Array(int n) : element_cnt(n) {    // construct a new array
         if (n > 0)
            elements = new T[element_cnt];
         else
            die("Bad Array size ", element_cnt);
   }
   T &operator [](int indx) const {   // overloaded subscripting operator
         if (indx < 0 || indx >= element_cnt)
           die("Bad Array index ", indx);
         return elements[indx];
   }
   int size() { return element_cnt; } // return the size of the array
};

main()
{
   Array<int> a1(6);
   for (int i=0; i<a1.size(); i++)
         a1[i] = i;
   for (int i=0; i<a1.size()+1; i++)    // error on a[7]
         cout << i << ". " << a1[i] << endl;

   Array<short> a2(6);
   for (int i=0; i<a2.size(); i++)
         a2[i] = i;
   for (int i=0; i<a2.size(); i++)
         cout << i << ". " << a2[i] << endl;
}
```

Example 3

```
//
// Embedded C++ example of a simple array class that does range checking
// on creation on an array, and on array subscripting operations.
//
// This time we use a template class for the array class and
// C++ exception handling to deal with runtime errors.
//
```

```
#include <iostream>

// The array class using templates

template <class T>
class Array {
private:
   T *elements;                    // array elements
   int element_cnt;                        // array size
public:
   class Range {                   // A nested class to deal with
   public:                                 // runtime errors
         int indx;
         char *msg;
         Range(char *m, int i) : msg(m), indx(i) {}
   };
   Array(int n) : element_cnt(n) {    // construct a new array
         if (n > 0)
            elements = new T[element_cnt];
         else
            throw Range("Bad Array size ", element_cnt); // runtime error
   }
   T &operator [](int indx) const {
         if (indx < 0 || indx >= element_cnt)
           throw Range("Bad array index: ", indx);      // runtime error
         return elements[indx];
   }
   int size() { return element_cnt; } // return the size of the array
};

main()
{

   // arrange to catch runtime errors related to the Array<int> class

   try {
         Array<int> a1(6);
         for (int i=0; i<a1.size(); i++)
            a1[i] = i;
         for (int i=0; i<a1.size(); i++)
            cout << i << ". " << a1[i] << endl;
   }
   catch (Array<int>::Range rng) { // deal with runtime errors here
         cerr << rng.msg << rng.indx << endl;
   }

   // arrange to catch runtime errors related to the Array<short> class

   try {
         Array<short> a2(-1);
         for (int i=0; i<a2.size(); i++)
            a2[i] = i;
         for (int i=0; i<a2.size(); i++)
            cout << i << ". " << a2[i] << endl;
   }
   catch (Array<double>::Range rng) {      // deal with runtime errors here
         cerr << rng.msg << rng.indx << endl;
   }
}
```

Maintaining Real-Time Performance In Highly Integrated Processors

Kevin Pope, General Manager at Quadros Systems Europe explains how control and data plane developers can use a common infrastructure while staying clearly within their own comfort zones of development environments that meet the needs of their very different types of application

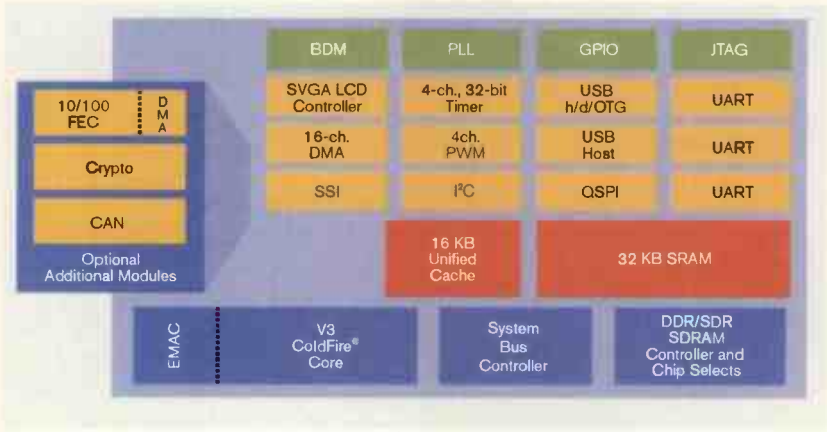

Figure 1: Example of a highly integrated convergent processor

Leading designers of advanced processor architectures, such as Freescale Semiconductor, ARM and Analog Devices (ADI), are constantly driven by market demands for more cost-effective silicon solutions, leading to ever-increasing levels of on-chip integration and complexity.

What is less obvious is the simultaneous trend for such processor designs to include more than one type of processing model. This is often driven by a desire to improve power consumption and silicon cost by merging the capabilities of a previously two-chip solution into a single device. An increasingly common approach is to augment an existing RISC/CISC control plane processing model with DSP (data plane) capabilities (or vice versa). Examples include Freescale's ColdFire microcontroller design, where the firm has enhanced the basic RISC processor model with DSP-like capabilities, and ADI's Blackfin, where they have developed a processor combining a full set of both RISC and DSP instructions in a single core.

However, such processor architectures, often referred to as "Convergent" designs, place new challenges on the embedded software developers to deal effectively with what now may be the conflicting demands of both data-plane and control-plane code executing on a single core, particularly in devices which also have a high level of I/O integration on-chip.

Addressing challenges

The markets general adoption of prioritised pre-emptive Real-Time Operating Systems (RTOS) generally addresses the challenge of highly integrated I/O in the control plane. However, the inclusion of DSP capabilities brings in a new integration challenge: How to integrate data plane code onto a single processor also running control plane code at the same time. In other words, how to support applications that were previously run on two separate processors, each dedicated to providing the specific execution environment for data-plane and control-plane code, now converged onto a single device.

In fact, to avoid losing processor efficiency when running such code together, these highly integrated convergent devices require a properly designed RTOS.

The specific problem with a conventional control-plane RTOS in a convergent processor is that the high priority data-flow processing tasks suffer from the same overheads as the lower priority control tasks. Every time an interrupt occurs, there is likely to be a context switch to a high priority task. The context switch is an operation performed by the RTOS to save the current task's registers and restore those of the high priority task that runs as a result of the interrupt.

Unfortunately, on a processor with additional DSP capabilities, such a context may be even larger and, thus, take still more time to save and restore. Between interrupts and their associated context switching, a lot of processor cycles are spent not on processing data but on housekeeping. The overhead of interrupts applies to all tasks and will almost certainly cause a burdensome load on the data-plane code that may not be expecting to suffer such a burden; on a dedicated DSP processor tasks

would normally expect to run to completion without interruption, and so have no concern for this issue.

Convergent and I/O intensive microcontrollers

In this article, we will take the architecture of the ColdFire microcontroller as an example of a new type of convergent device, combining both RISC and DSP processor models, coupled with a high level of on-chip I/O functionality.

The control plane processor of the ColdFire is based on a 68000 code-compatible Variable Length Instruction RISC engine. The addition of a MAC (Multiply Accumulate) unit to the processor is centered on the notion of providing a limited set of DSP operations that are being used in embedded code today, while supporting the current multiply instructions in the ColdFire architecture. These DSP functions are implemented in a four-stage execution pipeline that is optimised for 32 x 32 multiply-accumulate operations (referred to as the EMAC). This offers a very effective way to support specific functions such as complex fast Fourier Transform (FFT) algorithms (perhaps in a VoIP application), or servo control (for example in a hard disk controller) running on-chip entirely under software control.

The problem is further exacerbated when you consider that it is now common for there to be more I/O devices integrated onto a processor than there are levels of interrupt priority to differentiate the relative importance of each device or channel.

To understand the full impact of high integration within a microcontroller upon a system balancing data plane and control code, let us review a new but representative device, the ColdFire MCF532x. This device has USB host and USB OTG (On The Go) support, both of which could be channeling data at up to 12Mbps. It has Fast Ethernet Controller (FEC) that could be transmitting and receiving at up to 100Mbps, 3 additional UARTs, probably running up to 38.4kbps each, a Synchronous Serial Interface (SSI) most likely used to stream data to a codec, an I2C running at 100kbps, a Queued Serial Peripheral Interface that could be running up to 20Mbps, not to mention a whole set of timer devices.

Despite supporting most of these devices with buffering and DMA, we have the potential for a highly complex I/O sub-system that will put considerable load on even the most responsive RTOS. In this case, it becomes even more important to process interrupts efficiently so as to minimise unnecessary overhead that could cause data loss for the data plane code.

Partitioning the system

So how would a better approach work? An ideal RTOS would allow the application to be cleanly partitioned between the control application tasks, data plane tasks and Interrupt Service Routines (ISRs). The data plane tasks would be processed at a priority level higher than the lower priority control tasks but with minimised overhead when switching between tasks because it is often not necessary for such tasks to save the entire processor context. See Figure 2 for a diagrammatic representation of such a partitioning structure. Zone 1 in the diagram represents hardware priority levels set in the processor, zone 2 represents the set of highest software RTOS priorities used by lightweight data-flow threads and zone 3 the lower set of RTOS priorities used by the control plane tasks.

In this new model, a new event is still detected by the processor, which invokes an ISR. The ISR must save the registers that will be changed during processing of the interrupt as in any system. However, upon exit from completion of the ISR the scheduler may now change context to a high priority process to work on the data received. These high priority software processes are not part of the ISR but separate entities that run at the lowest hardware priority level, allowing other interrupts to be recognised.

The nature of those high priority processes imposes a minimised overhead on the system. They run only when there is a block of data available, with the consumption or production of each data block being considered a complete cycle of the process. Consequently, these processes have no need to restore a context upon their start and there is none to save upon their completion, a property that results in a significant reduction in CPU cycles spent on housekeeping overhead.

Thus, there is only one context to consider saving and/or restoring – that of the interrupted task. This

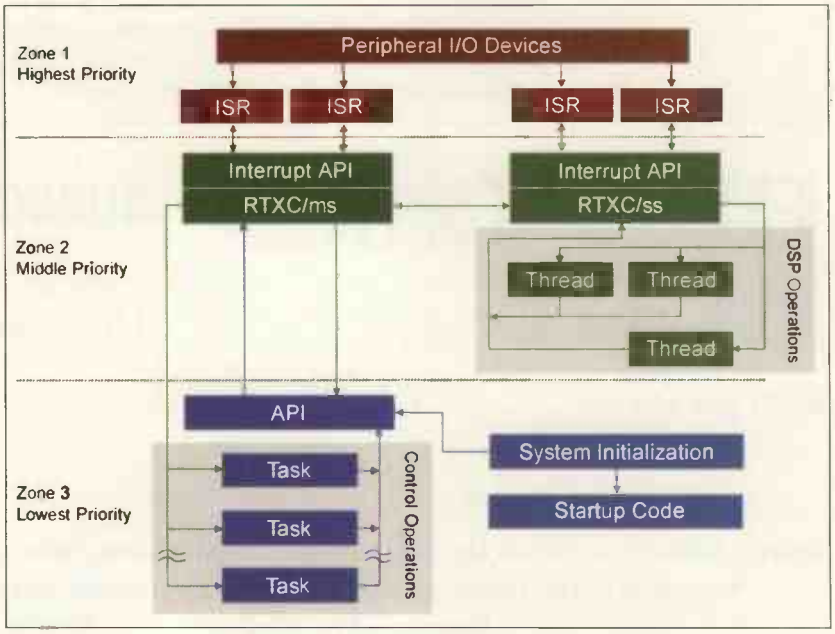

Figure 2: Example of dual mode RTOS partitioning in a system

means that valuable cycles are being saved every time a data plane process is run.

The benefit is compounded if the data-plane code using the EMAC is able to complete several iterations of cascaded processing within several algorithm processes (scheduled in a 'round robin' style at the same high priority level), before finally passing the resultant information to the control application tasks.

Saving cycles in this way not only increases the available processing power, but also reduces the time taken before processing of the captured data begins, resulting in a more responsive system.

Once the processed data reaches the lower priority tasks, the application will behave in much the same way as the traditional RTOS approach. However, this new and improved RTOS model will have prevented several context switches and stack operations and will have delivered the processed data to the control application in a fraction of the time. In interrupt-intensive applications, the saving in CPU cycles used in an entire system for processing EMAC data flows in ColdFire applications, as compared to the conventional RTOS model, has been measured to be as much as 40%. Because fewer processor cycles are required, a processor with a reduced clock rate could be selected, reducing power consumption and possibly a lower cost for the processor too. Alternatively, greater application functionality could be supported on the same device at a given clock speed.

The ideal convergent RTOS

Such "dual mode" RTOS models are already becoming available for ColdFire, Blackfin and other convergent processors. They generally combine a data flow single-stack RTOS model and an event-driven, multi-stack RTOS model, that allows the two processing models to work together efficiently in a single execution environment with a unified API.

The event-driven RTOS follows the traditional model of a priority-based pre-emptive scheduler. Each task has its own stack and each time one task pre-empts another, there is a context switch involving storing and restoring processor state. The data plane RTOS employs a simpler, specialised executive using lightweight tasks often called 'threads' (not to be confused with similarly named objects in POSIX, Linux or other operating systems).

Typically, one uses thread architecture on DSPs or resource-constrained MCUs, where threads share a single, common stack. Once started, a thread is designed to run to completion without the overhead of context saves and restores when changing from one thread to another at the same priority level.

Some RTOS vendors also include a configuration utility so that the ideal balance of code and data resources can be configured into a specific build of the RTOS, thus matching the RTOS with the needs of the applications it is supporting. In the same way that Freescale reduced the instruction set of the ColdFire to essential elements, the software designer can scale the RTOS to a set of essential elements that address the needs of the application.

Communication and synchronisation

Within this cohesive model, the three types of processing entities, tasks, threads and ISRs, all communicate and synchronise through RTOS objects and share a common API to the RTOS services. Thus, control and data plane developers can communicate using a common infrastructure while staying clearly within their own comfort zones of development environments that meet the unique and individual needs of their very different types of application and, most importantly of all, each can deliver high performance in a manner consistent with their slightly different views and definitions of performance.


PCBTrain
PTH PCBs with solder resist and legend from just £30.00 lot price.
Deliveries from 24 Hrs
Visit WWW.pcbtrain.com for full details.
▲ NEWBURY ELECTRONICS, the home of PCBTrain, the UK market leader for low-cost prototype PCBs
▲ Exclusive UK representatives for Circuit Mission Ltd (China) - manufacturers of low-cost high quality rigid PCBs
▲ Exclusive UK representatives for Sunflex Ltd (Taiwan & China) - manufacturers of flexible circuits
CML
SUNFLEX
+44 (0) 1635 40347
NEWBURY ELECTRONICS LTD
Faraday Road Newbury Berkshire RG14 2AD UK

HANDYSCOPE HS4

Four channels not enough

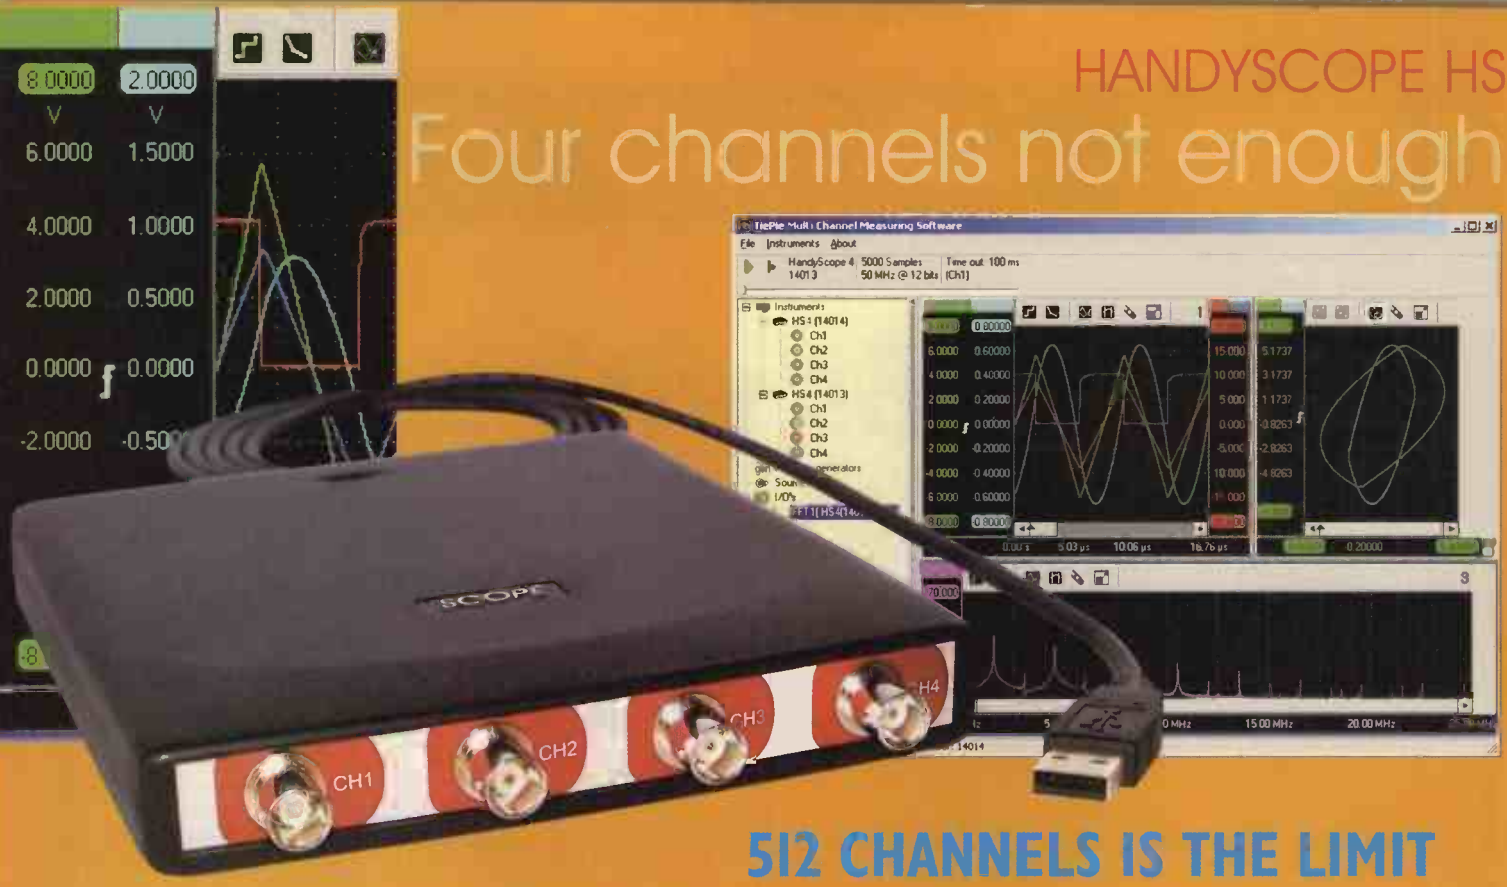

512 CHANNELS IS THE LIMIT

COMPLETE PACKAGE STARTING AT £ 440

The Handyscope HS4 (50MHs 12/14/16 bit) is a powerful and versatile four channel measuring instrument with extension

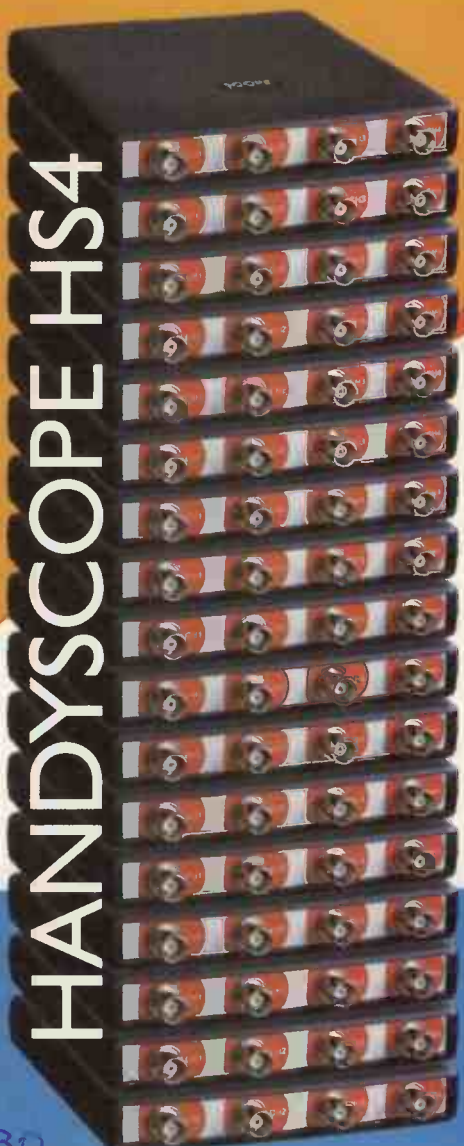

The Handyscope HS4 starts a new standard for multi channel measuring.

It offers perfect measure qualities and through the USB connection it is easy to connect to every PC. Because of the very versatile software it becomes simple to extend the instrument to 512 channels.

A four channel, 12-16 bit oscilloscope, spectrum analyzer, transient recorder and voltmeter created as a most compact instrument.

- Making virtual instruments (combine maximum of 128 instruments)
- Extension to maximum of 512 channels.
- USB 2.0 connection (USB 1.1 compatible)
- Sample speed up to 50 MHz per channel
- 12 to 16 bit resolution (6 μVolt resolution)
- 25 MHz bandwidth
- Input sensitivity from 200 mVolt up to 80 Volt
- Large memory up to 131060 samples per channel
- Four integrated measuring devices
- Spectrum analyzer with a dynamic range of 95 dB
- Fast transient recorder up to 100 kHz
- Several trigger features
- Auto start/stop triggering
- Auto disk function up to 1000 files
- Auto setup for amplitude axis and time base
- Auto trigger level and hysteresis setting
- Cursor measurements with setup read-outs
- Multi window signal display
- Multi channel display

WWW.TIEPIE.NL

for more information, demo software, software, source code and DLL's visit our internet page: http://www.tiepie.nl

TiePie engineering (NL)
Koperslagersstraat 37
8601 WL SNEEK
The Netherlands
Tel: +31 515 415 416
Fax: +31 515 418 819

TiePie engineering (UK)
28, Stephenson Road, St. Ives
Cambridgeshire, PE17 3WJ, UK
Tel: 01480 - 300695
Fax: 01480 - 461654

Data Centric Design for Networked Applications

In this article **Gordon A. Hunt**, principal engineer at Real-Time Innovations presents a data-oriented approach that enables seamless integration of different communication and data storage models in a real-time system

Today's embedded systems are becoming increasingly complex. Applications are becoming more distributed and individual systems (nodes) are becoming more heterogeneous. Additional complexity is added with real-time and dynamically-changing data requirements. Just to make the problem even more interesting, systems are required to enable seamless access to the data they contain through a variety of methods. Low-level messaging, publish/subscribe, data storage and SQL (Structured Query Language), and web service technologies are expected to be fully integrated, scalable and upgradeable in today's distributed applications.

By moving from a message-centric point-to-point solutions, which tend to be operating system specific and/or proprietary implementation oriented to standards-based data-centric technologies, we can develop systems that are inherently more robust, maintainable and upgradeable to meet changing customer and market requirements. Such a data-oriented approach decouples the system implementation in time, space and function, which significantly simplifies lifecycle development of the distributed system.

Recognising that it's the data that is critical in your system, by defining the data and its transient states you can completely define your system. You can then enable your application developers to use the development tools most familiar to them. Such an approach simplifies integration between nodes and addresses the issues of running on and connecting between heterogeneous real-time nodes and back-end Unix systems. For real-time access to data you can use standards-based publish/subscribe peer-to-peer technologies that facilitate high-speed deterministic connectivity, while your back-end system developers, which are more familiar with the enterprise space, can use SQL for their data processing needs and can access localised data or data on the real-time nodes. The following example will illustrate these concepts and ideas.

Example

The following diagram depicts a typical distributed system problem we are trying to solve, using this data-oriented approach. The goal in this example is to maintain the temperature in many buildings, using embedded controllers each hooked to a number of sensors. Each of these sensors and control processes are connected through a transport mechanism such as Ethernet, shared memory, or bus backplane technologies.

Basic protocols such as TCP-UDP/IP or higher-level protocols such as HTTP can be used to provide standardised communication paths between each of the nodes. To achieve data integrity and fail-over capabilities, multiple controllers and sensors can be deployed in each building. Additionally, depending on the size of the building, multiple controllers, each with appropriate backups could be distributed for the different zones. Controllers within a building need to collaborate and all data collected from the various sensors is stored real-time in web-accessible databases. With the inclusion of these distributed databases, we are providing a standards-based way for external applications to obtain, process and manipulate real-time sensor data without having to know the specifics of the real-time data infrastructure. The external access and monitoring applications can simply receive real-time updates from any sensor as well as issue commands to the various controllers via SQL, ensuring that optimal temperature is maintained. This simply stated example is surprisingly complex, containing many elements of real-time messaging, data integrity and failover capabilities, integration with databases, web services, as well as scalability and modularity concerns.

Data model

In order to simplify this example, we will only focus on the data the sensors send to their controller and how it can be distributed throughout the entire system. The first step in a data-centric approach is to carefully describe the data format in a standards-based way, either IDL or XML, and give it a "Topic" name. Topics are the element of the Data Distribution Service (DDS) middleware publish-subscribe standard (see sidebar) which identify the data objects and provide the basic connection between publishers and subscribers. Subscribers, in this case the Controllers, register Topics with the middleware they wish to receive. Publishers, the individual sensors in this example, register topics with the middleware they will send. If Topics do not match, communication will not take place.

Topics enable one to find specific information sources and sinks when architecting a loosely coupled system. A loosely coupled system is one in which you do not know a priori how many sensors or controllers there are going to be or where they all are. The controller can simply subscribe to "TempSensor", the Topic's name, and receive all the sensor updates for that building. Similarly, a sensor does not need to know if it is sending its data to one or multiple controllers.

Specification of the Topic's name is a key element in a data-centric approach to creating open real-time systems. One could name each sensor's topic based on its unique location in the building, "Floor12Room3Sensor14" for example, but the controller would then need to be configured every time a sensor is added or removed from the system. Topics (name and type) define the standard interface for the distributed system and should be chosen appropriately.

Data type

Specification of the Topic's data type is equally important as the Topic's name. For this example we are using Interface Definition Language (IDL) because it is an open standard and readily maps to XML and SQL semantics.

In the definition of the Topic's type, one or more data elements can be chosen to be a "Key". Keys provide scalability and the communication infrastructure can use the key to sort and order data from many sensors. In this example, without Keys, one would need to create individual Topics for each sensor. Topic names for these topics might be: Sensor_1, Sensor_2, and so on. Therefore, even though each Topic is comprised of the same data type, there would still be multiple Topics. With keys, there is only one topic, "TempSensor", used to report temperatures.

New sensors can be added without creating a new Topic. The publishing application would just need to set a new ID when it was ready to publish. An application can also have a situation where there are multiple publishers of the same Topic with the same Key defined. This enables the application to provide redundancy. Using our example, we can put two sensors in the same room, giving them the same Key value states so they are measuring the same piece of information. Managing the redundancy, should one or both sensors report to the controller, is accomplished though Quality-of-Service (QoS).

Data-centric QoS

Data-centric communication using DDS provides the ability to specify various parameters like the rate of publication, rate of subscription, how long the data is valid and many others. These QoS parameters allow system designers to construct a distributed application based on the requirements for, and availability of, each specific piece of data. A data-centric environment allows you to have a communication mechanism that is custom tailored to your distributed application's specific requirements yet remains a loosely coupled design and architecture.

The ability to set QoS on a per-entity basis is a significant capability provided by DDS. Being able to specify different QoS parameters for each individual Topic, Publisher or Subscriber, gives developers many options when designing their system. Through the combination of these parameters, a system architect can construct a distributed application to address an entire range of requirements, from simple communication patterns to complex data interactions.

The following briefly details how one might leverage a few of the QoS in DDS for this example.

➢ **Domain –** A Domain is the basic DDS construct used to bind individual publications and subscriptions together for communication. A distributed application can select to use single or multiple domains for its data-centric communications. In the example, different buildings map to different Domains. Domains isolate communication, promote scalability and segregate different classifications of data.

➢ **Partition** – The Partition QoS is a way to logically separate Topics within a Domain. The value is a string. If Subscriber sets this string, then it will only receive messages from Publishers that have set the same string. In the context of our example, Partitions can be used to group sensors on different floors. For example, we want to divide the building into different zones, where each zone is controlled by a dedicated controller, the sensor and controller could set the Partition to "Floor 1" and "Floor 1-6" respectively. Here, the controller will receive data from all sensors on floors 1 through 6. So, using Partitions make it easy to group the sensors that are 'hooked' to a controller. A controller can take over a different zone by changing or adding to its Partition list.

➢ **Ownership –** The Ownership QoS specifies whether or not multiple publishers can update the same data object and also how you achieve fault-tolerance using DDS.

Returning to our example, if we have multiple sensors in the same room and we only want to get data from the primary (as long as it is functioning) then the Ownership QoS policy is set to Exclusive, stating that only one sensor can update that keyed value. Setting the Ownership policy to Shared is stating that we can have multiple sensors in the same room all reporting the same piece of keyed data. In this case the controller would get all updates from all sensors and treat the values as the same measurement.

➢ **Durability –** The Durability QoS specifies whether past samples of data will be available to newly joining subscribers. Considering our example, if a controller were to reboot, rather than require all sensors to resend their data, or require the data to be sent at a periodic rate in case the system reboots, one simply gets the latest published value for every attached sensor. This effectively decouples the system in time and provides a high degree of data integrity.

➢ **History** – History specifies how many data samples will be stored for later delivery. Specifically, a rebooted controller may want the last five samples from its sensors, so that it can make sure that readings are consistent.

➢ **Reliablity –** Finally, the Reliability QoS may be set on a per Topic basis and informs the middleware that the Subscription should receive all data (no missed samples) from a Publication even over non-reliable transports. Generally, for periodic publications, Reliability doesn't need to be set since you can just get the updated value one sample period later. Although periodic sensor data doesn't need to be delivered reliably, synchronisation commands between Controllers in this example could be.

Integration with databases

The final element of our example system is the integration of real-time data and traditional relational databases. Since both these technologies are data-centric and complementary, they can be combined to enable a new class of applications. In particular, DDS can be used to enable a truly decentralised data structure for distributed database management system (DBMS),

while DBMS technology can be used to provide persistence for real-time DDS data.

Working with the example, each building can maintain the history of the various building sensors in a locally maintained database. The application that manages all buildings would have its own automatically maintained database of the specific data stores (tables) from each building that the application needed. Information is pushed to where it is needed, not senselessly replicated throughout the distributed system.

IDL data models can be automatically and cleanly mapped to SQL table schemas. For example, the Topic "TempSensor" becomes a table named "TempSensor" and the data contents, identified by the Key, become rows in the table.

Essentially, the database is simply another subscription to the sensors' update and automatically receives current data from all the distributed sensors. Changes to the database are pushed to entities that are interested in that particular topic/table name. Embedded applications don't need to know SQL or OBDC semantics, and the database applications don't need to know publish/subscribe semantics. This is a critical point when building large systems: get the data to where it needs to go in a format that is native to the developers.

For true data integrity and scalability, databases should be distributed as well.

RTI's SkyBoard implements a distributed shared database, where fragments of the shared database are kept in the data caches of the hosts in the network on an as-needed basis. Thus, the database becomes a combination of the data stores distributed throughout the system. When a node updates a table by executing a SQL INSERT, UPDATE or DELETE statement on the data cache, the update is proactively pushed to other hosts that access this table via real-time publish-and-subscribe messaging, enabling real-time replication and synchronisation of any number of remote data stores.

Finally, once data is automatically entered and maintained in a DBMS, using standard tools, one can build a web application that accesses and manipulates the database data. Thus, the web application does not need to know how many buildings, sensors or controllers there are in the system. Nor does the web application need to know the middleware specifics that the temperature control system is using to distribute data. The application can just use SQL and ODBC to read and change all of the available real-time data in the system decoupling implementation specifics across the system.

Summary

By starting with the data model and designing the systems following a data-centric approach, we demonstrated building a system that seamlessly integrates a variety of different communication trends and data store trends (database and embedded local data types), while still achieving a high degree of data integrity.

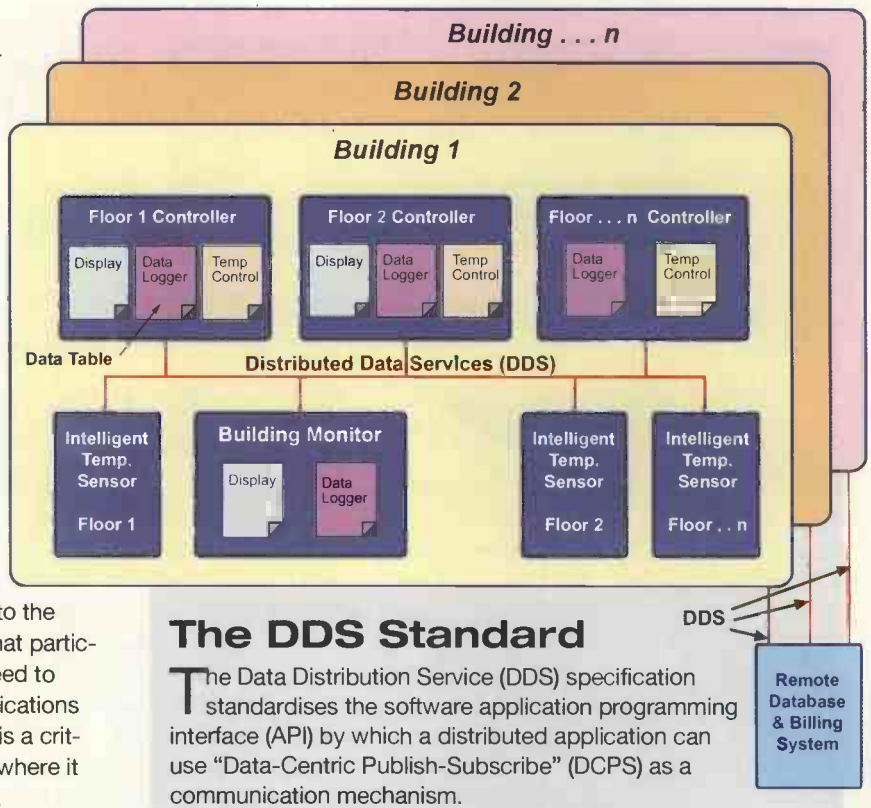

The DDS Standard

The Data Distribution Service (DDS) specification standardises the software application programming interface (API) by which a distributed application can use "Data-Centric Publish-Subscribe" (DCPS) as a communication mechanism.

The DDS standard has three main goals:

1. To define a model for communication as pure data-centric exchanges, where applications publish (supply or stream) data, which is then available to remote applications that are interested in it.

2. To provide a mechanism of specifying the available resources and providing policies that allow the middleware to align the resources to the most critical requirements, giving system designers the ability to control Quality of Service (QoS) properties that affect predictability, overhead and resource utilisation.

3. To permit systems to scale to hundreds or thousands of publishers and subscribers in a robust manner.

Since DDS is implemented as an "infrastructure" solution, it can be added as the communication interface for any software application.

Advantages of DDS:

- Based on a simple "publish-subscribe" communication trend
- Flexible and adaptable architecture that supports "auto-discovery" of new or stale endpoint applications
- Low overhead – can be used with high-performance systems
- Deterministic data delivery
- Dynamically scalable, efficient use of transport bandwidth
- Supports one-to-one, one-to-many, many-to-one and many-to-many communications
- Large number of configuration parameters that give developers complete control of each message in the system.

DDS provides an infrastructure layer that enables many different types of applications to communicate with each other. The DDS specification is governed by the Object Management Group (OMG), which is the same organisation that governs the specifications for CORBA, UML and many other standards. A copy of the DDS specification can be obtained from the OMG website at www.omg.org

Embedding Mobile Audio

Jeremy Copp investigates the standard audio capabilities in today's handsets and the requirements for embedding audio capabilities into chipsets

Since the success of monophonic ringtones in the late 1990s, the audio capabilities of the mobile phone have improved on yearly. Consumers now expect rich polyphonic audio as a standard feature on their handsets, in much the same way as they demand a colour screen or built-in camera.

This article will investigate the standard audio capabilities in today's handsets and the factors driving the recent surge in adoption. It will also consider the manufacturers' motivations, needs and requirements, in terms of embedding audio capabilities into their handsets, chipsets and reference designs.

Not just a ringtone

Whilst the 'crazy frog' ringtone was driving everybody crazy last year, it also placed the ringtone at the forefront of the collective consciousness. However, this represents only a small component of the advanced audio capabilities of mobile phones. Today, network operators deliver a wide range of additional value added services that rely on audio. Network operators drive such advances as they seek to increase revenues from non-voice services of which audio is a core component.

Nevertheless, the success of the mobile audio experience is highly dependent upon the quality of the sound generated. In contrast to the tinny, single note of the first generation mobile ringtones, today's handsets are able to produce audio playback of ever increasing quality. The ability to create 'natural' musical sounds with a wide variety of instrumentation and effects encourages personalisation, increases use of the phone and creates a more acceptable audio environment. A high fidelity audio experience is also essential if mobile phones are to be used for music playback applications.

One example is in enhancing the user experience for mobile gaming, an area in which the operators are concentrating a great deal of resources. Mobile gaming has rapidly evolved from simple games to hugely complex, immersive games offering hours of gameplay. Leading games publisher Glu Mobile was one of the first developers to make full use of the high quality audio synthesis within its games. In contrast to the first generation of mobile games that were silent, allowing programmers to access the audio engine from within the Java Virtual Machine (JVM), the user experience can be dramatically enhanced with high quality atmospheric background music and realistic sound effects.

Mobile audio standards

One can clearly see the potential applications for audio as a revenue generator in mobile applications. However, the adherence to industry standards is a key consideration. The wider audio industry has always benefited from adherence to standards. The successes of formats ranging from vinyl to the CD were all based on the industry deciding upon a fixed standard and then sticking to it. The benefits to such a policy extend across the value chain from the consumer, who can buy any music album safe in the knowledge that it will play perfectly on their hi-fi equipment, to the record company, who don't have to stock different versions of their music for Sony or Panasonic stereos, for example.

It is no different in the mobile phone world. Conformity to such standards benefits consumer, operator and content producer alike. The mobile value chain is already painfully aware of the drawbacks of non-standardised handset technologies. The problems are well illustrated through the current situation with the Java J2ME profile, which, whilst nominally a standard, is also highly fragmented. This lack of standardisation means that mobile games developers must create different versions for each and every handset manufacturer and often-different versions for different handset ranges from the same manufacturer.

One such standard is the MIDI format, which has been around for over twenty years and is published

and administered by the MIDI Manufacturers Association (MMA) and its Japanese affiliated organisation, AMEI. Both MMA and AMEI have been active in publishing audio file format standards that are designed for the wireless environment, specifically Scalable Polyphony MIDI (SPMIDI) and eXtensible Music Format (XMF). The technical groups within 3GPP have recently adopted both standards, the SPMIDI and XMF (in a format known as Mobile XMF), as formats for structured audio.

SPMIDI allows audio content authors to produce one file that will play optimally on a range of devices with varying maximum polyphony levels. This means that both the content and service providers have to only manage and deliver one file, rather than having to author and deliver a number of different files, each optimised for different phone models.

XMF is a multimedia container format that allows for custom instrument definitions, custom linear samples and multiple audio files to be delivered in a secure format to the device. This allows audio composers to include special instruments or sound effects in the content and also to sequence linear digital sample with MIDI data. Examples of these could include sound effects or vocal loops. XMF allows CD quality audio to be delivered in a small file size, together with the necessary copyright acknowledgement and rights protection information.

The latest generation mobile phones support both SPMIDI and Mobile XMF formats for ringtones and multimedia messaging audio clips, ensuring both a high fidelity audio experience for the user and an environment for open standards based services to be deployed.

Hardware vs. software implementation

It can be clearly seen that standards such as XMF and SPMIDI benefit the entire mobile value chain. As such, network operators will support and increasingly mandate both standards for their handset ranges.

However, whilst the standards have now been defined, the issue of incorporating support on the handset still exists. The early audio engines on mobile phones were comprised of an additional hardware chipset built into the device. This approach found favour when handsets did not support sufficient processing power to support a software solution. However, as Moore's Law plays its role on the mobile phone, the disadvantages of this approach have been highlighted. These include:

- **Cost:** a dedicated audio chip can add significantly to the bill of materials for a phone, costing several dollars per unit cost.
- **Power:** adding an additional device to the phone increases power consumption and decreases the time that the phone can run on battery charge. This has become an even greater issue as components such as colour screens and more powerful processors are very power hungry.
- **Physical volume:** hardware devices occupy additional space within the phone, making it harder to reduce the size of the device.
- **Flexibility:** the synthesiser has a fixed configuration and audio performance, meaning that manufacturers cannot differentiate within their product range nor against competitors and that the audio performance of the synthesiser cannot be optimised to match the characteristics of the audio output stage.
- **Supply chain:** a single source of devices can potentially represent a risk to a manufacturer's ability to maintain volume production if supplies are limited. This is especially important with mobile phones, which are a very high volume product.

The processor and memory capacities of mobile phone devices today can now support a full software wavetable synthesiser implementation, and many manufacturers are adopting this solution.

A pure software solution provides significant flexibility to the manufacturer for integrating and customising the synthesiser for a wide range of products. Device designers have the freedom to partition the audio engine to match the underlying hardware architecture, having the potential to split the processing tasks between a main microprocessor and coprocessor DSP(s), if required. The capabilities of the synthesiser can also be varied between products in the range, allowing for product differentiation and additional features to be included in high-end models.

A software implementation also allows easy expansion of the audio engine capabilities to support additional audio formats and new requirements without the need for any hardware redesign.

Gone soft

Another major benefit of a pure software approach is that the audio performance of the synthesiser can be fully customised and optimised to match the characteristics of the audio output stage, hence ensuring the highest quality sound output. With a fully configurable wavetable, the manufacturer can choose which instruments to populate the device with, the size and resolution of samples for each instrument definition, and optimise the instrument samples to match the frequency characteristics of the audio amplifier and transducer in the phone. This means that manufacturers can further differentiate their products by having a distinctive sound and instrumentation set. Crucially, the audio optimisation is carried out by the device manufacturer rather than being demanded from the content developers who have to do optimisation for all specific phone models, as is the case in a fixed hard-

ware synthesiser implementation. Indeed, such is the flexibility of a pure software wavetable implementation that the instrumentation set could be produced to exactly replicate the sound of a hardware FM synthesiser, if required.

Licensing software synthesiser technology also provides the manufacturer with a significant cost saving over a hardware solution, with costs per unit in the order of a few cents for high volume devices. This reduces the overall bill of materials for the phone and, as such, potentially increases the manufacturers' margin. A software solution is capable of using existing memory and processor resources (and so not requiring additional devices) and occupies no physical space in the device. It is possible for a manufacturer to invest in a single technology that can be applied across all devices in its product range, allowing any research and development costs to be amortised over all products.

A software implementation also eliminates a dependency upon the supplier to provide product to meet production deadlines and so reduces single sourcing risks (especially if the manufacturer has access or rights to the software source code).

Embedding

Audio software can be embedded directly in chipsets and reference designs. Optimisation of the audio engine can take place at the core of wireless chipsets, for example. The software is embedded into the ROM. This allows handset producers to develop mass-market handsets based on validated hardware and software reference designs, which support the complete range of advanced audio and multimedia standards. They also stand to gain from significant improvements in time to market as well as cost savings in the production of devices.

As such, embedding audio technology in software at the chipset level makes logical sense for today's handset manufacturers. Even though the phones of today have come on leaps and bounds from the first generation models of the late 1990s, they still represent a significantly constrained environment with not only limited processing resources (both the power of the microprocessor and the amount of memory available), but also power usage and physical volume. This means that any method of adding advanced functionality, such as audio, whilst working within physical and commercial considerations is sure to be attractive.

PicoScope 3000 Series PC Oscilloscopes

The PicoScope 3000 series oscilloscopes are the latest offerings from the market leader in PC oscilloscopes combining high bandwidths with large record memories. Using the latest advances in low power electronics, the oscilloscopes draw their power from the USB port of any modern PC, eliminating the need for mains power.

- High performance: 10GS/s sampling rate & 200MHz bandwidth
- 1MB buffer memory
- Advanced display & trigger modes
- Compact & portable
- Supplied with PicoScope (oscilloscope/spectrum analyser) & PicoLog (data acquisition) software.

PicoScope	3204	3205	3206
Bandwidth	50MHz	100MHz	200MHz
Sampling rate (repetitive)	2.5GS/s	5GS/s	10GS/s
Sampling rate (single shot)	50MS/s	100MS/s	200MS/s
Channels	2+Ext trigger	2+Ext trigger/Sig gen	2+Ext trigger/Sig gen
Oscilloscope timebases	5ns/div to 50s/div	2ns/div to 50s/div	1ns/div to 50s/div
Timebase accuracy	50ppm	50ppm	50ppm
Spectrum ranges	0 to 25MHz	0 to 50MHz	0 to 100MHz
Record length	256K	512K	1MB
Resolution accuracy	8 bits / 3%		
Ranges	10mV to 2V/div		
PC Connection/Power supply	USB2.0 (USB1.1 compatible)		

Tel: 01480 396395

www.picotech.com/scope206

pico® Technology Limited

PLEASE ENSURE YOU TELEPHONE TO CHECK AVAILABILITY OF EQUIPMENT BEFORE ORDERING OR CALLING.

SPECIAL OFFERS

OSCILLOSCOPES

TEKTRONIX 2247A 4 Channel 100MHz Counter/Timer/ Voltmeter £275
TEKTRONIX 2335 Dual Trace 100MHz Delay Sweep ... £125
TEKTRONIX 485 Dual Trace 350MHz Delay Sweep ... £300
IWATSU SS5711 4 Channel 100MHz Delay Sweep ... £150
PHILIPS 3065 2+1 Channel 100MHz Dual TB/Delay – Autoset £200
PHILIPS 3055 2+1 Channel 60MHz Dual TB/Delay – Autoset £150
PHILIPS PM3217 Dual Trace 50MHz Delay Sweep ... £125
KIKUSUI COS6100 5 Trace 100MHz Delay £150
TEKTRONIX 475A Dual Trace 250MHz Delay Sweep ... £175
TEKTRONIX 475 Dual Trace 200MHz Delay Sweep ... £150
TEKTRONIX 465B Dual Trace 100MHz Delay Sweep ... £125
TEKTRONIX 465 Dual Trace 100MHz Delay Sweep ... £95
PHILIPS PM3209 Dual Trace 40MHz Delay £125
PHILIPS PM3215 Dual Trace 50MHz £75
KENWOOD CS4035 Dual Trace 40MHz £50
PANASONIC VP5564A Dual Trace 40MHz £50
HITACHI V525 Dual Trace 50MHz Cusors £95
HITACHI V523 Dual Trace 50MHz Delay £80
HITACHI V425 Dual Trace 40MHz Cursors £75
HITACHI V422 Dual Trace 40MHz £60
HITACHI V223 Dual Trace 20MHz Delay £60
HITACHI V222 Dual Trace 20MHz £50
HITACHI V212 Dual Trace 20MHz £50
FARNELL DTV12-14 Dual Trace 12MHz £40

STORAGE

PHILIPS PM3320 Dual Trace 200MHz 250Ms/S £300
LECROY 9400 Dual Trace 125MHz £325
TEKTRONIX 468 Dual Trace 100MHz Delay Sweep Digital Storage £200
VELLEMAN HPS5 1MHz 5MHz Sampling. Handheld Unused £60

ANALYSERS

ADVANTEST R3265A 100Hz-8GHz £4500
TEKTRONIX 492P 50kHz-21GHz £2250
HP8560A 50Hz-2.9GHz Built In Tracking Gen £3250
HP 8560A 50Hz-2.9GHz £2950
HP 8569A 10MHz-22GHz £950
HP 8565A 10MHz-22GHz £750
HP 853A with 8559A 100kHz-21GHz £1100
HP 182T with 8559A 100kHz-21GHz £750
HP182T with 8558B 100kHz-1500MHz £600
HP 182T with 8557A 10kHz-350MHz ..... £300-£400
HP 140T with 8555A 10MHz-18GHz £500
ADVANTEST TR4131 10kHz-3.5GHz £950
WAYNE KERR SSA1000A 150kHz-1GHz £950
MARCONI 2382 200Hz-400MHz High Resolution £1250
MARCONI 2370 30Hz-110MHz £500
HP 8754A Network Analyser 4-1300MHz £500
MARCONI 6500A Amplitude Analyser with head £750
HP 334A Distortion Analyser 5Hz-600kHz £100

SIGNAL GENERATORS

HP 8350B Sweeper with 83592B 10MHz-20GHz £1500
HP 8350A Sweeper with 83592A 10MHz-20GHz £1250
HP 8350B Main Frame Only £125
HP 83525B RF Plug-in for 8350 0.01-8.4GHz £500
HP 83590A RF Plug-in for 8350 2-20GHz £800
HP 8660C Sig Gen 1.3GHz £450
HP 8660C Sig Gen 2.6GHz £750
HP 86603A RF Plug-in for 8660C 1-2600MHz ..... £.....
HP86631B Axillary Section for 8660C ..... £.....
HP86632B Modulation Section for 8660C ..... £.....
MARCONI 2017 0.01-124MHz Low Phase Noise £500
MARCONI 2019 Synthesised AM/FM 80kHz-1040MHz ... £325
FLUKE 6060B AM/FM Syn Sig Gen 10kHz-1050MMHz ... £300
LEADER LSG221B Sig Gen 25-950MHz £200
HP 8656B Synthesised 0.1-990MHz £500
HP 8656A Synthesised 0.1-990MHz £400
HP 8640A AM/FM 500kHz-512MHz £150
HP 8620C Sweep Osc with 86290B 2-18.6GHz £500
HP8620C Sweep Osc with 86222B 0.01-2.4GHz £400
HP8620C/B/A with any of the following plug-ins ... £150-£200
HP 86220A Plug in 10-1300MHz
HP 86230B Plug in 1.5-4GHz
HP 86235A Plug in 1.7-4.3GHz
HP 86240A Plug In 2-8.5GHz
HP 86240C Plug in 3-6-8.6GHz
HP 86245A Plug in 5.9-12.4GHz
HP86250B Plug in 8-12.4GHz
HP 86250D Plug in 8-12.4GHz
HP 86260A Plug in 12.4-18GHz
MARCONI TF2015 AM/FM 10-520MHz £95
MARCONI TF2018 AM/FM 10kHz-120MHz £95
PHILIPS PM5328 100kHz-180MHz with 200MHz Freq Counter IEEE £225
PANASONIC VP8117A AM/FM 100kHz-110MHz FM 0-100kHz Digital Display etc. Unused £225
HP 8165A Programmable Signal Source 1MHz-50MHz (Pulse/ Function) £325
HP 3325A Synthesised Function Gen 21MHz £350
HP 3312A Function Gen 0.1Hz-13MHz AM/FM Sweep/Tri/Burst etc. £200
WAVETEK 21 Stabilised Function Gen 11MHz £225
WAVETEK 23 Synthesised Function Gen 12MHz £275
EXACT 529 AM/FM Function Gen 20MHz £150
ANALOGUE 2030 Synthesised Multi Function Waveform £250
THANDER TG503 Pulse/Function Gen 5MHz £195
THANDER TG502 Sweep/Function Gen 5MHz £195
KRON-HITE 5200A Sweep Func Gen 0.00003Hz-3MHz £150
HP 3310B Info as 3310A + etc. £120
HP 3310A Func Gen 0.005Hz-5MHz Sine/Sq/Tri/Ramp/Pulse £80
PHILIPS PM5132 Function Gen 0.1Hz-2MHz £95
PHILIPS PM5131 Function Gen 0.1Hz-2MHz £75
FEEDBACK FG601 Func Gen 0.001Hz-1MHz £60
HP 8112A Pulse Gen 50MHz £750
HP 8111A Pulse Generator 20MHz £400
LYONS PG73N Pulse Gen 20MHz £50
LEADER LAG120B Sine/Sq Audio Generator 10Hz-1MHz .. £60
FARNELL LFM4 Sine/Sq Osc 10Hz-1MHz Low Distortion, TTL Output, Amplitude Meter £60
GOULD J3B Sine/Sq Osc 10Hz-100kHz Low Distortion . £50-£75
FARNELL LF1 Sine/Sq Oscillator 10Hz-1MHz £50
MARCONI SANDERS 6055C Signal Source 850-2000MHz £125
MARCONI SANDERS 6055B Signal Source 850-2150MHz £125
MARCONI SANDERS 6056B Signal Source 2-4GHz £125
MARCONI SANDERS 6057B Signal Source 4.5-8.5GHz ... £125
MARCONI SANDERS 6059A Signal Source 12-18GHz ... £125
MARCONI SANDERS 6070A Signal Source 400-1200MHz £125
FLUKE 6011A Synthesised 11MHz £125
PHILIPS 5514V Colour Bar Generator Video £195
BLACK STAR ORION Colour Bar Gen £50
BLACK STAR ORION Later Version Metal Case £75

FREQUENCY COUNTERS/TIMERS

EIP 371 Source Locking Microwave Counter 10Hz-18GHz ... £325
EIP 331 Autohet Microwave Counter 825MHz-18GHz £195
HP 5386A Counter 10Hz-3GHz £350
FEEDBACK SC230 Counter 1.3GHz £75
RACAL 9916 Counter 10Hz-520MHz £75
RACAL 9906 Universal Counter 200MHz £95
RACAL 9904 Counter Timer 50MHz £50
RACAL 1991 Counter/Timer 160MHz 9 digit £195
MARCONI 2431A Frequency Meter 200MHz £50
MARCONI 2437 Counter/Timer 100MHz £75
HP 5340A Automet Microwave Counter 10Hz-18GHz £250
HP 5316A Universal Counter 0-100MHz HPIB £95
THANDAR TF810 Frequency Counter 5Hz-200MHz Battery ..£60
THANDAR TF200 Frequency Meter 10Hz-200MHz 8 digit ... £40
BLACK STAR Meteor 100 Counter 5Hz-100MHz £50
BLACK STAR 1325 Counter Timer 1300MHz £150
BECKMAN UC10A Universal Counter 120MHz £60
LEADER LDC9043 Digital Counter 100MHz £125

DIGITAL MULTIMETERS ETC

SOLARTRON 7150 6½ digit True RMS IEEE £75
SOLARTRON 7150Plus As Above + Temp Measurement ... £100
DATRON 1065 5½ digit Autocal AC/DC Resistance IEEE £95
FLUKE 77 3½ digit Handheld £35
FLUKE 77 Series 2 3½ digit Handheld £45
FLUKE 8060A 4½ digit True RMS Handheld £75
BECKMAN HD110 3½ digit Handheld in Carry Case £30
TTI 1905A 5½ digit Bench £60
SOLARTRON 7045 4½ digit Bench £30
AVO DA116 3½ digit with Batteries & Leads £20
AVO 8 Mk6 in Ever Ready Case with Leads etc £75
AVO 8 Mk5 with Leads etc £50
RACAL 9301A True RMS RF Millivoltmeter £125
RACAL 9300 True RMS Millivoltmeter 5Hz-20MHz usable to 60MHz £30
RACAL 9300B as 9300 ..... £45
GOODWILL GVT427 Dual Chan AC Millivoltmeter 10mV in 12 ranges 10Hz-1MHz Unused £75
KENWOOD VT176 Dual Chan Millivoltmeter £40

POWER SUPPLIES

FARNELL XA35.2T 0-35V 0-2A Twice Digital £95
FARNELL LT30-2 0-30V 0-2A Twice £110
FARNELL B30/20 30V 20A Variable No Meters £110
FARNELL B30/10 30V 10A Variable No Meters £55
FARNELL LT30-1 0-30V 0-1A Twice £75
FARNELL L30.2 0-30V 0-2A £55
FARNELL L30.1 0-30V 0-1A £40
FARNELL E350 0-350V 0-200mA £125
FARNELL D30-2T 0-30V 0-2A Twice Digital £95
THURLBY PL330 0-32V 0-3A Digital (Kenwood badged) £75
THURLBY TS3021S 0-30V 0-2A LCD £65
THURLBY PL320 0-30V 0-2A Digital £55
TAKASAGO GM035-3 0-35V 0-3A 2 Meters £45
TAKASAGO TM035-2 0-35V 0-2A 2 Meters £35
ISOLATING TRANSFORMER - Yellow - 500VA with 13Amp Socket £35

STEWART of READING

EUROCARD MasterCard VISA

17A King Street, Mortimer, Near Reading RG7 3RS
Telephone: (0118) 933 1111. Fax: (0118) 933 2375
www.stewart-of-reading.co.uk
Open 9am-5.00pm Monday to Friday (other times by arrangement)

Used Equipment – GUARANTEED. Manuals supplied
This is a VERY SMALL SAMPLE OF STOCK. SAE or Telephone for lists.
Please check availability before ordering.
CARRIAGE all units £16. VAT to be added to Total of Goods and Carriage

Simulated Software Project Driven by Quality Gates

Daniel Lübke and **Thomas Flohr** from the Software Engineering Group at the University of Hannover in Germany present their experience from a simulated software project, driven by Quality Gates

Teaching software development skills is perhaps one of the most difficult tasks in IT education because many necessary skills cannot be taught but need to be experienced by the students in real-world projects. This is especially true for social restrictions and problems, like teamwork, communication with the customer and hard deadlines, which are difficult to teach in classes.

In order to prepare our students for their future job, we decided to conduct a software project with customers, quality assurance (QA) and restrictions concerning communication and time. We successfully used quality gates (QG) to structure the project, create constant time-pressure and to better control the whole process and quality. QG is a quality checkpoint in a software project at which certain quality aspects of existing project's deliverables are checked.

Here, we present our project setup, the aspects we simulated and the experiences we (including the students) had while conducting the project.

We also focus on the areas that need improvement in the future.

Difficult task

Teaching Software Engineering (SE) practices to students is a difficult task. Most students learn their first programming language in the first year of their studies. Within the programming classes, small examples are the only viable option to show techniques and facilitate abstraction abilities. With this limited programming experience, students attend SE lectures, in which techniques designed to successfully perform larger and even many parallel projects are presented. Such SE techniques have been developed by practiced people with hands-on experience in developing software.

Therefore, students need the appropriate experience as well, in order to understand the implications and advantages of these SE techniques. Otherwise, these formalised and, sometimes, very bureaucratic methodologies seem to unnecessarily slow down projects by introducing many additional activities besides the implementation.

The solution provided by many universities is a compulsory software project or a corresponding lab. However, the questions remain how to design such a course, how close can it resemble reality and how close should it resemble reality. The main problem with this is that larger projects need certain programming and software development practice for being successful. Therefore, there is an experience dependency cycle, which needs to be overcome (see **Figure 1**).

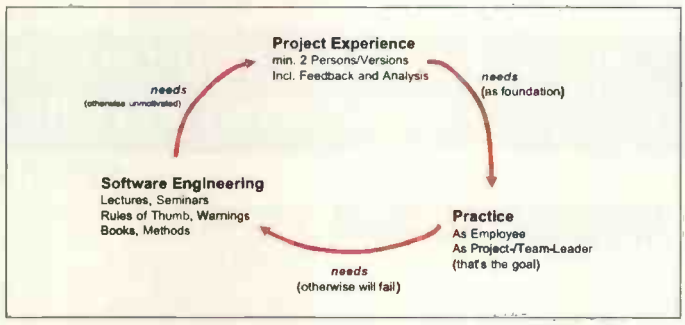

Figure 1: Dependency cycle in teaching SE skills

To overcome this dependency cycle, our university's software project is designed to convey selected experiences with chosen problems and techniques typical for software projects, like simulated time-pressure and QGs. We selected methods, which we think are the most important for the students to understand. The selection of methods allows students to break into the dependency cycle by improving their skills and gain practice and experience in certain methods. Later, they are able to improve their skill-set upon these experiences, because they are able to better understand implications of more complicated concepts, like eXtreme Programming, CMM and others.

For teaching classical project setups in our software project, we 'created' a virtual company in which nine student-teams should develop different projects. The teams had to obey certain rules. These rules should stipulate certain experiences and make the influences of certain SE techniques more transparent. Within our students' software-project we split up roles (customer, developers, quality assurance), tried to limit the time the developers (students) could spend with the customers (supervisors)

and supervised all teams with quality gates. The setup was chosen in a way that allowed students to learn teamwork, to communicate with customers and to learn QGs as a method of ensuring quality of the process and the deliverables.

Quality gates

Normally, a control committee performs the checks in a QG session, together with the project leaders. The check itself is solely based on a checklist containing the necessary deliverables and the aspects to prove. The checklist only contains formal checks, like if some properties, like versioning, specific chapters etc, are present. More in-depth inspects can be done before the QG by experts or within reviews. The positive outcome of such analysis becomes a check on the checklist.

If successful (i.e. all aspects are fulfilled or only small improvements are necessary), the project team passes the QG and can proceed to the next phase of the project. If the projects deliverables are insufficient, the team must improve or (re)create them. In the latter case, all deliverables are checked again in another QG session. At a QG, the project can also be stopped, if the deliverables show that the project is doomed to fail or other events occurred. Such outside events can include financial problems of the company or market changes making the project not marketable. The QG session's output is a protocol including the discovered flaws and the decision of the control committee.

The entire process of a QG is illustrated in **Figure 2**.

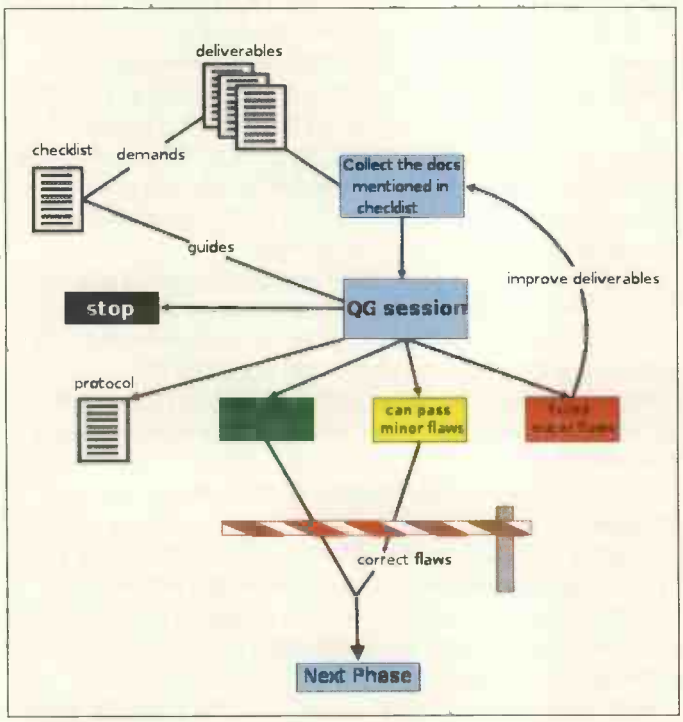

Figure 2: Procedure of a quality gate

As a methodology, the QG is used to ensure a "minimal" quality of the deliverables of a certain phase in projects' progress. The checklist is not only related to one special project, it can be seen as general checklist for different kinds of software projects or at least related projects.

Several QGs can be used in a software project to structure it. For example, at the end of any project's phase, a QG session

Checklist "Ready for Implementation"

Team data

Name of the team:			
Team member:	Name		
	...		
Submission of deliverables:	Time	Place	2nd QG?
			☐
Date for Quality Gate:			

Necessary deliverables

Software design document	☐
Schedule for tests and implementation	☐

Formal aspects

The deliverables are printed out.	☐
The deliverables are written in German or English.	☐
The deliverables are almost free of typing errors.	☐
All deliverables have a version number and a creation date.	☐
...	☐

Requirements on the software design document

The template for the software design document is used.	☐
General concepts are explained.	☐
There is graphical overview of the classes.	☐
...	☐
If possible, the UML standard is used for the diagrams.	☐
All interfaces are explained.	☐
The team successfully passed the walkthrough or review.	☐

Requirements on the schedule

The workload (e.g. modules, packages, classes) was spread on the team	☐
There is a schedule when to run the tests.	☐

Yes	☐
No	☐

The team passes the Quality Gate:

Figure 3: Extract of a quality gate checklist

can be performed. Then each QG has its own checklist, because each phase's deliverables differ from the deliverables of other phases. However, a project usually will not contain too many QGs because its progress will be slowed down by the necessary formalities.

The main advantage of a QG is its simplicity: The formal checks can be easily done and the checklist (see **Figure 3**) is a guideline that can be used even by inexperienced supervisors.

In contrast to project milestones, QGs are more general: They are used to simultaneously supervise many projects and specify what kind of deliverables have to be produced. Because in professional software organisations, all projects have to follow the same process, it is possible to define quality checks after process phases. Therefore, QGs are process-specific. Contrary, milestones are project-specific and specify the concrete attributes and contents the deliverables must have. Moreover, due to their nature, a project will have less QGs than milestones. Each QG imposed by the quality management will become a project's milestone, but projects can and normally will define many more milestones.

Because only three supervisors were available to control nine student teams, we decided to use QGs to control them, because the minimal effort was very appealing.

Project's setup

The main constraint for the software project was the available time for supervising students while maintaining a high quality level of teaching. Therefore, the use of QGs seemed appropriate and was used to structure the whole development process. The development process was divided into analysis, design and implementation. After each of these steps, a QG assured "minimal" quality of the deliverables. For example, it was checked if there are acceptance test-cases within the requirements document produced in the analysis phase. If a team failed a QG, it had one week to fix the found mistakes. After that week, a second QG checked the team's deliverables again. If the second QG failed again, the team's project was cancelled and the students would have failed. The process and time structure using QGs can be seen in **Figure 4**.

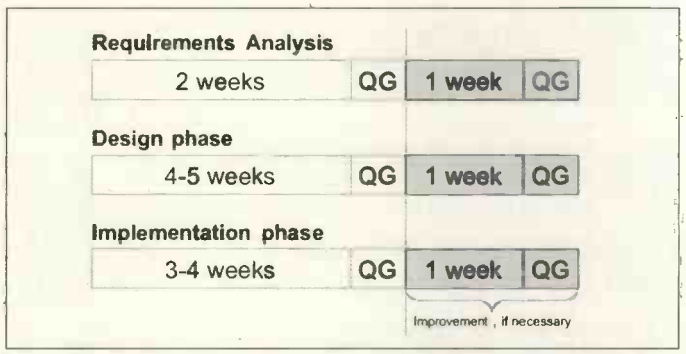

Figure 4: The software project's phases delimited by quality gates

A prerequisite for the QG were the approvals of "experts". For example, the customer needed to sign the requirements document and the design was checked by walkthroughs or reviews.

The QGs took place at a seminar room that was arranged like a courtroom, with the supervisors acting as "judges" in front of the developers (see **Figure 5**). One supervisor was going through the checklist, while the other two were checking the submitted deliverables for the given quality attributes. After the whole checklist was done, the judgement was made. If there were no failures, the team passed without having to do any changes. If there were some shortcomings in the deliverables, the supervisors decided whether the QG was failed or if it was sufficient to let the team proceed and correct the mistakes. These corrections had to be brought to the team's QA expert. Each QG session took about 20-30 minutes, with all nine taking approximately 4.5 hours.

Figure 5: A quality gate session

In addition to the QGs, roles for customer and quality assurance/software engineering expert were separated, and a voucher system limited the time students could spend with customers.

Experiences and feedback

While conducting the software project, we experienced lots of aspects of the QGs we had not thought about before. Furthermore, we got students' feedback after the whole project. To conserve these impressions and experiences, we used the LIDs approach (see below). LIDs is a light-weight technique for experience solicitation and can easily capture experiences. This is done in a discussion guided by a template in which a moderator types the arguments of the participants. Additionally, we handed out a questionnaire to the students, which also contained questions about their opinion of QGs.

Overall, the feedback was positive. They sought to structure a software project, and the checklists are useful to indicate which deliverables are necessary for the different project's phases. Nevertheless, some groups considered the QG method patronising, because the committee checked also small (formal) details, like missing page and version numbers of the deliverables. Since most groups were not aware of the QG procedure, this led to irritations at the first QG.

Students also seemed to be surprised that only the quality aspects are being checked for, mainly because they did not read the checklists carefully enough or, simply, misunderstood them. Most students had the expectation that also a document's content or product's functionality is tested and were not afraid of failing the QG. The same students also did not take the checklists very seriously and, therefore, ignored some aspects of the project. After the first QG session, students said that just writing something inside the documents should be enough to pass the QG, as long as the quality aspects are fulfilled. Two teams also mentioned that QGs are too bureaucratic.

Altogether two groups failed the first QG and had to attend a further iteration of the same QG. Nearly all other teams had minor deficiencies in their documents, but passed the first QG. All teams learned from the conduction of the first QG and mastered the next two QGs without any major problems, with the exception of one group, which failed the second quality gate in the walkthrough.

One team also was puzzled because the customer attended the quality gates in the role of a QA. This had a negative impact on the role play aspect of the software project but was a problem of our set-up, because all supervisors were attending all QGs.

Furthermore, the students answered questions about the QGs after the project. Altogether we evaluated 44 questionnaires and some results are shown in **Figure 6**.

Most students found the QGs helpful for structuring the project and think that the software project was more realistic because of their usage. Because of this, most students answered that for further software projects, QGs should be conducted again. Nearly half of the students answered, that by the imposed deadlines before the QGs, the workload was more evenly distributed across the whole project timeframe. However, nearly as many students felt that QGs were too bureaucratic. Only one student

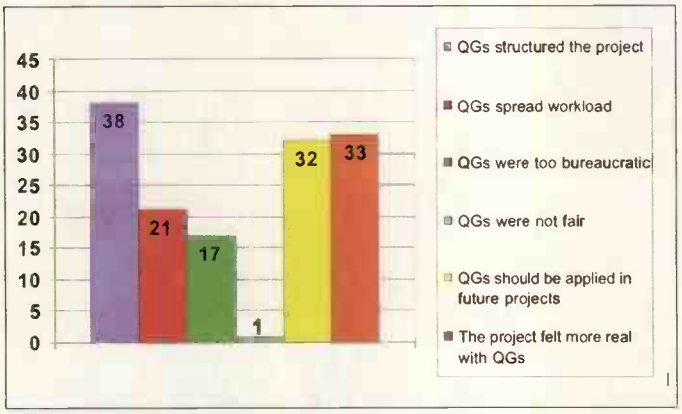

Figure 6: Students' opinion about the attributes of quality gates

thought that QGs were unfair. At this point, we would have expected a higher number of students agreeing with this option of the observations we made during the QGs.

After the second series of QGs, we decided to conduct a supervisor-only LIDs session to evaluate the results of the first two QGs.

We observed that the second QG was quite more successful than the first one, despite the fact that three staff members attended nine QG sessions (each 30 minutes) on one day and nine walkthroughs or reviews (each 120 minutes) in the same week. This workload was really exhausting and we observed that at one session we became less critical in judging the deliverables.

We also noticed that students better understood the QG concept after attending the first QG session. Prior to the second QG, the students asked the QA to check formal aspects in their documents. Quality obviously became better and, to our surprise, all teams successfully delivered software.

Improvements

From our point of view, quality gates proved as a valuable addition to software projects with students in university. QGs successfully spread the workload across the project's phases, outlined the essential deliverables of each phase and ensured a minimum quality. Nevertheless, some improvements will be done for the following software projects. Among them is editing the checklists to increase the number of software projects they can be applied to. So far, our checklists are only created for Java projects and could not be fully applied to PHP projects.

In future software projects, the customer will not be allowed to attend the QGs in order to avoid their confusion with supervisor roles.

Some minor corrections will be done to the checklists to fix some flaws. Some aspects were hard to prove (in the given time) and very few are not exactly quality aspects.

We must give a better introduction to the QG concept in the project's beginning. Although QGs were part of the SE lecture, students did not fully understand that only quality aspects are checked for or did not know that a QG is a strict barrier. Furthermore, students need to be told beforehand that QGs are very formalistic and the labour they invested to create a "perfect" document is not honoured.

Nine QGs on one day are far too many and this reduces the ability to judge critically. Because not all three scientific staff members are needed in a QG session, we will reduce that to only two persons, thereby further reducing effort to control parallel projects.

The lengths of the project's phases need to be planned by the students. Because teams had other tasks to solve, the same timeframe for e.g. doing design is unrealistic. Students will need to plan beforehand how long they might need to create the needed deliverables and then plan the QGs accordingly.

Furthermore, we are still planning to change the "pass with minor flaws" outcome of the QG. The requested changes need to be better tracked or QGs need to be strictly pass or fail. This would clearer emphasise the QG's checkpoint function. However, it will also increase student's frustration, if they fail because of some minor flaws, like a missing version number.

Outlook

Quality gates have proven to be a very useful technique to supervise many projects with minimal effort. Because of the positive feedback we got and the experiences we had ourselves, we plan to keep QGs in future software projects. For this, we will further improve the QG process. Checklists will be improved and we will reduce the number of supervisors attending the sessions. This will further reduce the effort we need to invest for supervising the teams.

Additionally, from a psychological point of view, we must better outline the difference between quality and functionality aspects to avoid disappointments. This is the aspect, which puzzled students the most.

All in all, all participants to the software project, including students and supervisors, found QGs to be a valuable tool to guide software projects.

2GHz full-performance RF Generator

£2,195

The new TGR2050 is a synthesised signal generator with a frequency range of 150kHz to 2000MHz.

Excellent frequency accuracy is matched with high stability, wide dynamic range, low phase noise and low leakage.

FM, phase and AM modulation modes are incorporated along with RS232 and GPIB (IEEE.488.2) interfaces.

The TGR2050 provides the most cost effective solution available for engineers needing a wide range high performance RF source.

- 150kHz to 2000MHz frequency range with 10Hz setability
- -127dBm to +7dBm amplitude range with 0.1dBm setability
- Better than 1ppm internal accuracy; external frequency lock
- FM, Phase and AM modulation, internal or external source
- Direct numeric entry or rotary control with user setable frequency and amplitude increments
- Non-volatile set-up memories for repetitive testing
- Full remote control via RS-232 and GPIB interfaces
- Unrivalled performance and features at this price level !

For applications with a narrower frequency requirement a 1GHz generator, the TGR1040, is available for around £1,200 .

For full details of both generators, please contact us directly, visit our website, or circle the reply number below.

TTi

Measurably better value

Thurlby Thandar Instruments Limited
Glebe Road, Huntingdon, Cambridgeshire PE29 7DR
Tel: 01480 412451, Fax: 450409, e-mail: sales@tti-test.com
Web: www.tti-test.com

THURLBY THANDAR INSTRUMENTS

A better solution in custom cables, coiled leads and terminated assemblies, takes more than a mere specialist. It takes the market leader.

Kalestead. The details are on our website.

Cable Division, Network House,
300-302 Cressing Road, Braintree, Essex CM7 3PG
Tel: 01376 349036 Fax: 01376 348976
www.kalestead.co.uk email: sales@kalestead.co.uk

KALESTEAD LIMITED

EECO

ELECTROMECHANICAL SWITCHES

Still one of the best man/machine interface components available.. Contact us for your C&K retrofits

ELASTOMER PADS

Elastomer keypads made to your design - in almost any shape, size or colour

EECO Switch is the European Sales Office for the EECO/XYE factory in Dongguan, China. The modern XYE factory of over 27,000 sq. mtrs. (300,000 sq. ft.), with over 1500 skilled workers and 50 engineers, is an experienced, full service CEM/CDM partner for electronics and other goods. Take the first step to sourcing offshore outsourced production by contacting ECCO Switch, UK.

CHINA MANUFACTURING

Tel. (+44) (0) 1954 781818 . Fax. (+44) (0) 1954 789305

www.eecoswitch.co.uk

Automatic Conversions Between Analogue and Digital Rational Transfer Functions

Aydin Kizilkaya of Pamukkale University and **Erkan Yuce** of Bogazici University, both in Turkey, present simple and fast algorithms, which automatically perform either continuous-time to discrete-time or discrete-time to continuous-time conversions for rational system transfer functions, without computing matrix inversions

Conversions between analogue and digital transfer functions (TFs) are often encountered, especially in control and filter design applications. Given an analogue transfer function $H_A(s)$, the corresponding digital transfer function $H(z)$ can be found very easily. Similarly, if the digital transfer function $H_D(z)$ is known, then the associated analogue transfer function $H(s)$ is very easy to obtain.

Tustin's approach in M. Gopal's "*Digital Control and State Variable Methods*", which is equivalent to the bilinear transformation based on the trapezoidal approximation, is generally used to derive a discrete-time description of a continuous, linear-time invariant system TF, and vice versa. However, the direct application of this approximation on higher order TFs is a highly time-consuming process in computer simulations. Furthermore, the use of Tustin's approximation directly may be a cause to make mistakes in the analogue-to-digital (A/D) or digital-to-analogue (D/A) conversions of system TFs by hand calculation. To overcome these drawbacks, some special conversion matrices have been proposed in N.K. Bose's "*Properties of the Q_n – Matrix in Bilinear Transformation*", and K. Ichige, N. Otsuka and R. Ishii's "*An Automatic Design Procedure of IIR Digital Filters from an Analogue Low-pass Filter*", (p. 228) for deriving the corresponding digital TFs from the analogue prototypes, which have replaced the hand computations of bilinear transformation by automatic procedures. Note that these procedures are only valid for the sampling period T = 2, in which the bilinear transformation is equivalent to the Tustin's approximation by taking T as T = 2.

To perform the transformations from s-domain to z-domain and vice versa, in "*The Bilinear Z Transform by Pascal Matrix and Its Application in the Design of Digital Filters*", by B. Psenicka, F. Garcia-Ugalde and A. Herrera-Camacho and B. Psenicka and F. Garcia-Ugalde's "*Z Transform from Lowpass to Bandpass by Pascal Matrix*", procedures based on the Pascal Matrix have been developed for the specific filter TFs.

In this paper, two algorithms are proposed to perform the A/D and D/A conversions automatically and fast, by using the derived conversion matrices for the arbitrary sampling period T. The first algorithm provides the conversion from s-domain to z-domain, A/D conversion, and the second one is used to perform the conversion from z-domain to s-domain, D/A conversion, for a given rational TF. These algorithms are very simple and fast, and do not involve any matrix inversions to derive the related system TFs.

Conversion from s-domain to z-domain (A/D conversion)

It is well known that a general analogue rational system TF is given by:

$$H_A(s) = \frac{a_0 + a_1 s + a_2 s^2 + \ldots + a_m s^m}{b_0 + b_1 s + b_2 s^2 + \ldots + b_n s^n}, \quad m \le n \tag{1}$$

The vectors a and b, including the parameters of numerator and denominator polynomials in (1), are defined in the forms:

$$\boldsymbol{a} = [a_0, a_1, \ldots, a_m]^T \tag{2a}$$

$$\boldsymbol{b} = [b_0, b_1, \ldots, b_n]^T \tag{2b}$$

where the coefficients $\{ai\ ;\ i = 0, 1, \ldots, m\}$ and $\{bj\ ;\ j = 0, 1, \ldots, n\}$ are real and denote the parameters of analogue system TF.

Converting an existing analogue TF into digital system TFs for desired T values is realised by the well-known conventional Tustin's approximation. This approximation performs the one-to-one mapping from s-plane to z-plane, with an advantage that the imaginary axis of the s-plane is mapped into the unit circle in the z-plane. The mapping is

defined by:

$$s = \frac{2}{T}\left(\frac{1 - z^{-1}}{1 + z^{-1}}\right) \tag{3}$$

where T stands for the sampling period. Substituting (3) into (1) gives the corresponding digital TF expressed by:

$$H(z) = \frac{c_0 + c_1 z^{-1} + c_2 z^{-2} + \ldots.. + c_n z^{-n}}{d_0 + d_1 z^{-1} + d_2 z^{-2} + \ldots.. + d_n z^{-n}} \tag{4}$$

Note that c_i and d_i (i = 0, 1,..., n) in (4) denote the coefficients of digital system TF and are the functions of sampling period T. From (4), we can get the following vectors c and d:

$$\boldsymbol{c} = [c_0, c_1, \ldots, c_n]^T \tag{5a}$$

$$\boldsymbol{d} = [d_0, d_1, \ldots, d_n]^T \tag{5b}$$

The digital TF coefficients defined in (5) can easily be obtained benefiting from the analogue TF coefficients given in (2). For this aim, the following conversion matrix is derived:

$$\boldsymbol{Q}_{kl}^{A} = [q_{ij}^{A}] \in R^{(k+1)\times(l+1)} \tag{6}$$

where the coefficients $\{q_{ij}^{A}\ ;\ i = 1, 2,\ldots, k+1$ and $j = 1, 2,\ldots, l+1\}$ correspond to the (i, j) elements of the conversion-matrix $\boldsymbol{Q}_{kl}^{A}$. For $i = 1, 2,\ldots, k+1$ and $j = 1, 2,\ldots, l+1$, the values of q_{ij}^{A} are calculated by the following equation:

$$q_{ij}^{A} = T^{n-j+1} 2^{j-1} \sum_{p=0}^{i-1} (-1)^{i-p-1} \left\langle \begin{matrix} j-1 \\ i-p-1 \end{matrix} \right\rangle \left\langle \begin{matrix} n-j+1 \\ p \end{matrix} \right\rangle \tag{7}$$

Transformation from analogue TF coefficients $\{a_i\ ;\ i = 0, 1,\ldots, m\}$ and $\{b_j\ ;\ j = 0, 1,\ldots, n\}$ into the corresponding digital TF coefficients ci and di (i = 0, 1,..., n) are simply performed by:

$$\boldsymbol{c} = \boldsymbol{Q}_{nm}^{A}\boldsymbol{a} \tag{8a}$$

$$\boldsymbol{d} = \boldsymbol{Q}_{nn}^{A}\boldsymbol{b} \tag{8b}$$

We observed that the formulae derived in (6)-(8) resemble the procedure given in K. Ichige, N. Otsuka and R. Ishii's "*An Automatic Design Procedure of IIR Digital Filters from an Analogue Low-pass Filter*", on page 228 for $T = 2$. But, it is worthwhile to note that the equations in (6)-(8) are the generalised case of the procedure given in that paper. Consequently, these formulae can be used to compute the associated digital TFs from an analogue one, for different sampling intervals T.

Conversion from z-domain to s-domain (D/A conversion)

A general digital rational system TF is defined by:

$$H_D(z) = \frac{e_0 + e_1 z^{-1} + e_2 z^{-2} + \ldots.. + e_m z^{-m}}{f_0 + f_1 z^{-1} + f_2 z^{-2} + \ldots.. + f_n z^{-n}}, \quad m \le n \tag{9}$$

where $\{e_i\ ;\ i = 0, 1,\ldots, m\}$ and $\{f_j\ ;\ j = 0, 1, \ldots, n\}$ are real and denote the coefficients of the digital rational system TF. As in the previous section, we can define the following vectors e and f:

$$\boldsymbol{e} = [e_0, e_1, \ldots, e_m]^T \tag{10a}$$

$$\boldsymbol{f} = [f_0, f_1, \ldots, f_n]^T \tag{10b}$$

For arbitrary values of sampling interval T, in order to obtain the analogue versions of the TF described in (9), we use the transformation given in (3) by inverting it with the inverse mapping defined by:

$$z^{-1} = \frac{2 - sT}{2 + sT} \tag{11}$$

Substituting (11) into (9) gives the associated analogue rational TF represented by:

$$H(s) = \frac{g_0 + g_1 s + g_2 s^2 + \ldots.. + g_n s^n}{p_0 + p_1 s + p_2 s^2 + \ldots.. + p_n s^n} \tag{12}$$

where the analogue TF coefficients g_i and p_i (i = 0, 1,..., n) are functions of the sampling period T. From (12), we can define the following vectors g and p:

$$g = [g_0, g_1, \ldots, g_n]^T \tag{13a}$$

$$p = [p_0, p_1, \ldots, p_n]^T \tag{13b}$$

The analogue TF coefficients in (13) can easily be computed by using the digital TF coefficients in (10). For this aim, the following conversion matrix is derived:

$$Q_{kl}^D = [q_{ij}^D] \in R^{(k+1)\times(l+1)} \tag{14}$$

where the coefficients $\{q_{ij}^D ; i = 1, 2,..., k+1$ and $j = 1, 2,..., l+1\}$ correspond to the (i, j) elements of the conversion matrix Q_{kl}^D. For $i = 1, 2,..., k+1$ and $j = 1, 2,..., l+1$, the values of are calculated by the following equation:

$$q_{ij}^D = T^{i-1} 2^{n-i+1} \sum_{p=0}^{i-1} (-1)^{i-p-1} \left\langle \begin{matrix} j-1 \\ i-p-1 \end{matrix} \right\rangle \left\langle \begin{matrix} n-j+1 \\ p \end{matrix} \right\rangle \tag{15}$$

Transformation from digital TF coefficients $\{ei ; i = 0, 1,..., m\}$ and $\{f_j ; j = 0, 1,..., n\}$ into the associated analogue TF coefficients g_i and p_i (i = 0, 1,..., n) are simply performed by:

$$g = Q_{nm}^D e \tag{16a}$$

$$p = Q_{nn}^D f \tag{16b}$$

As a result, for different T values, derived equations (14)-(16) can be used to produce the corresponding analogue rational system TFs.

Note that the expressions like $\left\langle \begin{matrix} x \\ y \end{matrix} \right\rangle$; in (7) and (15) are defined by:

$$\left\langle \begin{matrix} x \\ y \end{matrix} \right\rangle = \begin{cases} 0, & y > x \text{ or } y < 0 \\ \dfrac{x!}{y!(x-y)!}, & \text{otherwise} \end{cases}$$

Remarks on the proposed algorithms

The s-domain to z-domain, A/D, conversion method presented in this paper resembles and it follows closely the formulation given in K. Ichige, N. Otsuka and R. Ishii's "*An Automatic Design Procedure of IIR Digital Filters from an Analogue Low-pass Filter*", for $T = 2$, but it adds the generality that it is valid for any sampling period T. In control or signal processing applications, the sampling period T may be chosen as a different value, which usually is not equal to 2.

A by-product of this paper is a method suggested to obtain the conversion from z-domain to s-domain, D/A conversion, without requiring any matrix inversion. A practical application for such conversion may be the stability analysis of a discrete time system for which the sampling period is known. One method of the stability analysis of a characteristic polynomial in the z-domain is the Jury test, which is quite cumbersome. The proposed conversion method from z-domain to s-domain will yield a characteristic equation in s-domain that can be tested with the very simple Routh-Hurwitz test.

Computer simulations

Two examples are considered to show the capability of the proposed algorithms. Firstly, we shall transform a given analogue low-pass TF $H_A(s)$ to the associated digital TFs $H(z)$, by using the formulae derived in (6)-(8), for different values of sampling period T. Next, the corresponding analogue TFs $H(s)$ will be derived automatically from a given digital band-pass TF $H_D(z)$, by using the formulae derived in (14)-(16) for different T values.

Example 1: Conversion from s-domain to z-domain (A/D conversion)

Let the existing analogue rational system TF $H_A(s)$ be of the sixth-order Butterworth low-pass filter characteristic whose coefficients are computed by using the table given in S. Franco's "*Design with Operational Amplifiers and Analog Integrated Circuits*" on page 169, assuming that the angular resonance frequency equals one. Then the coeffi-

cient vectors and the corresponding TF for this filter are given as:

$$\boldsymbol{a} = [1]^T$$

$$\boldsymbol{b} = [1 \quad 3.8626 \quad 7.4624 \quad 9.1338 \quad 7.4624 \quad 3.8626 \quad 1]^T \tag{17}$$

and:

$$H_A(s) = \frac{1}{1 + 3.8626s + 7.4624s^2 + 9.1388s^3 + 7.4624s^4 + 3.8626s^5 + s^6}$$

In order to transform the analogue TF in (17) into the corresponding digital rational system TFs as in the form of (4), the coefficients ci and di (i = 0, 1,..., n) must be found. Thus, the conversion matrices whose elements are obtained by (7) are firstly established as in (6). In this example, we observed that the respective orders of the numerator and denominator polynomials are m = 0 and n = 6. Hence, the following conversion matrices named as Q^A_{nm}; and Q^A_{nn} are to be used for computing the coefficients of the corresponding digital TFs. They can be obtained as the function of sampling period T:

$$Q^A_{nm} = \begin{bmatrix} T^6 \\ 6T^6 \\ 15T^6 \\ 20T^6 \\ 15T^6 \\ 6T^6 \\ T^6 \end{bmatrix}; \quad Q^A_{nn} = \begin{bmatrix} T^6 & 2T^5 & 4T^4 & 8T^3 & 16T^2 & 32T & 64 \\ 6T^6 & 8T^5 & 8T^4 & 0 & -32T^2 & -128T & -384 \\ 15T^6 & 10T^5 & -4T^4 & -24T^3 & -16T^2 & 160T & 960 \\ 20T^6 & 0 & -16T^4 & 0 & 64T^2 & 0 & -1280 \\ 15T^6 & -10T^5 & -4T^4 & 24T^3 & -16T^2 & -160T & 960 \\ 6T^6 & -8T^5 & 8T^4 & 0 & -32T^2 & 128T & -384 \\ T^6 & -2T^5 & 4T^4 & -8T^3 & 16T^2 & -32T & 64 \end{bmatrix} \tag{18}$$

As an example, if the sampling period T is taken as T = 10 in (18), the corresponding digital TF coefficients can then be obtained easily from the expressions defined in (8a) and (8b):

$$\boldsymbol{c} = 10^6[1 \quad 6 \quad 15 \quad 20 \quad 15 \quad 6 \quad 1]^T$$

$$\boldsymbol{d} = 10^6[2.1574 \quad 9.6579 \quad 18.34 \quad 18.8525 \quad 11.0411 \quad 3.4876 \quad 0.4636]^T$$

Thus, if these coefficients are placed in (4), the digital rational TF corresponding to (17) is in the form of:

$$H(z) = \frac{1 + 6z^{-1} + 15z^{-2} + 20z^{-3} + 15z^{-4} + 6z^{-5} + z^{-6}}{2.1574 + 9.6579z^{-1} + 18.34z^{-2} + 18.8525z^{-3} + 11.0411z^{-4} + 3.4876z^{-5} + 0.4636z^{-6}}$$

Frequency characteristics of the analogue and associated digital low-pass TFs are depicted in **Figure 1**. In the same figure, four frequency characteristics of associated digital TFs, obtained by using the conversion matrices defined in (18) for T = 0.7, T = 1.8, T = 4 and T =10 are plotted.

Example 2: Conversion from z-domain to s-domain (D/A conversion)

In this example, we shall transform a given digital TF $H_D(z)$ to the associated analogue TFs $H(s)$ for arbitrary T values. The digital TF in the form of (9) has the characteristic of fourth-order band-pass filter whose coefficient vectors and the TF are given by:

$$\boldsymbol{e} = [0.3499 \quad 0 \quad -0.6998 \quad 0 \quad 0.3499]^T$$

$$\boldsymbol{f} = [1 \quad 0 \quad -0.2194 \quad 0 \quad 0.1801]^T \tag{19}$$

and:

$$H_D(z) = \frac{0.3499 - 0.6998z^{-2} + 0.3499z^{-4}}{1 - 0.2194z^{-2} + 0.1801z^{-4}}$$

Digital rational system TF in (19) is also considered in B. Psenicka and F. Garcia-Ugalde's "*Z Transform from Lowpass to Bandpass by Pascal Matrix*".

As in the previous example, firstly, the conversion matrices whose elements are calculated by using (15) are estab-

lished in the form of (14). Note that these matrices are the same since the orders of numerator and denominator polynomials are equal, i.e. $m = 4$ and $n = 4$ in the chosen example. In this case, conversion matrices are calculated as:

$$Q^D_{nm} = Q^D_{nn} = \begin{bmatrix} 16 & 16 & 16 & 16 & 16 \\ 32T & 16T & 0 & -16T & -32T \\ 24T^2 & 0 & -8T^2 & 0 & 24T^2 \\ 8T^3 & -4T^3 & 0 & 4T^3 & -8T^3 \\ T^4 & -T^4 & T^4 & -T^4 & T^4 \end{bmatrix} \quad (20)$$

As an example, if the sampling period T is taken as $T = 0.95$ in (20), the corresponding analogue band-pass TF co-efficients can easily be obtained from the expressions defined in (16a) and (16b):

$$g = [0 \quad 0 \quad 20.2102 \quad 0 \quad 0]^T$$

$$p = [15.3712 \quad 24.9250 \quad 27.1450 \quad 5.6237 \quad 0.7825]^T$$

Thus, if these coefficients are placed in (12), the analogue rational TF corresponding to (19) is in the form of:

$$H(s) = \frac{20.2102s^2}{15.3712 + 24.9250s + 27.1450s^2 + 5.6237s^3 + 0.7825s^4}$$

Frequency characteristics of the digital and associated analogue band-pass TFs are illustrated in **Figure 2**. In the same figure, four frequency characteristics of associated analogue TFs obtained by using the conversion matrices given in (20) for $T = 0.2$, $T = 0.52$, $T = 0.95$ and $T = 1.4$ are plotted.

Automatic and fast

In this paper, two simple algorithms have been proposed to perform the conversions from the s-domain to z-domain (A/D) and z-domain to s-domain (D/A) by choosing the sampling period T arbitrarily. These conversions are realised with the conversion matrices that are the functions of sampling period T. In the proposed algorithms, we do not need to calculate any matrix inversions to arrive at the relevant system TFs, so that we need only simple matrix operations. It is worthwhile to point out that the presented algorithms are simple and fast, and they provide great facilities to perform the A/D and D/A conversions, especially for higher order rational system TFs.

They are particularly useful in the conversions of higher order transfer functions and their common feature is that the sampling interval T can be chosen arbitrarily.

Several examples have been presented to demonstrate how the proposed algorithms work. Essentially, these procedures can easily be carried out on a computer automatically and fast.

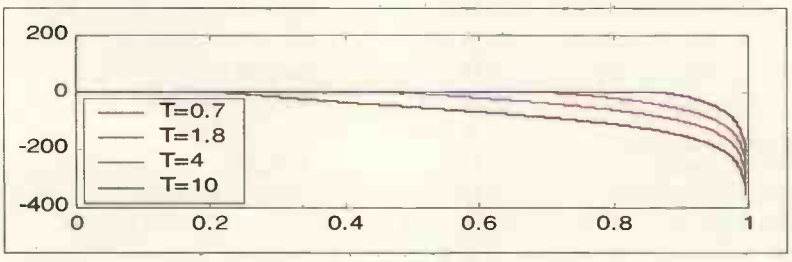

Figure 1: Frequency characteristics of the analogue low-pass transfer function and its corresponding digital versions for different T values

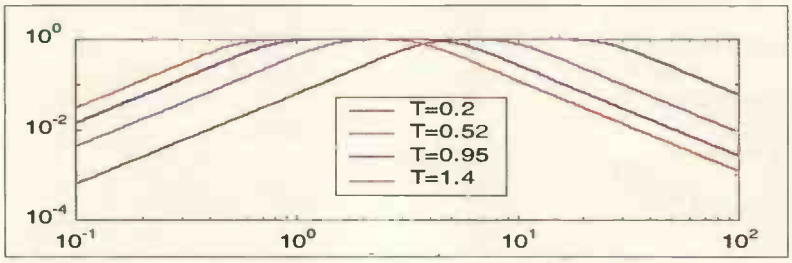

Figure 2: Frequency characteristics of the digital band-pass transfer function and its associated analogue versions for different T values

PICmicro Microcontroller DC Motor Control

Every motor control circuit can be divided into drive electronics and controlling software. These two pieces can be fairly simple or extremely complicated, depending upon the motor type, the system requirements and the hardware/software complexity trade-off.

Generally, higher performance systems require more complicated hardware. The tips below describe some basic circuits and software building blocks commonly used to control motors.

TIP 1: Drive software: Pulse-Width Modulation (PWM) algorithms

Pulse-Width Modulation (PWM) is critical to modern digital motor controls. By adjusting the pulse width, the speed of a motor can be efficiently controlled without larger linear power stages. Some PICmicro MCU devices have hardware PWM modules on them. These modules are built into the Capture/Compare/PWM (CCP) peripheral. CCP peripherals are intended for a single PWM output, while the Enhanced CCP (ECCP) is designed to produce the complete H-Bridge output for bidirectional Brushed DC motor control. If cost is a critical design point, and a PICmicro microcontroller device with a CCP module may not be available, software-generated PWM is a good alternative.

The following algorithms are designed to efficiently produce an 8-bit PWM output on the mid-range family of PICmicro microcontrollers. These algorithms are implemented as macros. If you want these macros to be a subroutine in your program, simply remove the macro statements and replace them with a label and a return statement.

Code 1: 1 output 8-bit PWM

```
pwm_counter     equ xxx  ;variable
pwm             equ xxx  ;variable
set_pwm macro A          ;sets the pwm
;setpoint to the         ;value A
MOVLW A
MOVWF pwm
endm
update_PWM macro         ;performs one update
                         ;of the PWM signal
                         ;place the PWM output
                         ;pin at bit 0 or 7 of
                         ;the port
MOVF pwm_counter,w
SUBWF pwm, w                      ;if the output
;is on bit 0
RLF     PORTC,f          ;replace PORTC with
;the correct port if
                         ;the output is on bit
                         ;7 of the port
                         ;replace the rlf with
                         ;rrf incf
                         ;pwm_counter,f
```

Code 2: 8 output 8-bit PWM

```
pwm_counter     equ xxx          ;variable
pwm0            equ xxx          ;
pwm1            equ pwm0+1
pwm2            equ pwm1+1
pwm3            equ pwm2+1
pwm4            equ pwm3+1
pwm5            equ pwm4+1
pwm6            equ pwm5+1
pwm7            equ pwm6+1
output          equ pwm7+1
set_pwm macro A,b;sets pwm b with
;the value A
MOVLW pwm0
ADDLW b
MOVWF fsr
MOVLW a
MOVWF indf
endm
update_PWM macro                 ;peforms one
                                 ;update of all 8
                                 ;PWM signals
                                 ;all PWM signals
                                 ;must be on the
                                 ;same port
MOVF            pwm_counter,w
SUBWF                   pwm0,w
RLF             output,f
MOVF            pwm_counter,w
SUBWF                   pwm1,w
RLF             output,f
MOVF            pwm_counter,w
SUBWF                   pwm2,w
RLF             output,f
MOVF            pwm_counter,w
SUBWF                   pwm3,w
RLF             output,f
MOVF            pwm_counter,w
SUBWF                   pwm4,w
RLF             output,f
MOVF            pwm_counter,w
SUBWF                   pwm5,w
RLF             output,f
MOVF            pwm_counter,w
SUBWF                   pwm6,w
RLF             output,f
MOVF            pwm_counter,w
SUBWF                   pwm7,w
RLF             output,w
MOVWF                   PORTC
INCF            pwm_counter,f
endm
```

TIP 2: Writing a PWM value to the CCP registers with a mid-range PICmicro MCU

The two PWM LSb's are located in the CCPCON register of the CCP. This can make changing the PWM period frustrating for a developer. Codes 3 to 5 show three macros written for the mid-range product family that can be used to set the PWM period.

The first macro takes a 16-bit value and uses the 10 MSb's to set the PWM period. The second macro takes a 16-bit value and uses the 10 LSb's to set the PWM period. The last macro takes 8 bits and sets the PWM period. This assumes that the CCP is configured for no more than 8 bits.

Code 3: Left-justified 16-bit macro

```
pwm_tmp equ xxx         ;this variable must be
                        ;allocated someplace
setPeriod macro a       ;a is 2 SFR's in
                        ;Low:High arrangement
                        ;the 10 MSb's are the
                        ;desired PWM value
RRF a,w                         ;This macro will
                        ;change w
MOVWF                   pwm_tmp
RRF             pwm_tmp,w
ANDLW                   0x30
IORLW                   0x0F
MOVWF                   CCP1CON
MOVF            a+1,w
MOVWF                   CCPR1L
```

Code 4: Right-justified 16-bit macro

```
pwm_tmp equ xxx         ;this variable must be
                        ;allocated someplace
setPeriod macro a       ;a is 2 bytes in
                        ;Low:High arrangement
                        ;the 10 LSb's are the
                        ;desired PWM value
SWAPF a,w               ;This macro will
                        ;change w
ANDLW                   0x30
IORLW                   0x0F
MOVWF                   CCP1CON
RLF             a,w
IORLW                   0x0F
MOVWF                   pwm_tmp
RRF             pwm_tmp,f
RRF             pwm_tmp,
```

Code 5: 8-bit macro

```
pwm_tmp equ xxx         ;this variable must be
                        ;allocated someplace
setPeriod macro a       ;a is 1 SFR
SWAPF a,w               ;This macro will
                        ;change w
ANDLW           0x30
IORLW 0x0F
MOVWF           CCP1CON
RRF     a,w
MOVWF           pwm_tmp
RRF     pwm_tmp,w
MOVWF           CCPR1L
```

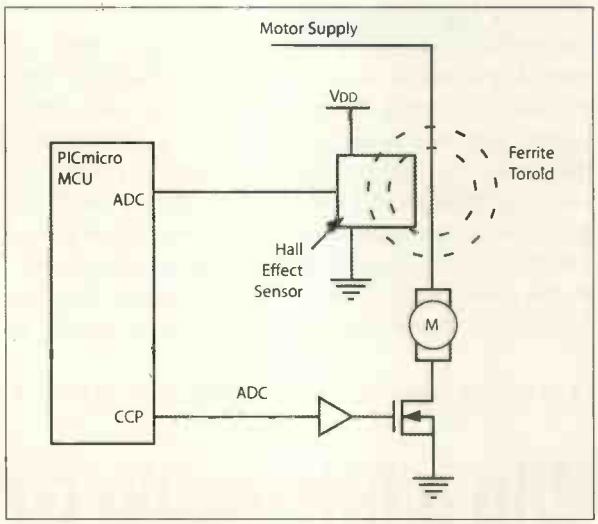

Figure 1: Resistive high-side current sensing

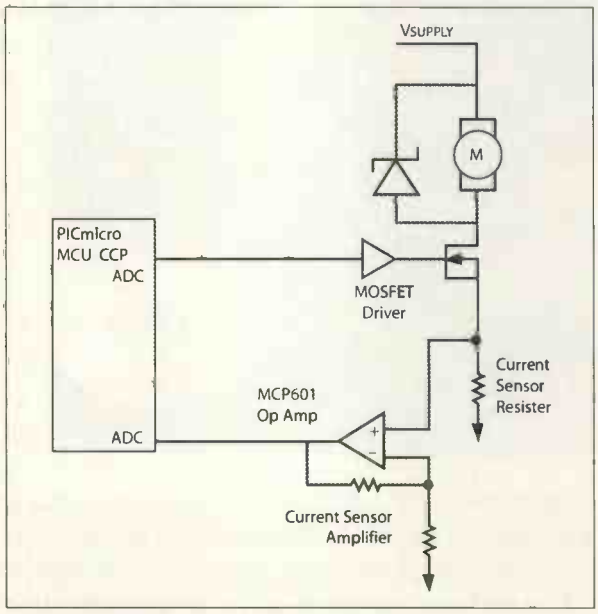

Figure 2: Resistive low-side current

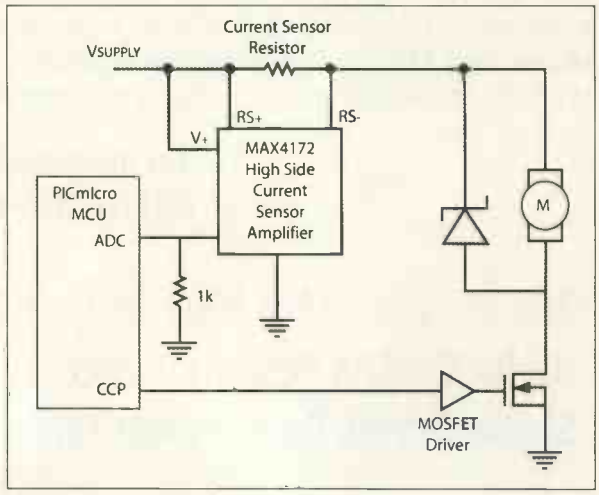

Figure 3: Magnetic current sensing

TIP 3: Current sensing

The torque of an electric motor can be monitored and controlled by keeping track of the current flowing through the motor. Torque is directly proportional to the current. Current can be sensed by measuring the voltage drop through a known value resistor or by measuring the magnetic field strength of a known value inductor. Current is generally sensed at one of two places: the supply side of the drive circuit (high-side current sense) or the sink side of the drive circuit (low-side current sense). Low-side sensing is much simpler but the motor will no longer be grounded,causing a safety issue in some applications. High-side current sensing generally requires a differential amplifier with a common mode voltage range within the voltage of the supply.

Current measurement can also be accomplished using a Hall Effect sensor to measure the magnetic field surrounding a current carrying wire. Naturally, this Hall Effect sensor can be located on the high side or the low side of the load. The actual location of the sensor does not matter because the sensor does not rely upon the voltage on the wire.

Win a Microchip rfPIC Development Tool

Electronics World is offering its readers the chance to win a Microchip rfPIC Development Kit. The kit provides an easy way to evaluate low-power RF communication links for embedded control applications. Designed to work in tandem with the PICkit 1 Flash Starter Kit, the rfPIC Development Kit 1 includes transmitter and receiver modules supporting frequencies of 315MHz and 433MHz.

The receiver modules feature an rfRXD0420 device that plug directly into the PICkit 1 development board. The modules are available separately so designers can create several prototypes based on the same module, without having to develop an actual RF design. All the design files are available, offering users the ability to migrate their module design into the application for lower cost volume production. Target applications for the rfPIC family include remote control, wireless sensors, automotive and home security.

The self-contained rfPIC12F675 transmitter modules are based on a PICmicro 20-pin Flash microcontroller that features an integrated UHF RF transmitter. The transmitter modules feature button inputs for remote control functions and analogue input that can be used for the evaluation of the microcontroller A/D converter peripherals. Code can be developed using Microchip's MPLAB Integrated Development Environment. Programming the microcontroller is easily accomplished by plugging the modules into the PICkit 1 Flash Starter Kit.

For the chance to win, log onto
www.microchip-comp.com/ew-rfpic

The winner of the Microchip PICDEM MC Development Board for Motor Control Design competion is:

Stefan Hansen, Development Technical Engineer, OJ Electronics, Denmark

RoHS
WHAT'S ALL THE FUSS ABOUT?

RoHS (the Restriction of the use of certain Hazardous Substances) and WEEE (Waste Electrical and Electronic Equipment) is the equivalent of Y2K for the electronics industry. The upcoming EU environmental directives are the most significant developments in electronics legislation to happen in many years and will completely revolutionise the way electrical and electronic products are designed, sold, recovered and recycled. Worryingly, many design engineers are still not fully aware that the upcoming legislation will affect them. For those who are, many questions remain unanswered. Complicated exemption rules, uncertainty about how the directives will be enforced, obsolescence and component availability has left engineers unsure of what they need to do and when. The clock is ticking. With only several months to go, there's no time to lose in the transition to RoHS. If compliant components aren't already part of the design cycle it could well be too late.

Q: What is the status of materials that have crystalline structure, like steel or granite, where some crystals may fail ROHS?

The use of homogeneous here seems at odds with common usage. For example, homogeneous granite means a collection of several samples of granite of like kind. There is an implied plural for comparison, and objects are not required to have a uniform composition.

Does "mechanically disjointed" really require granite or steel to be separated into crystals, or does "mechanically disjointed" really means taken apart with tools like screwdrivers? Use of tools like screwdrivers would allow homogeneous to retain its normal meaning.

Duncan Irvine, UK

A: *Many materials are not homogeneous if examined with a sufficiently powerful microscope but the definition given by the EC is provided to explain the intended meaning of "homogeneous materials". The term mechanical disjointing is used as it indicates methods where one material can be physically separated from another and the example techniques given are all simple methods – no microscope required. Therefore, it is the intention that a plastic containing particles of pigment, filler, etc is one homogeneous material. Granite would also be one homogeneous material as would a crystalline alloy but a coating on steel which can be mechanically disjointed (e.g. abraded off) would be a separate homogeneous material.*

Q: Are batteries covered by RoHS?

A: *No, although they will be recycled at equipment end of life. There is a dedicated directive in the pipeline for batteries and it is expected that batteries, such as NiCads, will be banned other than exemptions such as medical, security and handheld equipment.*

Q: Are battery chargers covered by the Battery Directive?

A: *No. The RoHS category a battery charger will fall into is dependent on its application... a charger for a mobile phone would be category 3, a hand drill category 6 etc.*

Q: Is there any impact on manufacturers that operate outside the European Union (EU)?

A: *A manufacturer outside the EU, who exports equipment into an EU country where they have a "presence" such as a subsidiary, sales office, distributor etc, has to register in that country and take care of the financial provision for the recycling of the equipment at end of life.*

Q: What is the current status on exemptions?

A: *At the time of writing, 14 exemptions are already in place, a further eight have been approved, 19 are under review and 23 new ones are currently out to tender.*

Q: What can I do with equipment that I have not sold on July 1st 2006?

A: *Providing the equipment is placed on the market, in a warehouse etc, it can be sold after July and also be maintained or repaired using non-compliant components. All equipment put on the market from July 1st 2006 must be RoHS compliant.*

Q: Will a certificate of compliance satisfy due diligence?

A: *A certificate will go towards satisfying due diligence but it is also recommended to have a robust "risk assessment" procedure in place covering both high-risk manufacturers and high-risk materials.*

Q: Who takes on the Producer liability when a piece of equipment is refurbished?

A: *If the equipment is refurbished, but otherwise unchanged, the original Producer remains responsible for recycling at end of life.*

Q: What about Military?

A: *RoHS does not apply to military/national security equipment, but only where the equipment's sole use is military/national security.*

Q: Which lead-free solder is recommended?

A: *Tin-silver-copper (SnAgCu) is the global alloy choice. Over time it is possible that more bismuth may be used to help bring the melting point temperature down (today it is important that there is no contamination with lead).*

Gary Nevison is chairman of the AFDEC RoHS team, board director at Electronics Yorkshire and head of product market strategy at Farnell InOne. As such he is our industry expert who will try and answer any questions that you might have relating to the issues of RoHS and WEEE.
Your questions will be published together with Gary's answers in the following issues of Electronics World.
Please email your questions to EWeditor@nexusmedia.com, marking them as RoHS or WEEE.

Designing Autonomous Mobile Robots

John M. Holland

Elsevier

I cannot remember if I have ever before read a technical book without skipping through pages to get to the next interesting bit. John Holland's work on Designing Autonomous Mobile Robots is the exception. I read, from cover to cover, every interesting, informative and often amusing paragraph in six sessions over a period of one weekend.

"Inside the Mind of an Intelligent Machine" is an alternative title given on the book's cover, which is more indicative of its true content. It is a dissertation on design philosophy whose main theme is to describe the manner in which one should organise the approach to the design of a complex project, autonomous mobile robots in particular. The physical creation of the necessary system hardware and software designs is left to other tomes.

The book has been organised into two parts:

Section 1 provides background information and introduces the types of software complexity needed to perform automatic control over an autonomous mobile robot. Its six chapters deal with system architecture, software concepts, real time software basics, fuzzy logic, closed loop control and communications. To illustrate some of the ideas, skeletal software examples are given in these chapters.

Section 2 comprises of thirteen chapters. The first chapter deals with basic navigation philosophy followed by a chapter on navigation using "live reckoning" as opposed to "dead reckoning". A comparison of the use of fuzzy logic against hard logic to perform navigation is made in chapter 11. A description of sensor types is followed by how sensor data can be misleading or contradictory and cause arbitration problems in navigation. A fascinating chapter how the equivalent of pain, fear and confidence can be induced into the robot's software is followed by an interesting discussion on how robot movement and timing of sensor data can cause erroneous positional information. Chapter 15 describes how to program a robot so that it can perform useful functions. This is followed by a chapter on various methods of how a robot can be commanded and monitored.

The chapter containing the subjects "The Law of Conservation of Defects" and "The "Art of Debugging" should be a compulsory subject for study by any budding design engineer. In this chapter Holland describes how he and his colleagues formed a hypothesis that defects cannot be created or destroyed and presents an exposé on how and why software and hardware bugs enter and leave systems. The reason why these suppositions never became a recognised theory is an unfortunate story and an object lesson why the choice of the research venue is important.

Holland also takes a close look at fault-finding and shows that there is more than one method of finding bugs in a system. He indicates that there are three basic approaches namely "SWAG", "divide and conquer" and "substitution". There is a comprehensive description of the various bug types, giving some indication of the length and direction of each of their garden paths. Those who have been involved in the design of a complex logical system will recognise the perplexities described in this chapter and can mutter with feeling "been there, done that, got the scars to prove it". The subject matter in this paragraph, although treated in an amusing manner, is of much significance when designing or debugging a system.

The penultimate chapter contains a discourse on why it is essential to keep a log of all events in order that accurate reports can be made when there is a problem. The final chapter gives an insight into the robotics industry and how technological, commercial and political factors have dictated the course of events. Holland bemoans the manner on which banks and venture capitalists have had a hand in constraining development. A brief history and present trends of the industry is given. A useful set of appendices closes the book.

A CD containing the full text of the book in "portable document (.pdf)" format and source code for the examples is included.

It is a well written, informative and often amusing dissertation, full of useful anecdotes. I would recommend it to any engineer.

Mike Button

Demystifying Switching Power Supplies

Raymond A. Mack, Jr.

Newnes (Elsevier)

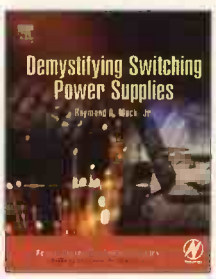

This is a title in the 'Demystifying Technology Series', described as "Books By Engineers, For Engineers". This book certainly fits that description: written from the perspective of an experienced and practising designer, it is aimed squarely at the design engineer whose field is not primarily in switching power supplies, but who is faced with the task of designing one. Theory is kept to a minimum; the emphasis is on the practical design choices, component selection and design calculations, required to get to the point of building a working design that can then be tested and refined in the lab. The mathematics should present no difficulty to most engineers – it is mostly the application of familiar relations e.g.,

$V = L(di/dt)$, $X_c = 1/[2\ pi\ f\ C]$.

Chapter 1 describes the basic operation of the circuit topologies most commonly used: buck, boost, inverting, buck-boost and various transformer-isolated converters (flyback, forward, half and full bridge). Inductorless charge-pump circuits are also described – this is a useful topic to include, as these can often be a more suitable alternative than the boost circuit that uses an inductor.

The resonant mode converter is briefly discussed in Chapter 2, but not treated further. There is a brief discussion (with a design example) of the SEPIC converter in Chapter 4.

Chapter 2 is devoted to the control circuits that provide such functions as closed-loop regulation of the output voltage, current limiting and soft-start. Control theory

(which could have occupied many densely mathematical pages) is mostly avoided – justifiably so, given the practical purpose of the book and the fact that most controller ICs used today are current mode, making loop compensation easier.

The author describes a practical technique (from a Linear Technology application note) he uses to optimise the loop response. This section of the book could have been made more helpful by including some screen-shots from an oscilloscope, showing what to recognise in the transient response during the adjustment procedure.

The input circuit (for off-line supplies) is discussed in **Chapter 3**. This gives brief coverage of input-voltage switching, EMI filtering, power-factor control, inrush current limiting, input hold-up time and safety agency requirements.

Chapters 4 and **5** contain one or two design examples for each of the topologies introduced in Chapter 1. Here, it is as if we are sat next to the author as he goes through the process of choosing components from manufacturer's websites, applying rules-of-thumb, jotting down calculations on a notepad and occasionally backtracking when things don't meet the objectives he is trying to achieve. Not the easiest bedtime reading, but this will be the most frequently referred-to part of the book when we have to do a design ourselves.

Of particular value is the discussion of alternative methods of getting a feedback signal from the output, across the isolation barrier, to the switch driver on the primary side of off-line supplies. Conventional circuits (using an opto-coupler) are also described.

An appropriate choice of components is of vital importance to achieving performance and reliability. **Chapter 6** describes the characteristics of various types of resistors and capacitors (including Niobium and Solid Polymer). **Chapter 7** describes semiconductor components: first the characteristics of various diodes are contrasted; then bipolar and Mosfet transistors are discussed, together with the circuits around them (drivers and snubbers) that are needed for performance and reliability. IGBTs are briefly discussed.

Chapter 8 introduces us to the selection and design of inductors, describing core materials and design calculations. Chapter 9 similarly treats transformers and includes design examples for a flyback transformer and a forward converter core.

The book concludes with two complete design examples, bringing together the earlier material: a "true sine wave" inverter/uninterruptible power supply, and an off-line supply for a PC.

Throughout the book, the author points us to many valuable application notes available on the Internet. For most of us, this will be a more easily accessed source of additional information than other books and is of more immediate value to the design task in hand. It would have been useful, though, if the references were gathered together in an annotated bibliography.

Another quibble I have is with the illustrations. In one figure, shaded areas that are referred to in the text (essential to understanding the explanation) are missing. In addition, the axes on some of the graphs should be better labelled.

More thorough exposition and detail can be found in the classic guides in this field: Pressman's 'Switching Power Supply Design' or Billings' 'Switchmode Power Supply Handbook'. This book, being recently written (2004/5), has the advantage of discussing newer components and is more focused on the essentials of getting designs up and running.

David Ashby

If you want to buy a copy of any of the books featured in Electronics World, you will receive a 15% discount, by writing to us at the usual address or emailing us on

EWeditor@nexusmedia.com

A Novel Single-Capacitor, Single-Current-Conveyor, Sinusoidal Oscillator

The well-known current-conveyor-based Schmitt-trigger square wave generator is shown to behave as a sinusoidal oscillator at relatively high frequencies. This is attributed to the widely ignored parasitic elements associated with terminal y of the current conveyor.

Figure 1 shows a simplified model for the second-generation plus-type current-conveyor (CCII+). This circuit can also be used for modelling the current-feedback operational-amplifier (CFOA) formed of a CCII+ followed by a voltage-follower.

At relatively low frequencies, the parasitic capacitances C_x, C_y and C_z, and the parasitic resistors R_x, R_y and R_z can be ignored and the CCII+ can be treated as an ideal device with $i_z = i_x$, $v_x = v_y$ and $i_y = 0$. At relatively high frequencies these parasitic elements cannot be ignored and usually they produce additional poles that may degrade the performance of the circuits built around the CCII+ or the CFOA.

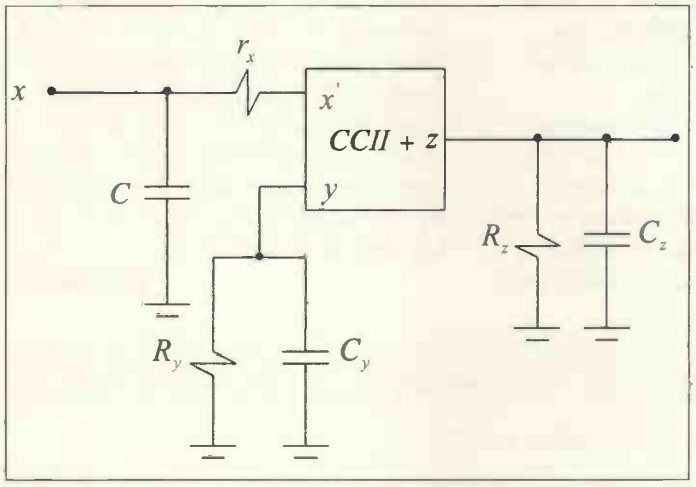

Figure 1: Simplified model for the AD844

$(r_x = 50\Omega, R_y = 10M\Omega, R_z = 3M\Omega, C_x = C_y = 2pF, C_z = 4.5pF)$

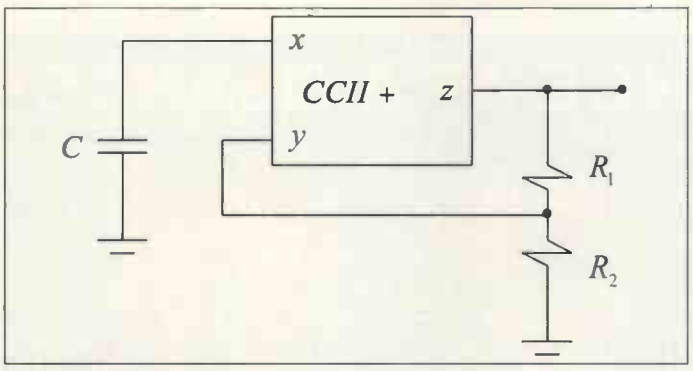

Figure 2: CCII-based square wave generator

However, by exploiting these parasitic elements, it is possible to develop new CCII- and CFOA-based circuits using reduced number of externally connected passive components. For example, resistance R_x has been exploited to develop a current-controlled current-conveyor and a new range of current-conveyor-based circuits has also been developed using only externally connected capacitors.

On the other hand, by exploiting the capacitance C_z and the resistance R_z, CCII- and CFOA-based sinusoidal oscillators using reduced number of externally connected passive components were developed. No attempt has been reported for exploiting to advantage the parasitic elements C_y and R_y.

On the other hand, the current-conveyor based Schmitt-trigger circuit shown in **Figure 2** is widely used as a square wave generator. In the analysis of this circuit, it is usually assumed that the current conveyor non-idealities are represented by the parasitic capacitance C_x and the parasitic resistance R_x associated with terminal x. Usually, the parasitic resistance R_y and the parasitic capacitance C_y associated with terminal y are ignored.

No attempt has been reported for studying the performance of the circuit of Figure 2 if the effect of the parasitic at terminal y is taken into consideration. It is, therefore, the major intention of this paper to investigate the performance of this circuit using the model shown in Figure 1.

Analysis

Using the CCII+ simplified model of Figure 1, routine analysis yields the characteristic equation of the circuit of Figure 2 that can be expressed as **Equation 1**:

$$sC_T R_1 (R_2 // R_y) - (1 + sC_T R_x)(R_1 (1 + sC_y (R_2 // R_y)) = 0$$

where $C_T = C + C_x$ and C are the externally connected capacitance. Therefore, by equating the real and imaginary parts of Equation (1) to zero, i.e. using the Barkhausen criterion, the frequency and condition of oscillation of the circuit of Figure 2 can be expressed as

$$\omega_o^2 = \frac{1}{C_T C_y R_x (R_2 // R_y)} \quad (2)$$

and

$$(R_2 // R_y)(C_T - C_y) = C_T R_x \quad (3)$$

Inspection of Equations (2) and (3) clearly shows that by proper selection of the externally connected capacitance C and resistance R_2, sinusoidal oscillation can be obtained from the circuit of Figure 2. Moreover, if $C_T >> C_y$ and $R_y >> R_2$, as is practically the case, then Equations (2) and (3) reduce to

$$\omega_o^2 = \frac{1}{C_T C_y R_x R_2} \quad (4)$$

$$R_2 \cong R_x \quad (5)$$

Inspection of Equations (4) and (5), shows that the frequency of oscillation can be controlled by adjusting the external capacitance C without disturbing the condition of oscillation. However, the

condition of oscillation cannot be controlled without disturbing the frequency of oscillation.

Experimental Results

The circuit of Figure 2 was tested using the AD844 configured as a CCII+. With additional built-in voltage buffer, the AD844 provides a low-output impedance sinusoidal signal for further processing. The results obtained with $R1 = 4.74\Omega$ and $R2 = 63.0\Omega$ are shown in **Figure 3** for different values of the capacitance C. In all cases the DC supply of the AD844 was $\pm 10V$ and the amplitude of the sinusoidal output voltage was around 7.7V. It appears from Figure 3 that, by exploiting the parasitic of the CCII+, the circuit of Figure 2 can provide sinusoidal oscillations with frequencies of the order of few MHz.

Muhammad Taher Abuelma'atti
and Munir Ahmad Al-Absi
King Fahd University of Petroleum and Minerals
Saudi Arabia

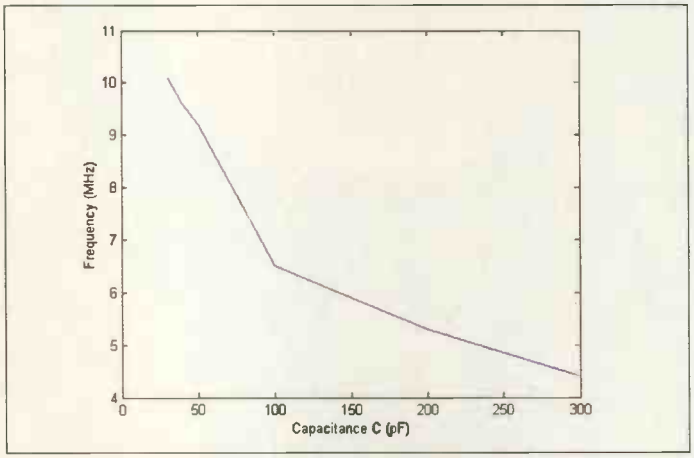

Figure 3: Variation of the frequency of oscillation with the capacitance C

$R_1 = 4.75\Omega$, $R_2 = 63\Omega$

Single Comparator Performs V-to-F Conversion

For transmission paths that are long and noisy, it is recommendable to convert an analogue information signal to a pulse train so that it can be sent and decoded more accurately. This process involves a voltage-to-frequency conversion (VFC). For applications where absolute accuracy is not required, the VFC shown in **Figure 1** offers a simple and low-cost alternative.

The circuit works as follows. Assuming that the input (Vi) is a constant dc, the circuit behaves essentially as an astable multivibrator. Frequency is determined by R, C, and the threshold voltages at the non-inverting input (V+). In the first case, if Vi is higher than the voltage at the inverting input (V-), then output (Vo) is high and the diode is reverse-biased. This allows C to charge via R while V+ is equal to Vi. At the point when the rising voltage at C exceeds V+, Vo goes low.

In this case, where V- is now higher than V+, the diode is forward-biased. This forces V+ to assume a voltage that is 0.7V above ground potential. Consequently C discharges via R until its voltage (=V-) falls below 0.7V (=V+). At that point, Vo snaps back to high and the cycle repeats. **Figure 2** illustrates the timing diagram for the circuit.

For an ideal op-amp and diode, it can be shown that the high and low times of the output are found using the following formula:

$$T_H = RC \ln\left(\frac{V_{CC} - 0.7}{V_{CC} - V_i}\right) \quad (1)$$

$$T_L = RC \ln\left(\frac{V_i}{0.7}\right) \quad (2)$$

where Vi is greater than 0.7 but less than Vcc. Then the period is given by $T=T_L+T_H$ and the output frequency is $f=1/T$. Note that f is determined by Vi. In fact, by trying out numerous values of Vi, it is found that f is inversely related to Vi. For the circuit in Figure 1, the output frequency varies from 10kHz to 40kHz when Vi is adjusted from 7.00V to 2.00V.

Arthur E. Edang
Don Bosco Technical College
Philippines

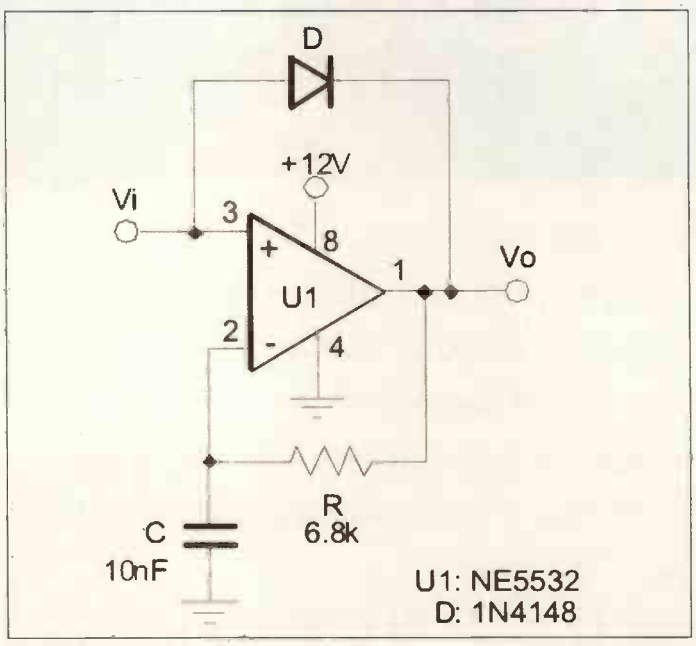

Figure 1

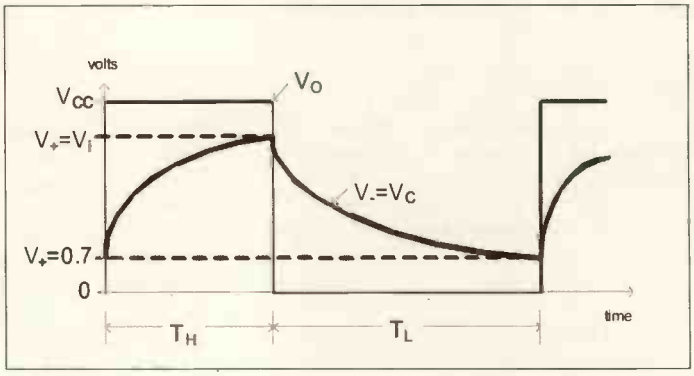

Figure 2

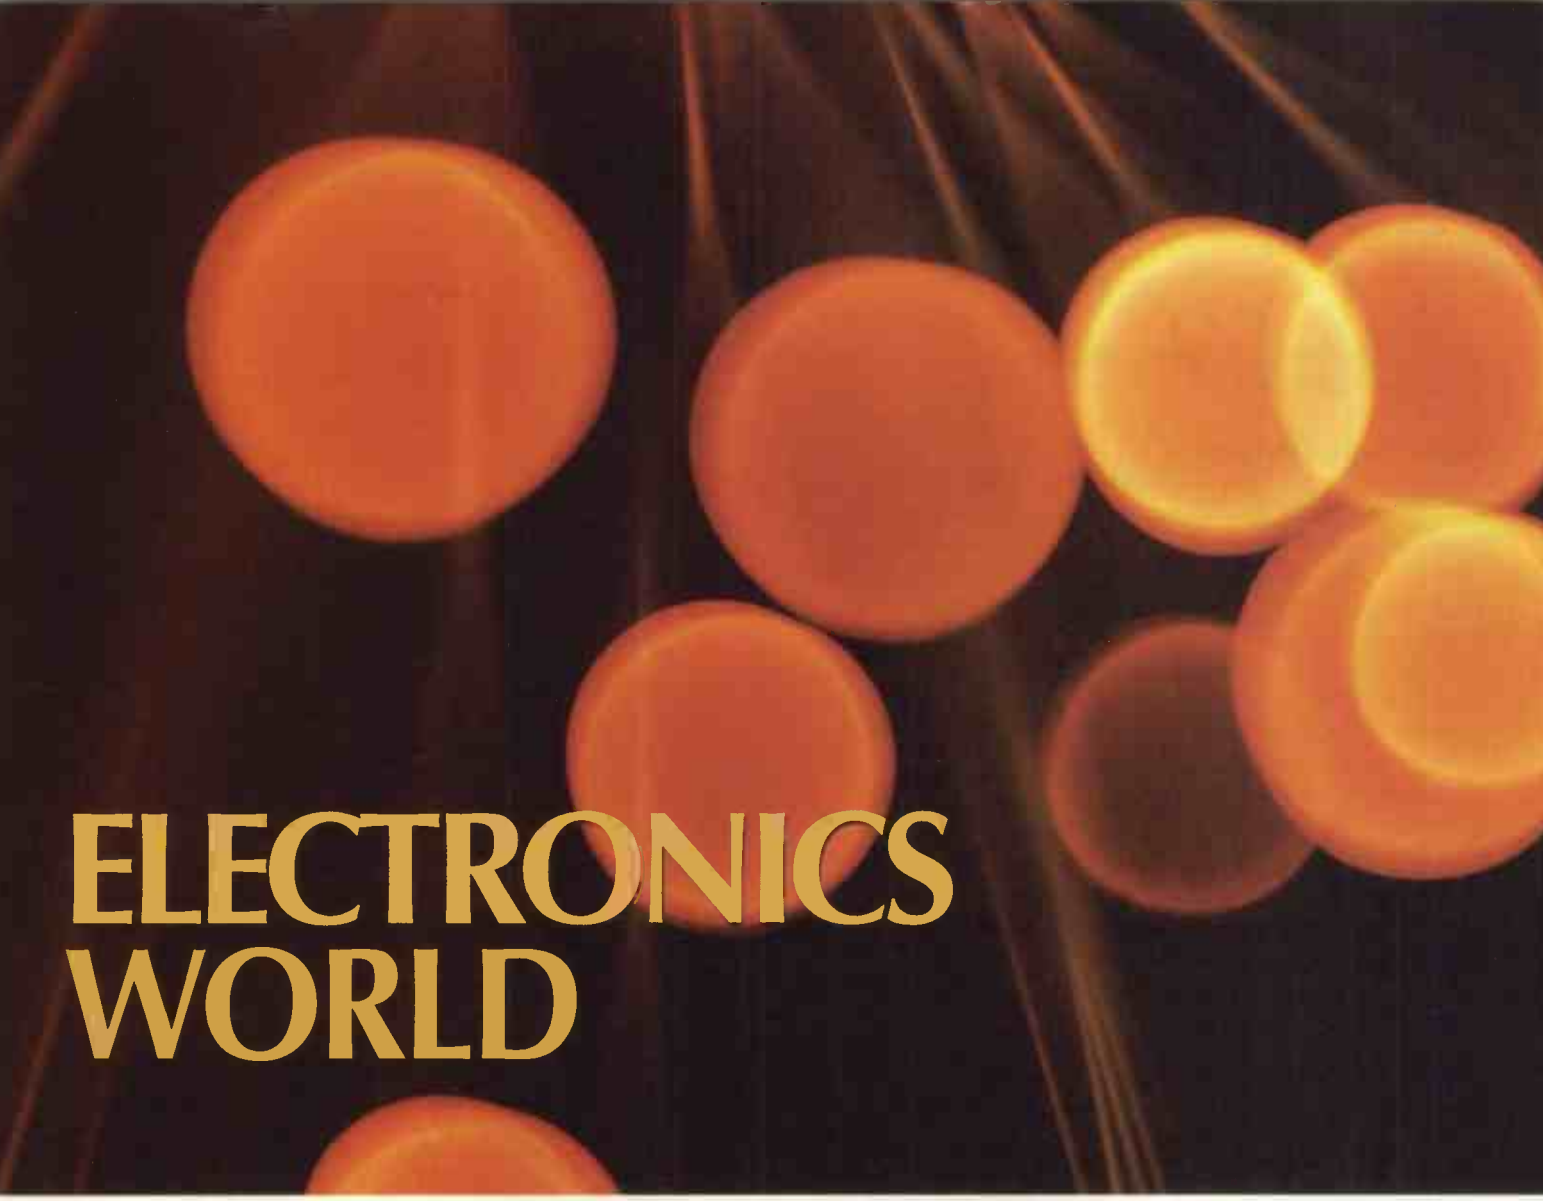

If you are interested in staying at the cutting-edge of electronic design then Electronics World is the magazine for you. Sign up and be inspired every month. A subscription guarantees your magazine arrives every month promptly and without fail.

Subscribe now!

TO MAKE SURE WE KEEP OUR EDITORIAL RELEVANT TO YOU, PLEASE TICK ANY OF THE FOLLOWING THAT APPLY TO YOU. THANK YOU.

1.
- A Professional engineer
- B Serious hobbyist

2. Your profession:
- A Management
- B Design
- C Purchasing/specifying
- D Service/maintenance
- E Research
- F Education

3. Your interest
- A Analogue
- B Digital
- C Software
- D Project management
- E EMC
- F RF/microwave

4. Your industry
- A Aerospace
- B Automotive
- C Broadcast
- D Computing
- E Communications
- F Consumer electronics
- G Industrial
- H Manufacturing
- I Medical
- J Military
- K Power
- L Semiconductors
- M T&M

If you prefer not to receive promotional mailing from other companies please tick box

All subscriptions run for their full term no credit will be given for cancellations during this period

Yes I would like to subscribe to Electronics World magazine for:
- One year £45 UK, £76 in Europe (€120) and Worldwide (US$145)
- Two years £80 UK, £140 in Europe (€215) and Worldwide (US$260) saving 10%
- Three years £105 UK, £185 Europe (€285) and Worldwide (US$345) saving 20%

Name

Job Title

Company

Address

Postcode

Tel

Fax

Email

PLEASE NOTE THAT BY GIVING YOUR EMAIL YOU ARE GIVING US PERMISSION TO CONTACT YOU WITH IT

PAYMENT CHOICE:

- I enclose a cheque made payable to Electronics World
- Please invoice my company
- Please debit my card Visa Mastercard Amex Switch Solo

Card No

Issue No (Switch/Solo only) Expiry date

Signature Date

Please complete payment and address details and return this order form to:

Ref:404

Debbie Jenner, Freepost SEA 11436, Electronics World, Media House, Swanley, Kent BR8 8BR (no stamp required)UK only

International address: Debbie Jenner, Electronics World, Nexus Media Communications, Media House, Swanley, Kent BR8 8HU United Kingdom

The Indago is a satellite navigation system with a difference; several features keep it streets ahead of other units on the market. A top-end product aimed at those who want superior build and design quality, it uses Navteq digital mapping, the most accurate software available. It is the only system to use the comprehensive Enigma database, giving advanced audio and visual warnings of UK and European accident problem spots, speed cameras, speed limits and schools. It also offers over 1,000,000 'points of interest' – from cash machines and petrol stations to tourist attractions and shopping centres – so you'll always know what's around the corner.

Around £599.95
www.indagosatnav.co.uk

With the Electronic Wine Warmer and Chiller, you can get your wine to the perfect temperature within minutes. This snazzy gadget doesn't simply warm all reds to one temperature and cool all whites to another, however, you select the precise type of grape you are drinking to take it to the exact temperature at which it tastes its best. The subtle differences in the grapes mean that they are at their optimum taste at slightly varying temperatures, so it uses Peltier Thermal technology to warm or chill your wine to the exact best temperature. It also maintains this temperature whilst you drink. The gadget needs no messy ice and doesn't use a noisy compressor, it is just quick, clean and simple to use.

Around £69.99
www.drinkstuff.com
or 08707 602 182

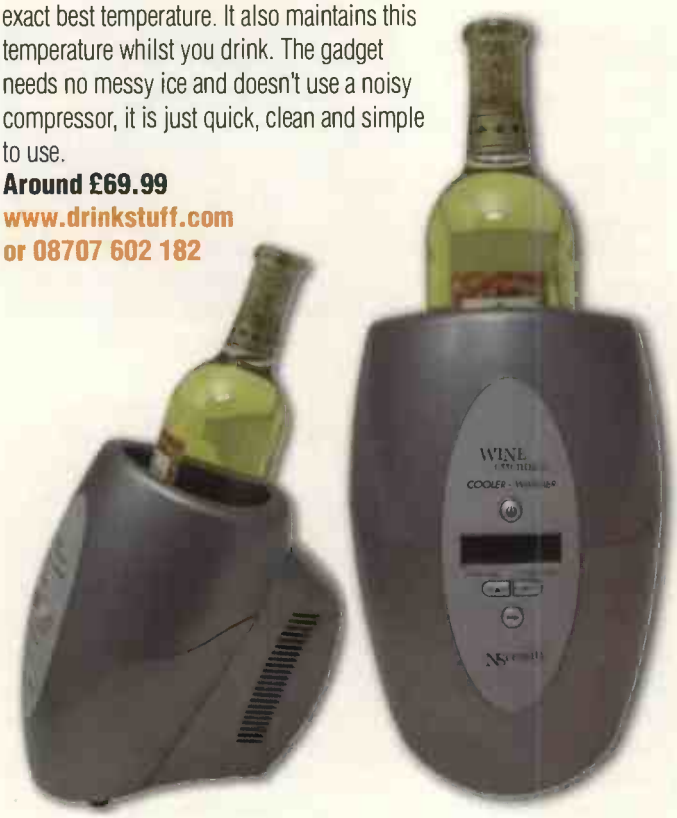

MagneticTime MT1 is a piece of software for iPods, PCs, PDAs and mobiles which allows you to listen to emails and Word documents in MP3 format while travelling, commuting or even exercising. Outlook, Outlook Express and Word documents can be accessed audibly since they are automatically converted to MP3 files. These audio files are created to take advantage of the standard functions of any MP3 player, so emails and documents can be played, paused, stopped, skipped back and forward, and filed in Playlists, etc.

Around £23
www.magnetictime.com

Netgear is making the world's first Skype wireless mobile telephone. The Wi-Fi phone will work wherever you are connected to a wireless Internet access point – be that in a home, office, cafe, or public hotspot etc. This makes headphones or USB phones plugged into a computer obsolete. Pre-loaded with Skype's software, all you need to do after turning on the phone is enter a Skype username and password. The Skype software pulls up your full contact list, displays the connection status on the phone screen and allows you to connect to any other Skype user for free. The phone will also allow you to connect to non-Skype users with the SkypeOut feature.

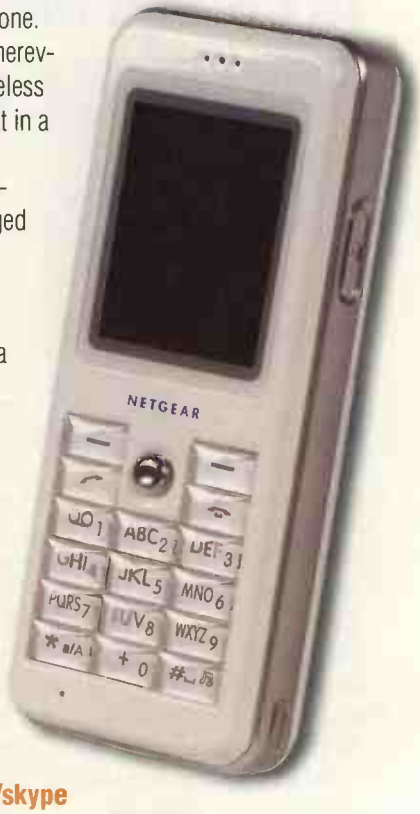

Price not yet known
http://tools.netgear.com/skype

Thermal imager offers high-res at low cost

Irisys has launched the IRI4010, a thermal imager aimed at the preventative maintenance market but at a lower entry-level cost. Its pixel resolution is 160 x 120 and measuring range of between -10°C and 250°C.

This portable device has an integrated colour LCD screen with an LED backlight. It has the capacity to store up to 1000 high quality images through an SD card.

Two measurement cursors highlight temperature and temperature differences, while a Class II laser focuses on the central measurement area for ease of identification of specific hot spots.

The imager has a 200 x 150 field of view lens and a maximum focus of 30cm.

The price of the unit costs £3500.

www.irisys.co.uk

Carrier-class synchronisation chip for ATCA and MicroTCA

Single-chip, ultra-low jitter synchroniser that solves the timing challenges posed by the popular AdvancedTCA (telecommunications computing architecture), AMC (advanced mezzanine card) and MicroTCA architectures is now available from Zarlink Semiconductor.

While easy to use and adaptable, AMCs do not provide redundant timing reference inputs to support carrier-grade timing. The holdover capability of the ZL30117 chip enables it to ride out the complete loss of its incoming reference, which can occur when switching from a failed clock unit to a back-up clock unit. The ZL30117 PLL continues to operate in full compliance with network requirements for several seconds after losing its reference, allowing time for the system to provide another reference source to the AMC.

The ZL30117 device accepts three reference inputs, supporting clock frequencies in any multiple of 8kHz up to 77.76MHz, as well as supporting 2kHz. The ZL30117 chip can directly lock to any of the standard clock input frequencies available to an AMC.

www.products.zarlink.com/profiles/

Body control MCUs for automotive uses

NEC Electronics announced a new line of microcontrollers (MCUs) – the Fx2 series of 8-bit MCUs and Fx3 series of 32-bit MCUs. The family is based on the company's 0.15 micron process technology with embedded flash memory, optimised for automotive body and safety control applications, for systems including air conditioning, windscreen wipers control, power windows, sliding doors, seat controls, smart keyless entry systems, anti-theft mechanisms and intelligent airbags.

The new product line-up is comprised of 29 devices: 13 based on the 78K0TM 8-bit CPU core and 16 devices based on the V850ESTM 32-bit CPU core. The devices are available in pin counts ranging from 44 to 176, and memory configurations from 32kB up to 1MB Flash.

The 78K0 series uses SuperFlash technology by Silicon Storage Technology; all products support the CAN and LIN protocols.

The products in the Fx3 series of MCUs are pin-compatible with NEC Electronics's previous-generation Fx2 series of MCUs.

www.necel.com

Two new dsPIC devices

The dsPIC30F5015 and dsPIC30F5016 are the latest two additions to Microchip's series of 16-bit dsPIC Digital Signal Controllers (DSCs). They feature an advanced PWM, designed for motor control, power conversion and lighting applications; a 1MSPS 10-bit A/D converter; 66kB of Flash memory and an operation of 30MIPS using an internal oscillator. These new devices are ideal for applications that drive power FETs and require advanced algorithmic processing.

The dsPIC30F5015 and dsPIC30F5016 will operate from 2.5 to 5.5 volts, which is valuable for analogue noise immunity or minimising voltage-translation logic. The devices are available to operate over an extended temperature range of -40°C to +125°C.

Additional key features include: 2kB of SRAM; 1kB of on-chip EEPROM; 8-output advanced PWM; 4 duty-cycle generators; 10-bit analogue-to-digital converter with up to 16 signal channels; 4 channel simultaneous sampling and PWM trigger option; quadrature encoder interface for motor control applications; five 16-bit timers and CAN, SPI, I2C and UART peripherals.

www.microchip.com/dspic

Omeg Custom Built Switched Potentiometers

PROTOTYPE SERVICE TO FULL PRODUCTION QUANTITIES

Terminations
Choice of solder tags or pcb pins.

Switch Action
Choice of rotary or push on / push off.

Rotary Switch Ratings
1 or 4 amp, 250 volts, single or two pole, on/off or changeover.

Push Switch Ratings
10 amp, 250 volts, single pole.

Rotary Switch Actuation
Choice of switch actuation at either end of the rotational arc.

See www.omeg.com

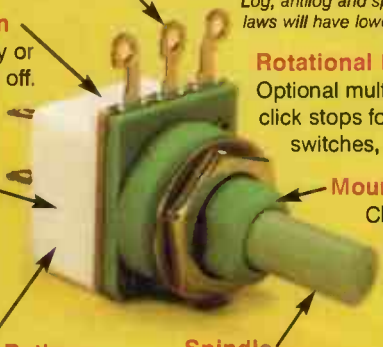

Potentiometer Power Ratings
0.25W – 16 mm pots (linear laws)
0.40W – 20 mm pots (linear laws)
Log, antilog and special non-linear laws will have lower power ratings.

Rotational Feel
Optional multi-position click stops for push switches, if required.

Mounting Bush
Choice of glass filled nylon or diecast zinc alloy.

Spindle
Choice of 4 mm, 6 mm or ¼" diameter, various lengths, with or without flats.

Resistance Ratings
Linear resistance laws : 1k – 1MΩ
Non-linear laws : 4k7 – 470kΩ

OMEG LIMITED IMBERHORNE INDUSTRIAL ESTATE, EAST GRINSTEAD, WEST SUSSEX, RH19 1RJ
Tel: 01342 410420 Fax: 01342 316253 Email: sales@omeg.com

Rechargeable batteries with solder tags.

NIMH		NICAD	
AA 2000mah	£2.82	AA 650mah	£1.41
C4Ah	£4.70	C 2.5A	£3.60
D9Ah	£7.60	D 4Ah	£4.95
PP3 150mah	£4.95		

Instrument case with edge connector and screw terminals

Size 112mm x 52mm x 105mm tall.
This box consists of a cream base with a PCB slot, a cover plate to protect your circuit, a black lid with a 12 way edge connector and 12 screw terminals built in (8mm pitch) and 2 screws to hold the lid on. The cream bases have minor marks from dust and handling.
Price £2.00 + VAT (= £2.35) for a sample or £44.00 + VAT (= £51.70) for a box.

866 battery pack originally intended to be used with an orbitel mobile telephone it contains 10 1.6Ah sub C batteries (42x22dia the size usually used in cordless screwdrivers etc.) the pack is new and unused and can be broken open quite easily £6.46 + VAT = £8.77

Please add £1.66 + VAT = £1.95 postage and packing per order.

JPG ELECTRONICS
Shaws Row, Old Road, Chesterfield S40 2RB
Tel: 01246 211202 Fax: 01246 550959
www.jpgelectronics.com
Mastercard/Visa/Switch
Callers welcome 9.30am to 5.30pm Monday to Saturday

High Voltage and RF Relays from Crydom

- Liquid Level Sensors
- Flow Sensors
- Reed Relays

crydom
Crydom Magnetics Ltd
7 Cobham Road
Ferndown Industrial Estate
Wimborne, Dorset,
BH21 7PE, United Kingdom
Tel: +44 (0) 1202 897969
Fax: +44 (0) 1202 891918
magnetics@crydom.com

KRP POWER SOURCE

MINIATURE POWER SUPPLIES 5-100 WATTS

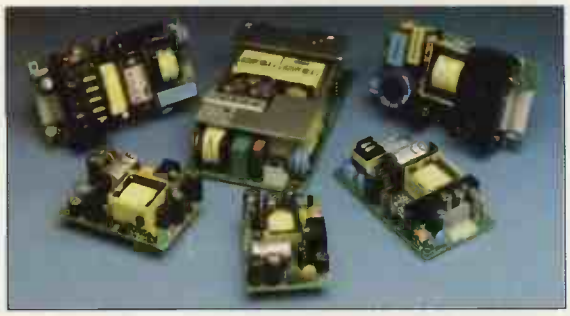

The Cincon CFM series of single output open card power supply modules provide a major breakthrough in the power to cost ratios. Available in 7 power sizes of 5, 10, 15, 20, 40, 60 and 100 watts, the largest unit has a plan area of only 127mm x 76.2mm and an exceptionally low height of 33.6mm. 40 watt dual and triple output models available.

- Universal input: 85–264 Vac
- High efficiency
- Fully regulated output: 3.3 Vdc – 48 Vdc
- PCB mounting/chassis mounting
- UL, Ulc, TUV and CE Approved
- Low profile, high MTBF

NEW 100 WATT MODEL

KRP POWER SOURCE U.K. LTD
Tel: (01635) 32585 • Fax: (01635) 32510
Email: sales@krp.co.uk • www.krp-power.co.uk

Website Directory

To reserve your website space phone Reuben Gurunlian 01322 611261 or email reuben.gurunlian@nexusmedia.com

Action Circuits

www.action-circuits.co.uk

Specialists in Chip Programming and SMD Reeling Services, Action Circuits (UK) Ltd offer subcontract services to the electronics industry second to none. We currently have trading partnerships with many Semiconductor manufacturers, distributors, CEM's and OEM's. Our component programming and tape & reeling service is available for high volume or sample quantities.

If you are looking for a service to tape and reel components or program your OTP or Flash programmable devices, then our highly experienced staff are committed to offering you a device programming service, laser marking and component taping service with a quick turnaround at a very competitive cost.

Conford Electronics

www.confordelec.co.uk

Lightweight portable battery/mains audio units offering the highest technical performance. Microphone, Phantom Power and Headphone Amplifiers. Balanced/unbalanced signal lines with extensive RFI protection. Specialists

Designer Systems Co

www.designersystems.co.uk

Electronic product design company with over a decade of experience promoting it's own product range and designing and manufacturing innovative products for client companies/individuals.

DB Technology

www.dbtechnology.co.uk/

Anechoic chamber and open area test site.

- Compliance Tests
- Fixes included. FCC Listed.
- Flexible hourly booking available.
- Rapid, accurate pre-compliance tests.

Forest Electronics

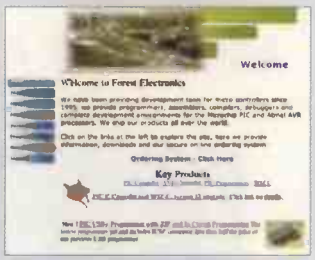

www.fored.co.uk

Electronic product design company with over a decade of experience promoting it's own product range and designing and manufacturing innovative products for client companies/individuals.

Harmonic Software Systems

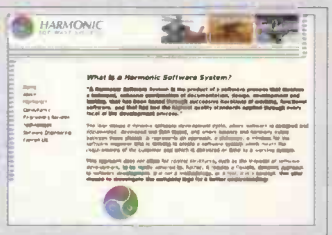

www.harmonics.co.uk/harmonic.htm

Harmonic Software Systems develops embedded real-time systems and DSP solutions for industrial, commercial and military customers. We supply:

- Vx Works Board Support Packages and device drivers. • Software Development services. • Consultancy at excellent rates.
- Safety critical systems up to SIL3.
- Digital Signal Processing algorithms and kernel development.

We are experienced in developing systems utilising a wide range of technologies, including

- MIL-STD-1553 • Firewire IEEE-1394
- Ethernet • CANBus • V RS232/422/485
- ARM, MIPS, PowerPC, x86 • TMS320, Sharc
- VME, cPCI, PC104 or custom hardware

Call us now on 01293 817635 to discuss your requirements.

Kalestead Ltd

www.kalestead.co.uk/cable.htm

At Kalestead we do not simply supply coiled leads from stock or "off the shelf", the reason for this being that they are made to each customer's specific requirements. We can, however, offer a quick turnaround using cable available from stock which can be coiled to you own specification.

It is because of the specialised nature of our business that we are unable to provide a price list on our web site. Please do not hesitate to contact a member of our sales team who will discuss with your and supply you with an individual quotation.

60

Mushroom Components

www.surpletronics.com

Surpletronics is a brand new electronics components e-Store with a difference. We specialise in popular electronic components from leading manufacturers at rock bottom prices that you won't find elsewhere.

New lines are arriving every day and many of our products are one-off specials that may not be repeated.

Our products are all new, unused and are fully guaranteed. Ordering is simple, with UK shipment free for orders over £50 and payment is completed credit card. Goods are shipped that day for prompt delivery.

Surpletronics - serious about surplus. Tel: 0870 850 5258

Colour screen shot and text will cost just £350 + vat for 12 months

www.electronicsworld.co.uk

Website Directory

To reserve your website space phone Reuben Gurunlian 01322 611261 or email reuben.gurunlian@nexusmedia.com

Newbury Electronics

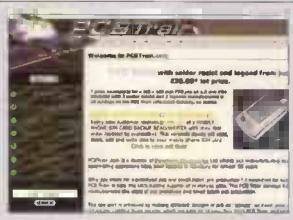

www..pcbtrain.com

PCBT Train is a division of Newbury Electronics Ltd whose pcb manufacturing and assembling operations have been located in in Newbury for almost 50 years.

Why pay more for a prototype pcb and small batch pcb production? Established for over 5 years, PCB Train is now the UK's leading supplier of prototype pcbs. The PCB Train concept has revolutionised the costs of pcb prototype and small batch pcb production.

The low cost is achieved by making different designs of pcb on "panels" on fixed production schedules yielding huge savings which we pass on to you. Buy from PCB Train and receive professionally UK manufactured (BS EN ISI 9001:2000 approved) prototype pcbs and small batch production runs of pcbs on time at prices a fraction of those typically quoted by our competitors.

Picotech

www.picotech.com

Established in 1991, Pico Technology is a worldwide leader in the design, development and manufacture of high performance PC Oscilloscopes and Data Loggers. During this time Pico has built up an impressive portfolio of products; including the PicoScope 3000 series of PC Oscilloscopes, the USB TC-08 Thermocouple Data Logger and the award-winning Automotive Diagnostics Kit.

We pride ourselves on offering innovative, high quality and affordable alternatives to traditional bench top test and measurement equipment, operating under the ISO9001:2000 quality system.

RD Research

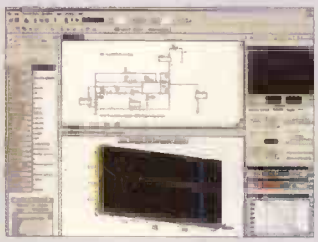

www.spice-software.com

B2 Spice Version 5.2

Low cost - high performance Spice software. Features only normally found in packages costing more than £1,000. We are so confident that B2 Spice V5.2 will impress you that we offer an unconditional 30-day evaluation of the full boxed version, complete with printed user manual. If it's not up to your expectations then simply return it.

• Fast and accurate simulations • Real-world parts • Circuit Wizard • Function generators, oscilloscope, ammeter, voltmeter, wattmeter, distortion meters, and more. • Real-time 'live' simulation adjustments.

All the power without the cost. Only £229 + VAT.

Call 01603 872331. www.spice-software.com

Telnet Ltd

www.telnet.uk.com

Suppliers of quality second-user test and measurement equipment at prices you can afford. Manuals and accessories supplied. If you would like a quote, please call. We also purchase your surplus test equipment. Please call us for the best offers.

TELNET
1 Stoney Court, Hotchkiss Way, Binley Industrial Estate, CV3 2RL
Tel: 024 76650702
Fax: 024 76650773
Mobile: 07860 400683
Email: dales@telnet.uk.com

Beta Layout Ltd

www.pcb-pool.com

The Best Value Online PCB Prototyping Service Available:

Instant on line quotations & ordering (no pre registration).

We offer the following:

• No minimum quantity. • No tooling charges. • No drill Limitations. • 1 to 6 layers (prototype quantities) • 2 layers (small production batches) • Fr4,1.6mm,35 um,HASL (Pb free/Rohs/Weee) • Soldermask / Silkscreen (optional) • Leadtimes from 2 -8 working days • Full DRC on all orders (we manually check every file !!!). • PCB-POOL® accepts direct outputs from 15 major layout softwares *(see our website for details)*

Download our fully functional PCB LAYOUT software FREE of charge.

Free Phone : 0800 3898560

Stewart Of Reading

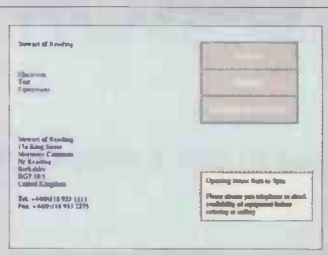

www.stewart-of-reading.co.uk

Stewart of Reading is a small business with over 30 years experience in the surplus electronic test equipment market.

We are well known within the industry, have a good reputation and successfully distribute stock worldwide.

All equipment is fully tested and guaranteed before distribution.

Stewart of Reading
17A King Street, Mortimer, Nr. Reading, Berkshire. RG7 3RS United Kingdom
Tel: +44(0)118 933 1111
Fax: +44(0)118 933 2375

61

Colour screen shot and text will cost just £350 + vat for 12 months

www.electronicsworld.co.uk

SPECIAL OFFER – 6 for the price of 5

ELECTRONICS WORLD CD-ROM 1999 – 2004 £30 each inc p&p (UK only)

Add £1 postage for Europe or £5 rest of the world

• easy to use • easy to browse • full text and diagrams of all articles, circuit ideas, letters etc •

Please send the following CD-ROMs:

- ❑ 1999
- ❑ 2000
- ❑ 2001
- ❑ 2002
- ❑ 2003
- ❑ 2004
- ❑ CDs @ £30 each = £.....
  add postage £.....
  Total £.....

❑ I enclose a cheque payable to Electronics World

Charge my: ❑ Visa ❑ Mastercard ❑ Amex ❑ Switch ❑ Delta Issue No. (Switch/Delta) ❑❑

❑❑❑❑❑❑❑❑❑❑❑❑❑❑❑❑ Expiry date:.....

Name:.....

Company name:.....

Address:.....

.....Postcode:.....

Telephone Number:.....Fax Number.....

Email:.....

Post to: Electronics World, Nexus Media Communications, Media House, Swanley, Kent BR8 8BR

Quasar Electronics Limited
PO Box 6935, Bishops Stortford
CM23 4WP, United Kingdom
Tel: 0870 246 1826
Fax: 0870 460 1045
E-mail: sales@quasarelectronics.com
Web: www.QuasarElectronics.com

Postage & Packing Options (Up to 2Kg gross weight): UK Standard 3-7 Day Delivery - £3.95; UK Mainland Next Day Delivery - £8.95; Europe (EU) - £6.95; Rest of World - £9.95 (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 17.5% VAT.
Call now for our FREE CATALOGUE with details of over 300 kits, projects, modules and publications. Discounts for bulk quantities.

QUASAR electronics
Get Plugged In!

0871 CREDIT CARD SALES 717 7168

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 details.

NEW! Bidirectional DC Motor Controller

Controls the speed of most common DC motors (rated up to 32Vdc/5A) 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: 3166KT - **£15.95**
Assembled Order Code: AS3166 - **£25.95**

DC Motor Speed Controller (5A/100V)

Control the speed of almost any common DC motor rated up to 100V/5A. Pulse width modulation output for maximum motor torque at all speeds. Supply: 5-15Vdc. Box supplied. Dimensions (mm): 60Wx100Lx60H.
Kit Order Code: 3067KT - **£12.95**
Assembled Order Code: AS3067 - **£20.95**

NEW! PC / Standalone Unipolar Stepper Motor Driver

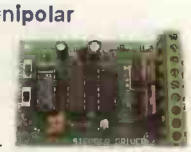

Drives any 5, 6 or 8-lead unipolar stepper motor rated up to 6 Amps max. Provides speed and direction control. Operates in stand-alone or PC-controlled mode. Up to six 3179 driver boards can be connected to a single parallel port. Supply: 9Vdc. PCB: 80x50mm.
Kit Order Code: 3179KT - **£11.95**
Assembled Order Code: AS3179 - **£19.95**

NEW! Bi-Polar Stepper Motor Driver

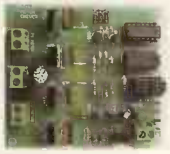

Drive any bi-polar stepper motor using externally supplied 5V levels for stepping and direction control. These usually come from software running on a computer. Supply: 8-30Vdc. PCB: 75x85mm.
Kit Order Code: 3158KT - **£14.95**
Assembled Order Code: AS3158 - **£29.95**

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 PSU345 £9.95

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 - **£39.95**
Assembled Order Code: AS3180 - **£47.95**

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 38x38mm. Powered by PC. Includes one DS1820 sensor and four header cables.
Kit Order Code: 3145KT - **£17.95**
Assembled Order Code: AS3145 - **£23.95**
Additional DS1820 Sensors - **£3.95 each**

NEW! DTMF Telephone Relay Switcher

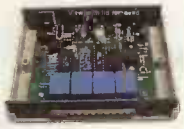

Call your phone number using a DTMF phone from anywhere in the world and remotely turn on/off any of the 4 relays as desired. User settable Security Password, Anti-Tamper, Rings to Answer, Auto Hang-up and Lockout. Includes plastic case. Not BT approved. 130x110x30mm. Power: 12Vdc.
Kit Order Code: 3140KT - **£44.95**
Assembled Order Code: AS3140 - **£64.95**

Serial Isolated I/O Module

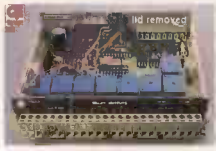

Computer controlled 8-channel relay board. 5A mains rated relay outputs. 4 isolated digital inputs. Useful in a variety of control and sensing applications. Controlled via serial port for programming (using our new Windows interface, terminal emulator or batch files). Includes plastic case 130x100x30mm. Power Supply: 12Vdc/500mA.
Kit Order Code: 3108KT - **£54.95**
Assembled Order Code: AS3108 - **£64.95**

Infrared RC Relay Board

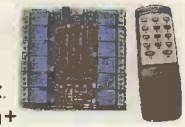

Individually control 12 on-board relays with included infrared remote control unit. Toggle or momentary. 15m+ range. 112x122mm. Supply: 12Vdc/0.5A
Kit Order Code: 3142KT - **£42.95**
Assembled Order Code: AS3142 - **£52.95**

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) £15.00
18Vdc Power supply (PSU010) £19.95
Leads: Parallel (LDC136) £4.95 / Serial (LDC441) £4.95 / USB (LDC644) £2.95

NEW! 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 Plug A-B lead not incl.
Assembled Order Code: AS3128 - **£39.95**

"PICALL" PIC Programmer

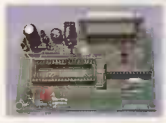

"PICALL" will program virtually all 8 to 40 pin *serial-mode** AND *parallel-mode* (PIC16C5x family)* programmed PIC micro controllers. Free fully functional software. Blank chip auto detect for super fast bulk programming. Parallel port connection. Supply: 16-18Vdc.
Assembled Order Code: AS3117 - **£24.95**

ATMEL 89xxxx Programmer

Uses serial port and any standard terminal comms program. 4 LED's display the status. ZIF sockets not included. Supply: 16-18Vdc.
Kit Order Code: 3123KT - **£24.95**

NEW! USB & Serial Port PIC Programmer

USB/Serial connection. Header cable for ICSP. Free Windows software. Wide range of supported PICs - see website for complete listing. ZIF Socket/USB Plug A-B lead not incl. Supply: 16-18Vdc.
Kit Order Code: 3149EKT - **£37.95**
Assembled Order Code: AS3149E - **£52.95**

No.1 for KITS

www.quasarelectronics.com

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

Products & Services

ARTICLES WANTED

TOP PRICES PAID
For all your valves, tubes, semi conductors and ICs.
Langrex Supplies Limited
1 Mayo Road, Croydon, Surrey CR0 2QP
TEL: 020 8684 1166 FAX: 020 8684 3056

EXCESS STOCK

SURPLUS WANTED
WE BUY: IC's, Memory, Relays, Caps, PSU's, Semiconductors, Populated Boards, Computers & Test Equipment
ANYTHING CONSIDERED
For our wide range of Semiconductor + Passives List, please ring, fax or e-mail
MAIL ELECTRONICS LTD
TEL: 0161-761 4520 / FAX: 0161-763 6863
EMAIL: andrew@mailelectronics.com
Web: www.mailelectronics.com

POWER

TOROIDAL TRANSFORMERS

- Product Range 10va - 6kva
- Short Lead Times
- Large or Short Production Runs
- Very Fast Prototype Service
- All products manufactured in UK

TIGER TOROIDS LTD
1 The Maltings Road,58
Victoria Road, Diss, Norfolk
IP22 3JE
Tel: (01379) 650580
Fax: (01379) 650585

SERVICES

POWER SUPPLY DESIGN
Switched Mode PSU
Power Factor Correction
designed to your specification
Tel/Fax: 01243 842520
e-mail: eugen_kus@cix.co.uk
Lomond Electronic Services

COMPONENTS WANTED

Components wanted
Excess/Surplus
www.mushroom.co.uk
01234 363611
will collect anywhere in Europe
Instant Cash Settlement

TAPE REELING & SERVICES

action circuits (UK) LIMITED
device programming & reeling specialists

SMD TAPING & REELING SERVICES
Tel: 00 44 (0)1582 412323
Email: sales@actioncircuits.com
www.actioncircuits.com

PCB MANUFACTURE

PCB Train
No tooling charge
Prototype PTH PCBs from just £30.00
Prices published at www.pcbtrain.com
The leading low-cost source for Prototype PCBs from 1-6 layers
+44 (0) 1635 40347
f +44 (0) 1635 36143 e pcbtrain@newbury.tcom.co.uk

DEVICE PROGAMMING SERVICES

action circuits (UK) LIMITED
device programming & reeling specialists

I.C. PROGRAMMING & LASER MARKING SERVICES
Tel: 00 44 (0)1582 412323
Email: sales@actioncircuits.com
www.actioncircuits.com

FOR SALE

PRINTED CIRCUIT BOARDS
DESIGNED & MANUFACTURED
agar Circuits

- Prototype or production quantities
- Fast turnround available
- PCBs designed from circuit diagrams
- Almost all computer files accepted
- PCB assembly - mechanical assembly
- Full product design-manufacture-test-repair

Unit 5, East Belfast Enterprise Park
308 Albertbridge Rd, Belfast BT5 4GX
TEL 028 9073 8897 FAX 028 9073 1802
info@agarcircuits.com

This magazine reaches over 10,000 qualified, potential customers. To advertise please contact Reuben - 01322 611261

SERVICES

PCB-POOL®

- Prototypes at a fraction of the cost
- Tooling and setup included
- Industry standard quality
- Any contour

FREE LAYOUT SOFTWARE!

- Follow up Services Runs
- CAM/CAD Consulting
- ONLINE Quotations
- ISO-Approved

with first order (while stocks last)
Freephone 0800-3898560
Beta LAYOUT
Tel.: +353 (0)61 701170
Fax: +353 (0)61 701164
sales@beta-layout.com
WWW.PCB-POOL.COM

EMBEDDED SOFTWARE DEVELOPMENT

HARMONIC SOFTWARE SYSTEMS
EMBEDDED SOFTWARE DEVELOPMENT

- ✔ Real Time Systems Design, Development & Testing
- ✔ Digital Signal Processing - Tms320, Sharc, Dsp56K
- ✔ Device Drivers, BSPs, Libraries, Firmware
- ✔ Safety Critical to SIL3, Misra-C:2004, POSIX, VxWorks
- ✔ VME, cPCI, PC104 or Custom Hardware

Tel: 01293 817635
Web: www.harmonicss.co.uk
Email: sales@harmonicss.co.uk
Premium quality at a competitive price

SURPLUS COMPONENTS WANTED

new parts @ surplus prices
www.surplectronics.com

PROTEUS

ELECTRONIC DESIGN

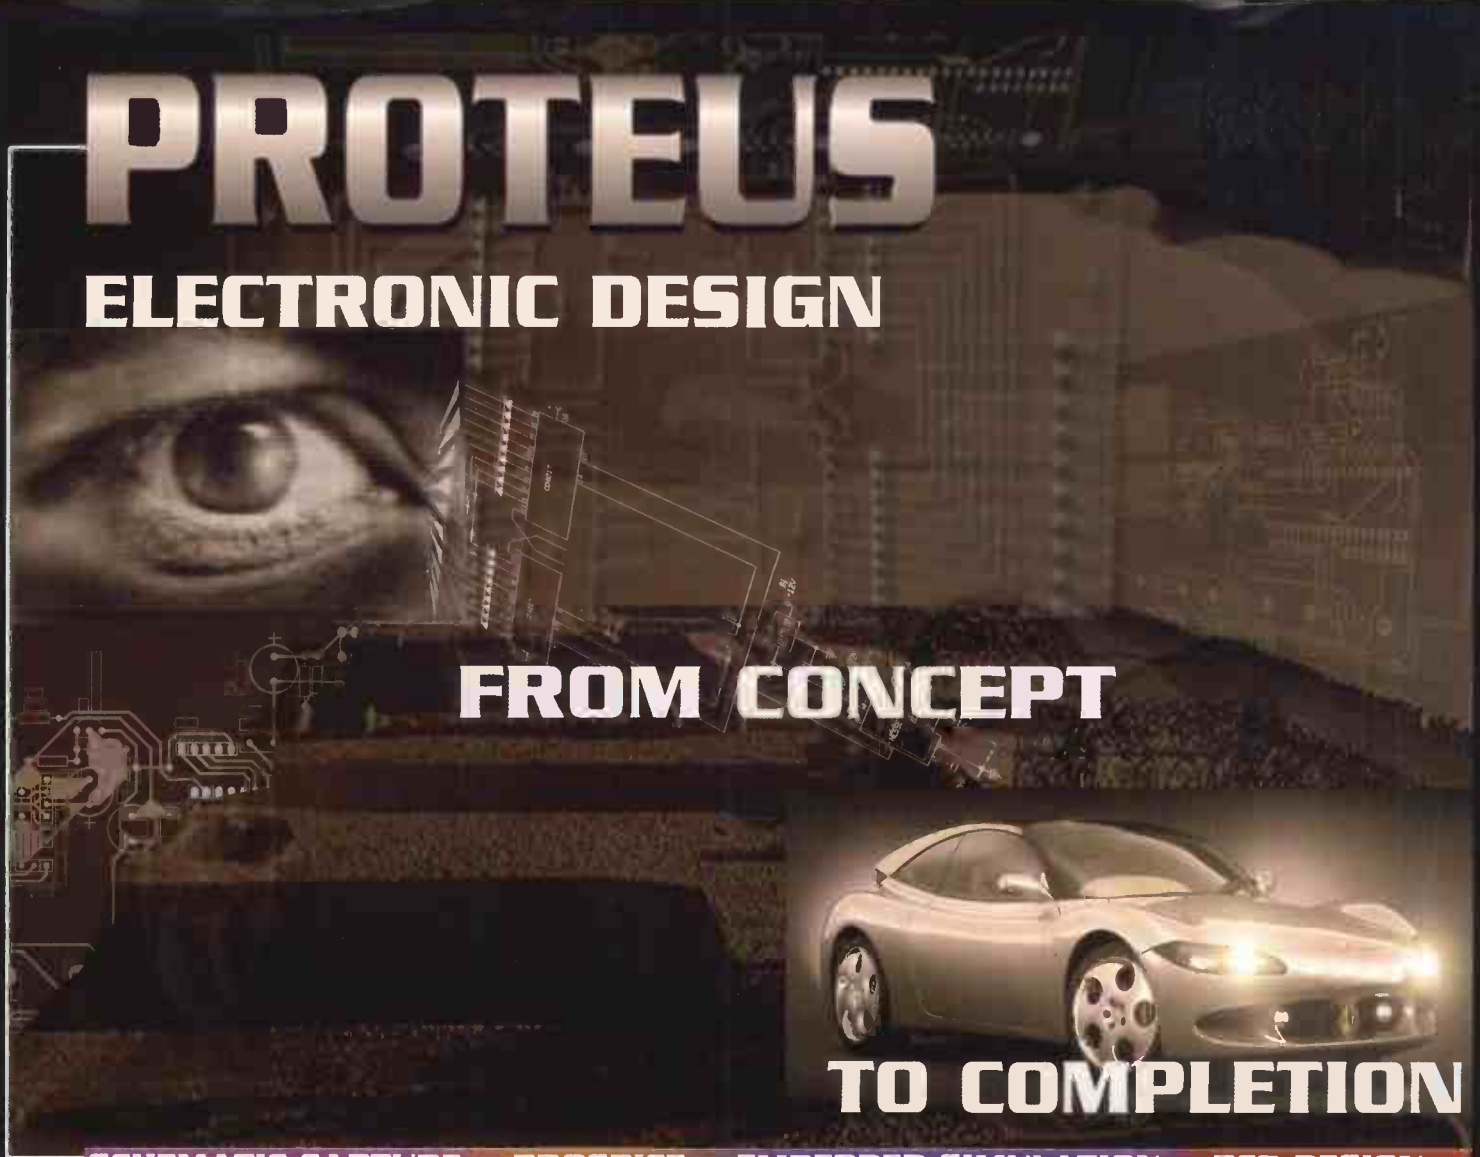

FROM CONCEPT

TO COMPLETION

SCHEMATIC CAPTURE | PROSPICE | EMBEDDED SIMULATION | PCB DESIGN

ISIS SCHEMATIC CAPTURE

A powerful capture package tailored for todays engineer and designed to allow rapid entry of complex schematics for simulation and PCB Layout.

PROSPICE MIXED MODE SIMULATOR

A customised implementation of the industry standard Berkeley SPICE 3F5 engine with extensive optimisations and enhancements for true mixed mode simulation and circuit animation.

VSM VIRTUAL SYSTEM MODELLING

The worlds first and best schematic based microcontroller co-simulation software. Proteus VSM allows you to simulate the interaction between software running on a microcontroller and any analog or digital electronics connected to it. This streamlines the project lifecycle and obviates the need for expensive hardware analysis tools.

ARES PCB DESIGN

A modern and professional layout package which seamlessly integrates with the ISIS capture software. Features such as autoplacement and autorouting, interactive DRC and an intuitive interface all serve to maximise productivity and reduce time to market.

LABCENTER ELECTRONICS LTD.

A technology pioneer in the EDA industry since 1988.
Technical support direct from the program authors.
Flexible packages and pricing tailored to customer requirements.

CONTACT US NOW

to discuss requirements or request a **FREE** evaluation copy.

labcenter Electronics

www.labcenter.co.uk

Tel: 01756 753440
Fax: 01756 752857
Email: info@labcenter.co.uk

Labcenter Electronics Ltd., 53-55 Main Street, Grassington, North Yorks, BD23 5AA. Registered in England 4692454

TEST EQUIPMENT SOLUTIONS

Quality Second User Test Equipment

With 12 Months Warranty

The Industry's Most Competitive Test Equipment Rental Rates

Sample Stock List - If you don't see what you want, please CALL!

	Sale (GBP)	Rent (GBP)
AMPLIFIERS		
AT/HP 8349B 2-20GHz +15dB >50mW Amplifier	2500	100
AT/HP 8447E 1.3GHz Power Amplifier	950	48
AT/HP 8447F 1.3GHz Pre/Power Dual Amplifier	1250	50
Ophir 5022B 1.4-2.4GHz 50W RF Amplifier	8350	334
Ophir 5027 20-1000MHz 40W RF Amplifier	4675	187
Ophir 5106B 1.4-2.4GHz 25W RF Amplifier	7250	290
FREQUENCY COUNTERS		
AT/HP 53131A 225MHz 10 Digit Universal Counter	950	39
AT/HP 53131A/001 DC-225MHz 10 Digit Universal Counter	1150	41
AT/HP 5342A/011/002 24GHz Frequency Counter	1150	42
AT/HP 5370A 100MHz Universal Time Interval Counter	1250	38
AT/HP 5371A 500MHz Frequency/Time Interval Analyser	1650	50
AT/HP 5372A 500MHz Frequency/Time Interval Analyser	2575	78
Marconi 2440 20GHz Microwave Counter	1550	47
FUNCTION GENERATORS		
AT/HP 3312A 13MHz Function Generator	750	29
AT/HP 3325A/001 21MHz Function Generator	850	32
AT/HP 3325B 21MHz Function Generator	950	30
AT/HP 3335A/001 81MHz Function Generator	1450	44
AT/HP 8116A 50MHz Function Generator	1495	45
AT/HP 8165A 50MHz Function Generator	1275	51
AT/HP 8904A/001/002/003 600kHz Function Generator	1950	61
LOGIC ANALYSERS		
AT/HP 16500C Logic Analyser Mainframe	1650	55
AT/HP 16510B 100MHz Timing 35MHz State 80Ch Card	890	35
AT/HP 1652B 100MHz Timing 35MHz State 80Ch with DSO	2150	75
AT/HP 16533A 1GS/s DSO Card For 16500 Series	950	40
AT/HP 16550A 102Ch 500MHz/100MHz Timing/State Mod	1550	51
AT/HP 1660AS 500MHz/100MHz Timing/State 136Ch + DSO	3525	144
AT/HP 1661A 500MHz Timing 100MHz State 102Ch Log Ana	2450	80
AT/HP 1662A 500MHz Timing 100MHz State 68Ch Log Ana	2350	71
AT/HP 1670G 500MHz Timing 150MHz State 136Ch Log Ana	4650	140
NETWORK ANALYSERS		
Advantest R3767BH 40MHz-8GHz Network Analyser	8250	330
Advantest R3767CH 40MHz-8GHz Network Analyser	10950	440
AT/HP 3577A 5Hz-200MHz Vector Network Analyser	3950	119
AT/HP 3589A 150MHz Network/Spectrum Analyser	4950	198
AT/HP 41951A Impedance Test Kit For 4195A	1700	52
AT/HP 4195A 500MHz Vector Network/Spectrum Analyser	6950	209
AT/HP 85107B 45MHz-50GHz Network Analyser System	47250	1926
AT/HP 8510B 50GHz Microwave Network Analyser	4250	170
AT/HP 8712B/1C2 1.3GHz Vector Network Analyser	6250	250
AT/HP 8714C 300kHz-3GHz Vector Network Analyser	6950	209
AT/HP 8752C 1.3GHz Network Analyser	5750	230
AT/HP 8753D/002/010 3GHz Vector Network Analyser	11500	345
AT/HP 8753D/006 6GHz Vector Network Analyser	13750	413
AT/HP 8757E Scalar Network Analyser	4650	140
AT/HP 89441A/AY7/8/9/AYA/AYB/UFG/UG7 DC-2650MHz VSA	18750	565
Anritsu MS4623B/3B/4/11SF/13 6Ghz VNA System	15950	481
Anritsu S251B/10A 2.5GHz Dual Port Scalar Network Ana	5250	158
Anritsu S331C 4GHz Scalar Network Analyser	3950	119
OSCILLOSCOPES		
AT/HP 54200A 2 Channel 50MHz 200MS/s Digitising Scope	850	33
AT/HP 54501A 4 Channel 100MHz 20MS/s Digitising Scope	1100	46
AT/HP 54502A 2 Channel 400MHz 400MS/s Digitising Scope	1100	34
AT/HP 54503A 4 Channel 500MHz 20MS/s Digitising Scope	1695	52
AT/HP 54510B 2 Channel 300MHz Digitising Scope	2650	106
AT/HP 54603B 2 Channel 60MHz 20MS/s Digitising Scope	850	32
AT/HP 54642D 2 Ch 500MHz 2GS/s Scope + 16 Logic Chs	4950	149
AT/HP 54645A 2 Channel 100MHz 200Ms Digital Scope	1500	45
AT/HP 54810A 2 Channel 500MHz 1GS/s Digitising Scope	3950	119
AT/HP 54825A 4 Channel 500MHz 2GS/s Digitising Scope	5250	159
AT/HP 54835A 4 Channel 1GHz 4GS/s Digitising Scope	7250	301
AT/HP 54845A 4 Channel 1.5GHz 8GS/s Infinium Scope	9350	374
Fluke 99 2 Channel 50MHz Handheld Scope	1150	48
Lecroy 9424E 4 Channel 350MHz Digitising Scope	1675	51
Tek 2205 2 Channel 20MHz Analogue Scope	300	24
Tek AM503S/03/A2 Current Probe System (inc.A6302 Probe)	1350	49
Tek CSA7404/4M/JA3 4 Channel 4GHz Comm Signal Analyser	28000	1120
Tek TDS340 2 Channel 100MHz 500MS/s Digitising Scope	1050	32
Tek TDS380/14 2 Channel 400MHz 2GS/s Digitising Scope	2150	69
Tek TDS754D 4 Channel 500MHz 2GS/s Digitising Scope	7250	290
Tek TDS794D/1M 4 Channel 2GHz 4GS/s Digitising Scope	10500	315

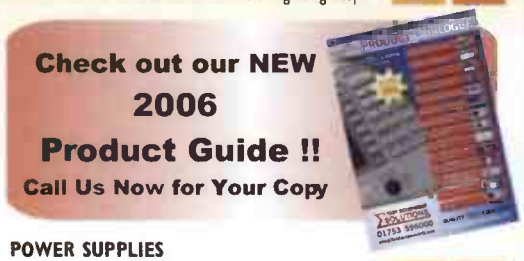

	Sale (GBP)	Rent (GBP)
POWER SUPPLIES		
ACDC Electronics EL750B 750W Electronic Load	1350	49
Wide Range of AT Supplies in Stock from 40W-2000W and up to 600V - prices range from	550	20
Farnell 3502A 35V/2A Dual Output DC Power Supply	325	26
Farnell AP100/90 100V 90A 3KW Power Supply	2550	77
Farnell AP60/50 60V 50A Power Supply	1750	58
Farnell LI2-10C 12V 10A DC Power Supply	450	22
Farnell LT30/2 30V 2A Power Supply	165	17
Farnell PSD3510A 35V 10A DC Power Supply	365	30
Farnell TOPS 3D 17V 1A Triple Output DC Power Supply	120	12
Kikusui PLZ-300W 300W Electronic Load	995	40
Racal 9232 30V/2A Dual Channel PSU	195	20
Tek 1101A Dual Probe Power Supply	150	15
Thandar PL310 (Dual) 32V 1.1A Dual Output PSU	159	16
Thandar PL310 32V 1.1A DC PSU	100	10
Thandar PL320 32V 2A DC PSU	135	15
Thandar PL330DP 32V 3A Dual Ch PSU with GPIB/RS232	450	27
PULSE GENERATORS		
AT/HP 8012B 50MHz Pulse Generator	695	33
AT/HP 8112A 50MHz Pulse Generator	1450	44
AT/HP 8130A/020 300MHz Pulse Generator	5500	220
AT/HP 8131A 500MHz High-speed Pulse Generator	5250	210
AT/HP 8133A 33-3000MHz Pulse Generator	11650	470
AT/HP 8160A 50MHz Pulse Generator	1350	41
SIGNAL & SPECTRUM ANALYSERS		
Advantest R3261A 2.6GHz Spectrum Analyser	3450	104
Advantest R3271 100Hz-26.5GHz Spectrum Analyser	9550	382
Advantest R4131C 3.5GHz Spectrum Analyser	2950	118
Advantest R9211A 10MHz-100kHz Dual Channel FFT Analyser	2950	89
Advantest U3641 3GHz RF Spectrum Analyser	3250	100
AT/HP 3561A 100kHz Dynamic Signal Analyser	2250	68
AT/HP 3562A 100kHz Dual Channel Dynamic Signal Analyser	2450	74
AT/HP 35660A 102.5kHz Dual Channel Dynamic Signal Ana	3250	130
AT/HP 35665A 102.4kHz Dual Channel Dynamic Signal Ana	3750	150
AT/HP 35B5B 40MHz Spectrum Analyser	4450	178
AT/HP 53310A 200MHz Modulation Domain Analyser	2350	99
AT/HP 53310A/001/030 2.5GHz Modulation Domain Analyser	3850	160
AT/HP 8562A 22GHz Spectrum Analyser	7600	309
AT/HP 8562E 13.2GHz Spectrum Analyser	11950	359
AT/HP 8591E/021/101/105 1.8GHz Spectrum Analyser	3900	156
AT/HP 8593E/004/041/130 22GHz Spectrum Analyser	10950	329
AT/HP 8594E/004/041/101/105/J62 2.9GHz Spectrum Ana	4950	149
AT/HP 8594E/004/041/105/151/163 2.9GHz Spectrum Ana	4750	143
AT/HP 8594E/010/041/101/105 2.9GHz Spectrum Analyser	6450	194
AT/HP 8594E/041 2.9GHz Spectrum Analyser	4500	135
AT/HP 8594E/041/102/130 2.9GHz Spectrum Analyser	4750	143
AT/HP 8595E 6.5GHz Spectrum Analyser	6000	181
AT/HP 8596E/021/101/105 12.8GHz Spectrum Analyser	8550	260
AT/HP 8901B 1.3GHz Modulation Analyser	1750	70
Anritsu MS2601B 2.2GHz Spectrum Analyser	2950	90
Anritsu MS2602A/01/02 100Hz-8.5GHz Spectrum Analyser	5750	174
Anritsu MS2661C/2/8 3GHz Spectrum Analyser	4250	128
Anritsu MS2663C/1/2/4/6/7/10/12 9kHz-8GHz Spectrum Ana	6550	262
Anritsu MS2711A 3GHz Handheld Spectrum Analyser	2950	89
Anritsu MS2711B 3GHz Handheld Spectrum Analyser	3450	104
R&S FSEA20/B4/B7/B8 9kHz-3.5GHz Spectrum Analyser	9500	285
Racal 9008A Modulation Meter	220	18
SRS SR760 100kHz Single Channel FFT Spectrum Analyser	1950	65
SIGNAL GENERATORS		
AT/HP 83731B/1E5 1-20GHz Synthesised Signal Generator	12500	375
AT/HP 8648A 100kHz-1GHz Synthesised Signal Generator	2500	100
AT/HP 8648D 9kHz-4GHz Synthesised Signal Generator	5950	179
AT/HP 8657A 1GHz Synthesised Signal Generator	1600	48
AT/HP 8657D/001 1GHz DQPSK Synthesised Sig Generator	1350	41
AT/HP E4431A/UN3 2GHz Digital Signal Generator	5250	158
AT/HP E4431B 250kHz-2GHz Synthesised Digital Signal Gen	5950	179
AT/HP E4432A/1E5/1EH 3GHz Synthesised Signal Generator	6250	188
AT/HP E4433A/1E5 250kHz-4GHz Synthesised Signal Gen	8100	324
Anritsu 68047C/2A/16 10MHz-20GHz CW Generator	11250	451
Anritsu 69059A 10MHz-26.5GHz CW Generator	9300	372
Anritsu MG3601A/02 1GHz Signal Generator	1600	48
Marconi 2019A 1GHz Synthesised Signal Generator	900	45
Marconi 2022E 10kHz-1.01GHz Synthesised Signal Generator	1250	38
Marconi 2024 9kHz-2.4GHz Synthesised Signal Generator	3750	150
Marconi 2030 1.35GHz Signal Generator	3950	158
Marconi 2032 10kHz-5.4GHz Signal Generator	8050	322
Marconi 2051 10kHz-2.7GHz Digital & Vector Signal Gen	5500	222
Marconi 2052 5.4GHz Digital & Vector Signal Generator	8950	360
National VP-7201A 500kHz RC Oscillator	485	28
Philips PM5330 180MHz Signal Generator	550	33
R&S SMHU58/B3 4.32GHz Signal Generator	6950	278
R&S SMIQ03A/B1/11/12/14/15/20/42/45/50 3.3GHz VS Gen	9950	299
WIRELESS		
AT/HP 8902A 1.3GHz Measuring Receiver	7950	240
AT/HP 8920A/103 1GHz Radio Comms Test Set	3250	98
Anritsu ME4510B Digital Microwave System Analyser	4750	145
IFR 2935 GSM Test Set [TRI BAND]	4950	149
Marconi 2955 1GHz Radio Comms Test Set	1950	71
Marconi 2955B/06 1GHz Radio Comms Test Set	3250	98
R&S CMS52/B22/B26/B53/B55 1GHz Radio Comms Test Set	4500	180
Racal 6104/001/002/003/006/014/04T Digital Mobile RTS	7250	220
W&G 4106 GSM/DCS1800/PCN1900 Mobile Phone Tester	2250	70
W&G 4107S Mobile Phone Tester	2450	74
Wavetek 4201S Triband Digital Mobile Radio Test Set	3650	110

The only company with the confidence to offer 12 month warranty as standard, with 24 month available!

FREE Palm Zire22 shipped with every order over £5K

Prices shown are for guidance in £UK GBP, exclusive of VAT and Ex-Works. All items subject to prior sale. Rental prices are per week for a rental period of 4 weeks. Free carriage to UK mainland addresses on sale items. Rental or non UK deliveries will be charged at cost. This is just a selection of equipment we have available - if you don't see what you want, please call. All items are supplied fully tested and refurbished. All manuals and accessories required for normal operation included. Certificate of Conformance supplied as standard. Certificate of Calibration available at additional cost. Test Equipment Solutions Ltd Terms and Conditions apply. All E&OE.

01753 596000

Fax: 01753 59 6001

www.TestEquipmentHQ.com

email: info@TestEquipmentHQ.com

Gio Feb 06