May 2011
Volume 117
Issue 1901
£4.60

Electronics WORLD

THE ESSENTIAL ELECTRONICS ENGINEERING MAGAZINE

A New Perspective in Oscilloscope Signal Analysis –

WaveRunner 6 Zi series

www.electronicsworld.co.uk

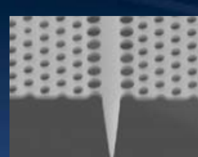

Technology
New developments in photonic technology

Nanometer Challenges
Issues that are slowing down the rush

Products
Wireless, active and passive components and other systems

PLUS

SPECIAL REPORT:
Designing Portable Systems

9 770959 833110 05

REGULARS

12
Focus

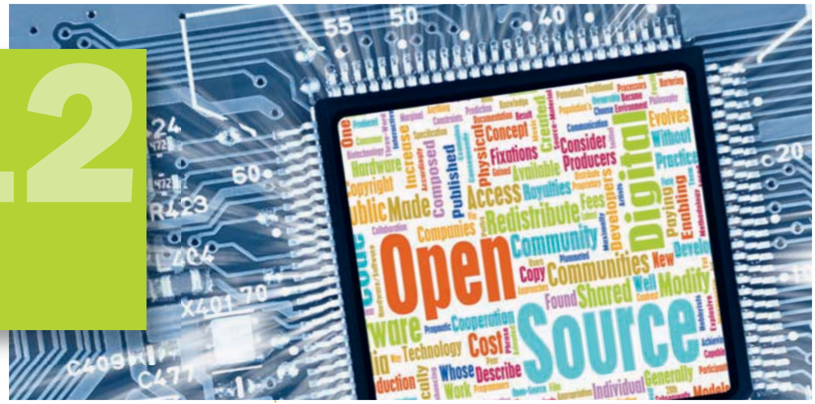

39
Warning System

FEATURES

Cover supplied by
LECROY
More about WaveRunner 6Zi
on pages 8-9

Dear Readers,
We'd like to inform you that ***Electronics World*** *magazine will no longer be available to purchase through news trade means commensurate from the June 2011 issue. Nevertheless, the magazine will continue to be available to order – as usual – via subscription directly from our publishing house, either in print or digital version, as our commitment to provide a quality product to our readers continues unabated. For further information and pricing please visit our website at* ***www.electronicsworld.co.uk***
Svetlana Josifovska*, Editor, Electronics World*

Disclaimer: *We work hard to ensure that the information presented in Electronics World is accurate. However, the publisher will not take responsibility for any injury or loss of earnings that may result from applying information presented in the magazine. It is your responsibility to familiarise yourself with the laws relating to dealing with your customers and suppliers, and with safety practices relating to working with electrical/ electronic circuitry – particularly as regards electric shock, fire hazards and explosions.*

2011 IS A WINDOW OF OPPORTUNITY FOR TECH M&A

Compared with the cyclical low in 2009, technology sector deal volumes recovered over the course of 2010

The alignment of cash-rich trade buyers, private equity sponsors with funds to invest and continued improvement in financing liquidity is opening a window of opportunity for technology merger and acquisition (M&A) activity, says a new report by PricewaterhouseCoopers (PwC).

"More high quality assets are now being brought to market as sellers seek to take advantage of the presence of cash-rich corporates with a renewed appetite for mega deals, a resurgence of private equity buyers and improving capital markets," said Andy Morgan, partner at PwC. "We are seeing a change in deal dynamics as appetite returns. The number of pre-emptive deals is on the rise with buyers originating deals rather than competing in traditional auction processes. Growth is firmly back on the corporate M&A agenda."

Compared with the cyclical low in 2009, technology sector deal volumes recovered over the course of 2010, while deal values rebounded more strongly. Some 393 deals completed in 2010 (up 32% from 2009), with a total value of €75.7bn – more than double the value completed in 2009 (€36.8bn). The recovery in M&A activity remains relatively cautious with both deal volumes and values still some way off the highs of 2007 and 2008.

Hot spots for deals in 2011 include:

FINANCIAL TECHNOLOGY – HOT

The regulatory push and a greater focus on transparency and risk management is fuelling demand for more highly integrated software solutions. Some banks will design and develop their own IT platforms but many are looking for external providers, and the area is likely to receive sharp attention from overseas players this year, especially from India. Misys's acquisition of software vendor Sophis for Euros 435m last November is a good example of consolidation in this sector.

CLOUD CLEARING

The development of virtualization and security technologies will underwrite continued interest in Cloud technologies. The prospect of significant cost savings and ability to handle large data volumes is making companies sit up and consider third parties handling their data and applications, and take responsibility for their security.

Deal multiples in the Cloud space are high. Cincinnati Bell and Montagu Private Equity paid double digit EBITDA (earnings before interest, taxes, depreciation and amortization) multiples for Cyrus One and Host Europe, but these look moderate compared with the prices paid by HP and Dell in the data storage space.

MOBILE DATA – GETTING SMARTER

Mobile application deals will come of age this year with all key areas of mobile technology attracting attention: advertising, payments, messaging, gaming and content. This sector has historically been highly fragmented but larger players have started acquiring in all areas of mobile data as the growth of the smart phone continues apace.

Google and Apple both made mobile advertising acquisitions last year and have started to look at mobile payments. In mobile gaming too, Electronic Arts has made a move recently acquiring Angry Birds developer Chillingo, while Disney acquired games publisher Tapulous in July.

UK PUBLIC SECTOR – OPPORTUNITIES IN THE CUTS

Big changes in the public sector will also mean opportunities for some. "Saving costs on the scale that the [UK] government requires will necessitate a new approach to outsourcing and ICT. We expect increased collaboration between government departments and local authorities which will drive the demand for master data management technology," said Philip Shepherd, partner at PwC.

PRIVATE EQUITY BOUNCES BACK

The heavily leveraged private equity business model of the credit boom years came to grief when the supply of cheap debt evaporated in the wake of the sub-prime crisis. However, private equity is back – the combined value of all PE deals in 2010 was Euros 10.7bn, up almost 60% from 2009.

Pent-up demand from funds with committed capital to invest and signs of improvement in the fundraising market are underpinning a more aggressive presence from private equity. Improving financing liquidity, meanwhile, is allowing private equity to complete larger deals.

*PricewaterhouseCoopers provides industry-focused assurance, tax and advisory services (**www.pwc.com/uk**).*

EDITOR: Svetlana Josifovska
Email: svetlanaj@stjohnpatrick.com

PRODUCTION MANAGER: Tania King
Email: taniak@stjohnpatrick.com

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

PUBLISHER: Wayne Darroch

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

SUBSCRIPTION RATES:
1 year: £46 (UK); £67.50 (worldwide)

MISSING ISSUES:
Email: electronicsworld@circdata.com

NEWSTRADE:
Distributed by Seymour Distribution Ltd,
2 East Poultry Avenue, London, EC1A 9PT
Tel: +44 (0) 20 7429 4000

PRINTER: William Gibbons Ltd

ISSN: 1365-4675

St John Patrick Publishers

Light-Speed Developments in Photonic Crystal Technology

European scientists and engineers are working together on a project called COPERNICUS to develop cutting-edge photonic crystal technology that has the potential to make electronic devices much faster, smaller and more efficient.

Photonic technology uses light – instead of electric current – to send and receive signals at extremely high speeds, and photonic crystals are nanoscale materials, enabling unprecedented control of light and the miniaturisation of key functions. They represent a so-called 'disruptive' technology, which means they have the potential to completely change the way things are currently done in this field.

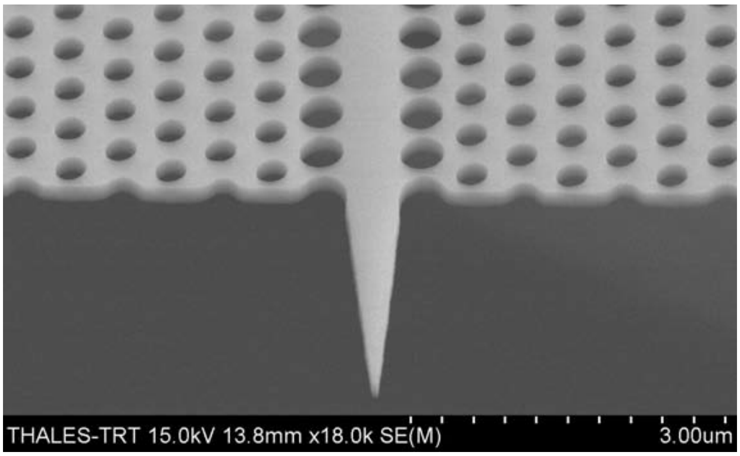

A photonic crystal membrane with a tapered mode adaptor, which allows light from an optical fibre to be coupled into the membrane

The efforts of the eight European academic and industrial partners working on the COPERNICUS project are expected to place Europe at the forefront of this field, which according to the Photonics21 European Technology Platform, world market reached €270bn in 2008, of which €55bn was produced in Europe – a growth of nearly 30% since 2005.

"We believe that our approach has all the hallmarks of a highly disruptive technology with the potential to place Europe at the forefront of photonics," said Alfredo de Rossi, Project Coordinator at Thales Research and Technology in France, one of the member companies.

A key aim of the COPERNICUS project is to develop very high speed, compact demultiplexing receivers, used to separate optical signals that have been transmitted together. These can be used where several light signals of different wavelengths or colours are transmitted together and must then be separated by the receiver so that the signal can be reconstructed. To achieve this, the consortium will target technological breakthroughs in ultra-compact integrated optical devices including switches, filters and detectors. Photonic devices will address the pressing need for low-power, ultra-high bandwidth data links in server farms, optical storage networks and on-board Internet/entertainment systems, where demand is driving the data bandwidth and technology integration level rapidly upwards. Next generation telecom systems will also benefit from these devices.

COPERNICUS, which runs until the end of 2012, has received funding worth nearly €3m from the European Commission's Information Society Technologies Programme.

Flash Disk Security Questioned

There have been recent reports stating that certain types of flash disks lack a secure deletion facility. However, storage systems specialist firm Origin Storage has taken this opportunity to distance magnetic drives from such reports as their architecture and, therefore, approach to security is very different to that of solid state storage devices (SSDs).

"A lot of companies have made the understandable mistake of presuming that flash drives are a slot-in replacement for magnetic drives, when in fact nothing could be farther from the truth," said Andy Cordial, managing director of Origin Storage.

Researchers at the University of California have discovered that the electronic data shredding procedures, also known as data sanitisation, do not always work the same on SSDs as on magnetic drives. This is due to the complex electronics on some of the latest generation of SSDs, which intercepts a data delete request and often only deletes the header, rather than the full data clusters that go to make up a given file on a magnetic drive. As such the so-called 'disk doctor' programs, which allow data retrieval on a sector-by-sector basis without resorting to requiring header data – as an operating system normally does, can effectively undelete supposedly sanitised data files on an SSD.

The bottom line, says Cordial, is that 'conventional' data overwrite commands that have worked well on magnetic drives since the earliest days of PCs in the 1980s cannot be relied upon to function in the same manner with a flash drive.

"As the university researchers found, the erase procedures provided by manufacturers should be verifiable as well, so that users could easily check post-sanitisation that their data had been removed," he said. "This is why we recommend SSDs for specific applications and magnetic drives for other uses. It's also why, where high levels of security are required, we recommend magnetic drives with additional levels of security, such as PIN/password entry system."

NEWS IN BRIEF

A new scheme in the UK will offer funding to small companies to enable them to research and develop new products, processes and services that could stimulate economic growth.

Managed by the Technology Strategy Board, the Grant for Research and Development scheme is already in effect.

The Court of Justice of the European Union (CJEU) has rejected the proposal for a unified patent litigation, deeming it incompatible with the provisions of the Treaty on European Union and the Treaty on the Functioning of the European Union, the two principle treaties on which the EU is based.

The proposal included a single European court structure which would handle patent litigation for several contracting states of the European Patent Organisation, potentially including all EU states as well as other European countries.

NSF-FUNDED ANTENNA PROJECT AIMS TO RELIEVE WI-FI SIGNAL INTERFERENCE

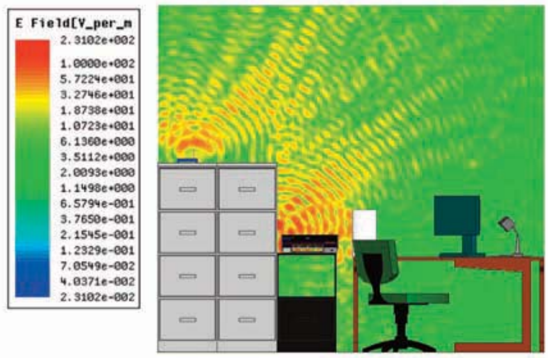

Typical Wi-Fi signal interference issues in an office environment

Gonzaga University in Spokane, US, received a $1.2m funding to develop a Smart Antenna and Radio Laboratory in part to investigate more reliable high-bandwidth wireless communications via Wi-Fi.

One of the Gonzaga research projects is aimed at overcoming the growing problem of wireless signal interference, as many users try to communicate simultaneously over the 2.4GHz band used for Wi-Fi. The smart technologies developed by the team – headed by Steven Schennum, an electrical engineering professor – will enable antennas to focus on one user signal at a time. For example, for a Wi-Fi user working on a laptop with a weak or cross-polarized signal, a smart antenna system would utilize algorithms to optimize the signal to that individual laptop. Using engineering simulation, Schennum and his team will develop new multi-antenna techniques that improve both the efficiency and bandwidth of wireless communications.

"We're creating a state-of-the art anechoic chamber for testing our physical antenna prototypes, but even the best antenna test chambers are limited in their size and shape, the performance of their absorptive materials and the range of frequencies they can accommodate," said Schennum.

The simulation software has been provided by ANSYS. "By simulating electromagnetic fields and currents in a virtual environment using ANSYS software, we can test the performance of our antenna designs for any location, plane or geometry – and over a limitless range of frequencies – before moving to the prototype stage," added Schennum.

The National Science Foundation (NSF) grant provides Gonzaga with funding for dedicated computers running ANSYS software that simulates smart antenna circuits and electromagnetic fields in three-dimensional structures.

Silver-Diamond Composite Offers Unique Capabilities for Cooling Powerful Defense Microelectronics

Researchers at the Georgia Tech Research Institute (GTRI) are developing a solid composite material to help cool small, powerful microelectronics used in defense systems. The material, composed of silver and diamond, promises an exceptional degree of thermal conductivity compared to materials currently used.

The research focuses on producing a silver-diamond thermal shim of unprecedented thinness – 250 microns or less. The ratio of silver to diamond in the material can be tailored to allow the shim to be bonded with low thermal expansion stress to the high power wide-bandgap semiconductors planned for next generation phased-array radars.

Thermal shims are needed to pull heat from these high power semiconductors and transfer it to heat dissipating devices such as fins, fans or heat pipes. Since the semiconductors work in very confined operating spaces, it is necessary that the shims be made from a material that packs high thermal conductivity into a tiny structure.

Diamonds provide the bulk of thermal conductivity, while silver suspends the diamond particles within the composite and contributes to high thermal conductivity that is 25% better than copper. To date, tests indicate that the silver-diamond composite performs extremely well in two key areas: thermal conductivity and thermal expansion.

"We have already observed clear performance benefits – an estimated temperature decrease from 285^{0}C to 181^{0}C – using a material of 50% diamond in a 250-micron shim," said Jason Nadler, a GTRI research engineer who is leading the project.

Diamond is the most thermally conductive natural material, with a rating of approximately 2,000 watts per meter Kelvin, which is a measure of thermal efficiency. Silver, which is among the most thermally conductive metals, has a significantly lower rating of 400W per meter K.

Nadler explained that adding silver is necessary to bond the loose diamond particles into a stable matrix; to allow precise cutting of the material to form components of exact sizes; to match thermal expansion to that of the semiconductor device being cooled; and to create a more thermally effective interface between the diamonds.

Microscope image showing diamond particles [Credit: Jason Nadler]

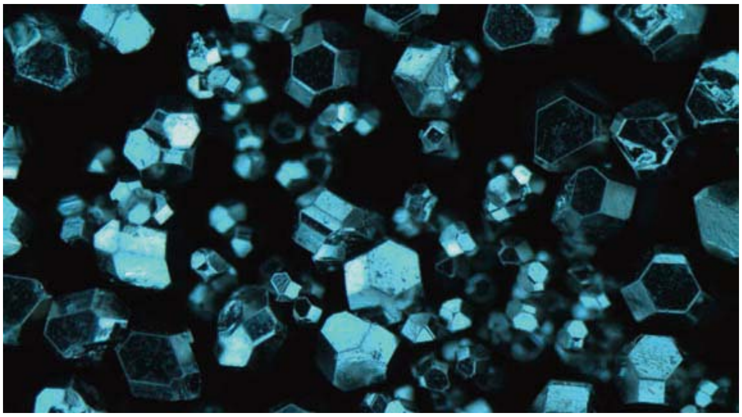

A NEW PERSPECTIVE
in Oscilloscope Signal Analysis

New LeCroy WaveRunner 6 Zi Oscilloscopes provide class leading performance, speed and versatility available from 400MHz to 4GHz in combination with the first rotating touch screen for horizontal and vertical signal display. Industry first 12-bits Oscilloscopes with 400 MHz and 600 MHz to follow in May 2011

LeCroy's new line of WaveRunner 6 Zi oscilloscopes define superiority in a test instrument with a powerful feature set including a wide range of application packages, advanced triggering to isolate events, a user interface developed for quick and easy navigation, a wide range of probing options, and lightning-fast performance. Still the price of the new WaveRunner wasn't changed compared to the previous generation.

The WaveRunner 6 Zi oscilloscope family features a pristine signal path that offers unmatched signal fidelity with low noise. This performance is augmented by a huge offset and timebase delay adjustment to allow easy signal and amplifier performance assessment and zooming on vertical and horizontal signal characteristics. With deep memory up to 128 Mpts, up to 40 GS/s sample rate, low noise, fast operation, a full line of probing solutions, and extensive serial data analysis tools, they are the most versatile oscilloscopes in the 400 MHz to 4 GHz class. In addition, LeCroy has designed the WaveRunner 6 Zi with a pivoting display that permits viewing signals vertically as well as horizontally to obtain more detail for optimum analysis. The comprehensive toolset includes everything engineers need to validate a design, debug errors at board bring up, and characterize an embedded system. The WaveRunner 6 Zi is the ultimate debug machine.

LeCroy is also announcing for May the first product line of 12-bit oscilloscopes with 400 MHz bandwidth, 2 GS/s sample rate and up to 256 Mpts.

FIRST ROTATING DISPLAY

The 12.1" high resolution WXGA wide screen is designed to provide the best view of any signal type on the display. The widescreen is ideal for a variety of signals where long records are required and zooming or scrolling Rotate the screen 90° degrees to optimize the display for viewing digital signals, jitter tracks, eye diagrams, and frequency plots. The screen image will adjust automatically when rotated. Tilt the display up or down in either orientation to minimize reflections or glare.

FAST THROUGHPUT WITH X-STREAM II STREAMING ARCHITECTURE

WaveRunner 6 Zi makes long acquisition memory a pleasure to use with X-Stream II fast throughput streaming architecture, available exclusively from LeCroy. This architecture uses variable waveform segment lengths to improve CPU cache memory efficiency. Traditional digital oscilloscopes that process full-size waveforms in one length pay a performance penalty due to overly simplistic architectures that are inattentive to the CPU architecture. The result – WaveRunner is much faster when processing long records compared to other oscilloscopes. X-Stream II is augmented with an Intel® Dual Core™ processors, high-speed data buses, a 64-bit Windows7TM OS, and up to 4 GB of RAM.

Serial Data Analysis Tools

		Decode	Trigger	Measure/Graph	ProtoSync	QualiPHY
Embedded	I^2C	•	•	•		
	SPI	•	•	•		
	I^2S	•	•	•		
	UART, RS-232	•	•	•		
Automotive	CAN	•	•	•		
	LIN	•	•	•		
	FlexRay	•	•	•		
Military & Avionics	ARINC 429	•		•		
	MIL-STD-1553	•	•	•		
Handset Cellular Mobile	DigRF 3G	•		•		
	MIPI D-PHY /CSI-2/DSI	•				•
	DigRF v4	•		•		
Storage / Peripherals Interconnects	8b/10b	•	•			
	Fibre Channel	•			•	
	SATA (1.5 & 3 Gb/s)	•	•		•	•
	SAS (1.5 & 3 Gb/s)	•			•	
	PCIe (Gen1)	•	•		•	•
	USB 2.0	•	•		•	•
	LPDDR2					•
	DDR2					•
	Ethernet					•

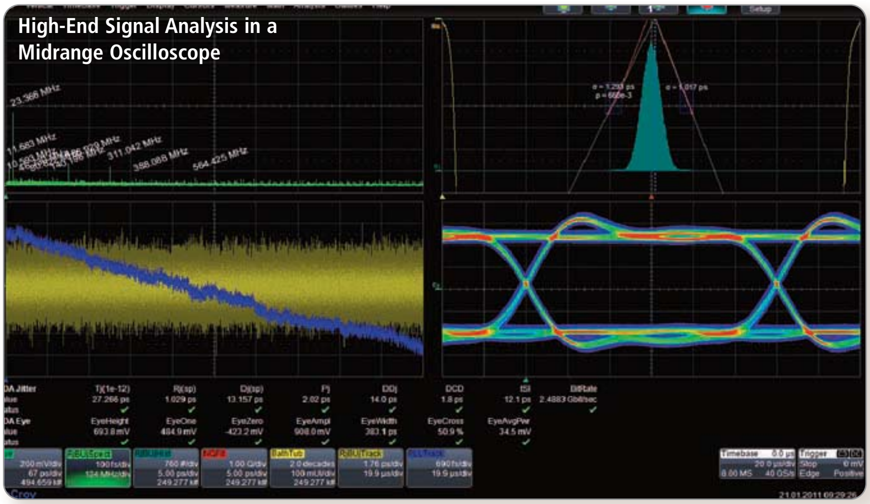

A NEW WAY TO NAVIGATE

The WavePilot control area provides convenient operation of Cursors, Decode, WaveScan,™ History, LabNotebook,™ and Spectrum by their respective function buttons on the front panel. The SuperKnob is a joystick-like knob in the center of the WavePilot control area used to easily navigate through tables, zoom and position waveforms, and quickly document and annotate your setups.

MOST COMPREHENSIVE SERIAL DATA ANALYSIS

Solving serial data problems requires intimate knowledge of the protocol to get started. With the WaveRunner 6 Zi, the oscilloscope is the expert. Simply connect your probes or cables and the scope can provide the correct level of detail needed to view, debug, and analyze the serial data signals. Solutions address the Embedded, Military and Avionics, Handset/Mobile/ Cellular, and Storage/Peripherals/ Interconnects, with a combination of decode, trigger, measure/ graph, ProtoSync, and compliance tools. Whether the protocol under test is a new emerging standard requiring jitter end eye diagram testing, a mature standard requiring compliance testing, or an embedded standard requiring protocol and measurement and timing analysis, with over seventeen trigger, decode, and compliance solutions, WaveRunner 6 Zi has it all.

The new JITKIT Jitter package makes it simple and easy to understand the basic system jitter performance of clock signals and clock-data activities, including period, half period, cycle-cycle, skew, amplitude, differential voltage crossing, slew rate, and a wide variety of other common jitter measurements.

MORE TRIGGER CAPABILITY ISOLATES MORE PROBLEMS MORE QUICKLY

A powerful combination of high bandwidth edge and 10 different SMART triggers, four stage cascade triggering, measurement trigger, and Triggerscan are all standard, and will isolate the problem quickly to permit rapid focus on the cause. The measurement trigger offers a powerful option to qualify a trigger event based on a qualified measurement with great resolution. A high-speed serial trigger enables triggering on up to 3 Gb/s serial patterns of up to 80-bits in length. A full range of serial triggers (I2C, SPI,UART, RS-232, Audio (I2S, LJ, RJ, TDM), CAN, LIN, FlexRay, MIL-STD-1553, SATA,PCIe, 8b/10b, USB2 and many others) are also available.

CUSTOMIZED TOOLS

Only LeCroy completely integrates third party programs into the oscilloscope's processing stream by allowing you to create and deploy a new measurement or math algorithm directly into the oscilloscope environment and display the resulton the oscilloscope in real-time. Use C/C++, MATLAB,® Excel, Jscript (JAVA), and Visual Basic to create your own customized math functions, measurement parameters, or other control algorithms.

HISTORY MODE

History mode lets you scroll back in time to isolate those anomalies, measure them with parameters or cursors, and quickly find the source of the problem. History mode is always buffering waveforms, so no user action is required to save traces, only to invoke the viewer.

TriggerScan uses high-speed hardware triggering capability with persistence displays to capture only the signals of interest and provide answers up to 100 times faster than other methods. Traditional fast display update modes work best on frequent events occurring on slow edge rates while TriggerScan excels in finding infrequent events on fast edge rates. In addition to the general purpose WaveShape Analysis tools, application specific solutions are available for Serial Data Compliance, Embedded Design, Digital Design, and Automotive. These options extend the LeCroy standard measurement and analysis capabilities and expand the WaveRunner 6 Zi's utility as needs change.

As for all its oscilloscope, LeCroy offers 3 years warranty for the WaveRunner 6Zi series.●

Landscape and Portrait Signal View with the Industry First Rotating Oscilloscope Display

ABOUT LECROY

LECROY CORPORATION IS A WORLDWIDE LEADER

in serial data test solutions, creating advanced instruments that drive product innovation by quickly measuring, analyzing, and verifying complex electronic signals. The Company offers high-performance oscilloscopes, serial data analyzers, and global communications protocol test solutions used by design engineers in the computer and semiconductor, data storage device, automotive and industrial, and military and aerospace markets. LeCroy's 45-year heritage of technical innovation is the foundation for its recognized leadership in "WaveShape Analysis"—capturing, viewing, and measuring the high-speed signals that drive today's information and communications technologies. LeCroy is headquartered in Chestnut Ridge, New York. Company information is available at http://www.lecroy.com.

Whatever Happened To Security Of SUPPLY?

MALCOLM PENN IS CEO AND CHAIRMAN OF MARKET ANALYST FIRM FUTURE HORIZONS, BASED IN THE UK

Some natural disasters change history; the tragic March 11 earthquake and tsunami in Japan may well turn out to be one of them. With luck, a much healthier and more soundly based electronics industry will emerge.

We have long voiced concern over what we saw as the industry's offhandedness towards its rationale and business practices whereby, with ever increasing momentum, common sense and sound practice were being progressively cast aside for the sake of near-term profits or "increased shareholder (stake holder) value". Where once company presentations focused on vision, intellect, creativity and new products, nowadays it seemed ROI (return on investment), ROE (return on equity), RONA (return on net assets), cash flow and margin dominated the agendas.

Of course profit and cash flow is important to long-term survival but our 'old-fashioned' belief has always been these were diagnostic – they came from doing the job right – not by manipulating the balance sheet. Business wheezes included selling and leasing back assets, off-shoring, lean operations and outsourcing everything and anything, reaching such monumental proportions that, with rare exceptions, the corporate tail is now wagging the dog, with the balance sheet driving the business not the business driving it.

We raised this issue in this column six months ago "*The Broken Industry Model – Back To Business Basics*"; and were prophesying then an industry-wide business model disaster waiting to happen. The supply chain mess the industry now finds itself in would have happened anyway; the earthquake and tsunami merely accelerated its arrival. We wish that it had not been this way, not just because of the suffering in Japan but because it has given executives and industry a messenger to blame rather than being forced to face up to its past bad decisions.

We have always counselled about the semiconductor industry's strategic nature and the fact its impact stretches way beyond its US$ 300bn intrinsic value, ultimately driving 10% of world GDP. What makes this even more important however was its absolute uniqueness throughout the whole of the food chain; from the raw materials it uses to the manufacturing equipment it needs to turn the sand into devices. Everything is special.

> *The supply chain mess the industry now finds itself in would have happened anyway; the earthquake and tsunami merely accelerated its arrival*

That in turn means the electronics industry that makes end products out of ICs, and in turn their customers and the system and service providers, all owe their existence and livelihood to an exclusive and near-term irreplaceable resource. Yet it was increasingly taken for granted, both by the chip industry itself – "wafer manufacturing is not a strategic competence, outsource it", to the chip industry's customers – "how much is your price today?"

Just-in-time, on-demand, lean and batch and outsourced manufacturing models took over from inventory, WIP, production lines and multiple sourcing, despite the fact that the whole manufacturing process takes six months from wafer build to end product delivery.

Cracks first appeared last September when Nissan was forced to shut down production due to a lack of engine management modules, itself the result of a shortage of ICs. No one was managing the overall supply chain; no one appreciated the long supply lead-times involved. Incidents like this were set to increase dramatically anyway in the second half of this year, due to the fact the supply chain had been squeezed too far. The earthquake and tsunami merely forced industry's hand.

Yet industry still seems to be in denial; the real implications have yet to sink in, with the issues being treated more as a disaster recovery plan rather than the need to fundamentally rethink the supply chain model.

Past squeezes on the supply chain, always adversarial in nature, have in many cases reduced the number of suppliers to just a few, deemed to be a desirable and natural result of the industry's maturity and consolidation. This may well be true, but does it make it wise? Is it beneficial to the industry as a whole that so few can effectively determine the fate of up to 10% of the world's GDP should disaster – or even bad (e.g. monopolistic) management – strike, especially when contingency planning and supply chain security are deemed irrelevant in the balance sheet driven world that the industry had boxed itself into?

Now disaster has struck, in the short term there is enough work in progress and inventory for firms to muddle through, but by the end of April, these will have been exhausted. This theoretically gives some breathing space to secure alternative suppliers, except for the fact this will be easier said than done, given the fact that all sectors of the supply chain were working flat out with little spare capacity before the disaster hit. A shortage-driven run on components will deplete the inventory much more quickly.

Restarting Japanese production is made worse by the fact that even factories not directly hit in Japan, are struggling to maintain their pre-earthquake production levels due to power shortages, rolling blackouts, infrastructure difficulties and problems with their employees getting enough food to eat and fuel to get to work. Full shift operation is impossible especially at Japan's crystal pulling and wafer fab plants. If you cannot pull the crystals to make the wafers (which requires 48 hours uninterrupted power), you cannot feed the wafer fabs, even if they could get sufficient guaranteed power to run – and keep running – the wafer processing gear. This will mean an on-going period of supply shortage hitting home in two to three months' time.

The wafer shortage will affect every fab in the world, not just those in Japan, as will the plastics and other materials shortages used e.g. in the back end processing. Assuming things get back to normal by the June-July time frame, it will then take three months for normal supplies to start resume, meaning IC shipment business as normal in Q4-2011 at the earliest, minus the output from any factories that were damaged beyond repair.

We have already seen OEM plant closures as a result of difficulty in getting supplies; this is more likely to get worse as the slow process of recovery gets underway. Fabs need power – lots of it and reliable – no one is going to risk restart their fab or wafer plant until such supplies can be relied upon. And with 40% and 22% of the world's wafer production and wafers processing respectively done in Japan, with no spare capacity, there is zero chance for this slack to be taken up elsewhere, even once getting raw materials, distribution, shipping product and employee absences because of transportation disruption are resolved.

We have already seen OEM plant closures as a result of difficulty in getting supplies; this is more likely to get worse as the slow process of recovery gets underway

So much then for the "wafers aren't strategic" school of thought punted by the IDMs going fabless/fablite and OEMs switching chip suppliers at the drop of a price reduction! So much too for a world where supply chain security played second fiddle to balance sheet gerrymandering.

I for one hope that the longer-term outcome of this tragedy is a return to business, social and moral responsibility and the sort of capitalism promoted by Stanford University and espoused by the likes of Dave Packard and Bill Hewlett at the old Hewlett Packard, a company making great things with corporate tradition of decent behaviour and long-term focus. It is time for companies to serve the long-term interests of their customers and the world in general; and time for semiconductors to be treated as the strategic industry it is. ●

Managing Open Source Licensing for Semiconductors

KAMAL HASSIN, DIRECTOR OF R&D AT PROTECODE, EXPLAINS THAT THERE ARE A NUMBER OF APPROACHES TO LICENSE MANAGEMENT, RANGING FROM DOING NOTHING TO FULLY AUTOMATED REAL-TIME SCANNING OF SOFTWARE TO DETECT AND REPORT LICENSE OBLIGATIONS

Software is a massive enabler for the semiconductor industry: drivers, real-time operating systems (RTOS), software development kits (SDKs), networking and security, administration, media formatting and compression and, of course, an endless list of applications. The semiconductor industry now spends more on software development than on all other R&D aspects. But how do developers build all of this software at a reasonable cost?

Increasingly, developers supplement custom coding with open source software, enabling rapid development at a reduced cost with a high degree of flexibility, security and stability. Open source opens up new opportunities that did not exist even a few years ago. As with every significant opportunity comes the need for responsible practices to ensure fairness and sustainability. In simple cases, managing third party content and the associated licensing obligations manually is possible. However, in larger projects, this becomes cumbersome and inaccurate, leading to potential quality or licensing issues.

There are a number of approaches to license management, ranging from doing nothing to fully automated real-time scanning of software to detect and report license obligations. All approaches can be viewed from a cost perspective, with the aim of maximizing developer productivity while minimizing risk of licensing issues.

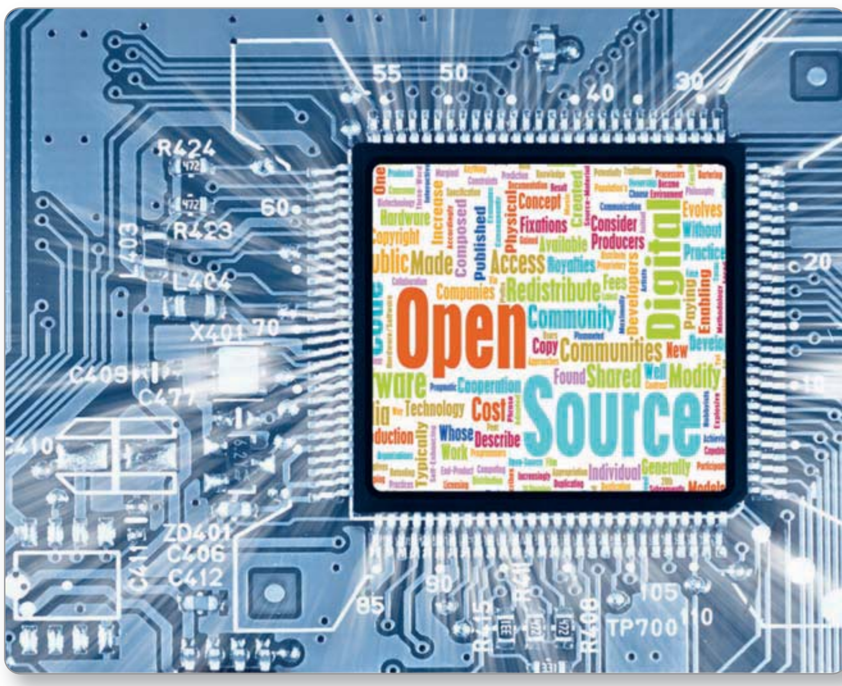

The cost of managing software license obligations is analogous to managing defects in the development process. It is well understood that the earlier a defect is identified and corrected, the less expensive it is – and the same is true for licensing obligations. Likewise, the earlier a development organization identifies “licensing bugs” that attract unacceptable license obligations, the less expensive it is to adjust the software to achieve licensing compliance.

Software auditing is required to ensure production code conforms to the licensing policy

Principle Aspects of License Management

A complete approach to assuring open source license compliance will generally include three major aspects:

1. Definition of a licensing policy which must be met by all software projects.
2. Auditing of all project software to detect any third party source code including open source that is unacceptable based on the licensing policy.
3. Corrective processes to ensure that all released software conform to the licensing policy.

Figure 1: How the cost of detecting and correcting "licensing bugs" escalates in later stages of the development process

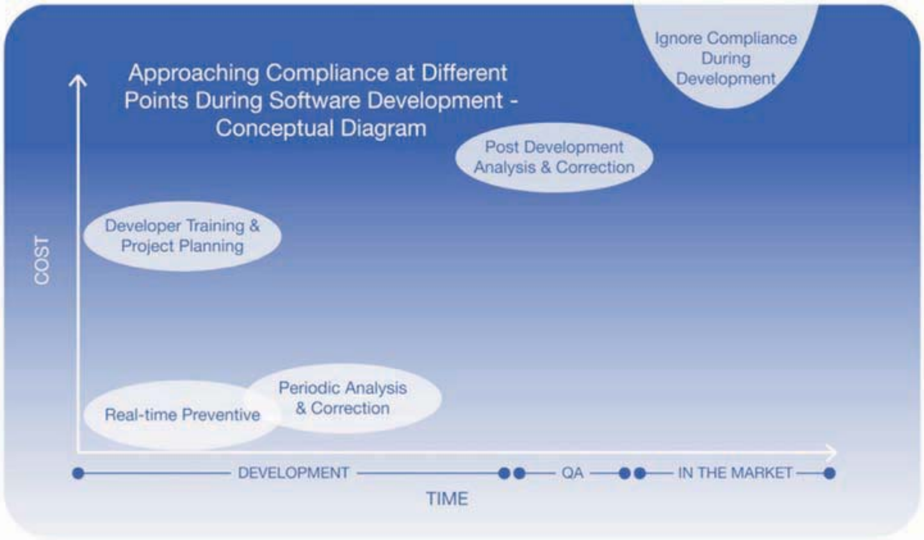

The creation of the licensing policy is an important step as it forms the basis upon which decisions regarding acceptance of open source software will be made. The licensing policy should be defined in accordance with both the business goals of the organization as well as its engineering processes, and generally requires the involvement of business and engineering managers, as well as legal counsel.

Software auditing is required to ensure production code conforms to the licensing policy, although the audit implementation can take a number of forms depending on the preference of the development organization. Audits can range from ad hoc developer training and post-development cycle auditing to proactive automated approaches such as periodic and real-time auditing.

Options for Managing Open Source Licensing

The following options are available to address license compliance at different points in the development process.

1. **Do nothing:** This option ignores the compliance issue because it carries the lowest up-front cost, but imposes the highest business risks and largest corrective costs as a product moves closer to launch.
2. **Developer training:** Some companies consider developer training and project planning is sufficient enough. This however, can be an expensive labour-intensive option given the increasing diversity of software licenses, the high cost of developer training and the constant churn within the development environment. With this option, compliance depends solely on busy developers and is prone to human error.
3. **Post-development license audits:** Auditing late in the project lifecycle does not impact the development workflow and can be implemented manually or using automated software tools. This option does, however, lead to more expensive rework due to additional system retesting cycles.
4. **Periodic assessment:** Licensing analysis during development allows for corrections along the way if license violations are detected. This type of analysis can be automated and tends to be less expensive than post-development assessment since changes and re-tests are always easier to undertake earlier rather than later in the cycle.
5. **Real-time preventive assistance at the developer workstation:** The most proactive approach is to audit software in real-time at the developer workstation. The development process is not disturbed and the cost of corrections is minimal as there is no impact to system integration and testing. This process can be automated and generally requires very little developer training.

Regardless of the type of auditing approach, the overall goal is to minimize the time and cost of correcting the final software release so that it meets all functional, quality and license compliance requirements. Figure 1 illustrates how the cost of detecting and correcting "licensing bugs" escalates in later stages of the development process.

Each organization must consider their approach, balancing the short term cost of developer training and tools versus the potential longer term cost of post-release legal problems.

Automating Open Source License Management

Tools are available to automate open source license detection and analysis. They can operate on demand, on a periodic schedule or in real-time within the development process. Generally such tools find compliance problems sooner, lowering the overall cost of license compliance and significantly reducing any risk of legal issues that could affect the deployment of semiconductor-based products. ●

To find out more about Managing Open Source Licensing for Semiconductors go to **http://protecode.com/november_news.php#semi**

Working from HOME

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

The traditional model of employment for electronic engineers has always been organised around the "lab", where a number of engineers (how many depends greatly on the company size and task at hand), covering a range of ages and levels of experience, are employed to work in a central location, where a pool of equipment and support services (ranging from drawing office staff to canteen facilities) are available to them. This site is frequently (although not always) the company's head office and/or its main production facility.

While this method of working has obvious benefits (staff immediately on-hand for meetings and discussions, sharing of scarce or expensive equipment, co-operative working and healthy "team" environments), it is not the only way of employing engineering staff. With the continued improvements in high-speed Internet access and the increased reliance on PC-based design tools and resources, it has become feasible to employ engineers in some disciplines as home-workers. In the low-power wireless module sector this is becoming more and more common. I have been working this way since 2003 and it can be highly successful.

Instead of commuting daily to a company lab, the engineer works from a small "lab" set-up in their home, using equipment provided by the employer. Communication with the rest of the company (management, production, fellow engineers) is by telephone, email and (relatively infrequent) progress meetings, while heavy use is made of postal/courier services for supply of parts and delivery of finished prototypes and samples.

This work-model is not appropriate to all engineering disciplines, but in low power RF it can be made to work very successfully.

As a working method, this has much to commend it, but it will not succeed without some diligence, both on the part of the engineer and the employer. Trust and mutual respect are vital components, combined with a good deal of flexibility. The engineer will find that many of the familiar parts of a working day (the enforced rhythm, the camaraderie) are absent and need to be replaced with self-reliance and willpower, while the employer must make the initial investment in test and computing equipment and then accept that the necessary work will be done, without intimate hour-by-hour supervision.

I've been working this way for eight years, which to me is proof enough.

PITFALLS:

Working alone is very different from working in a busy lab, as part of a team. It takes a different mindset and considerable willpower to set the work/life balance correctly (it's as easy to end up working excessive hours "because the lab is there"as it is to "goof-off" for days on end, because there's no-one to motivate you). You need to assign sufficient space in your home (a whole bedroom or study is a minimum: not just a corner or a kitchen table), and work schedules will impose on your existing home-life rhythm.

THERE IS THE RISK OF TURNING INTO A STRANGE, DE-SOCIALISED HERMIT – JUST LIKE ME

BENEFITS:

The time and money wasted in commuting is eliminated, as is the ecological impact of road or train travel. The company no-longer needs to find (expensive) space for the lab on-site, and it is possible to work solidly without interruptions. Lab equipment is always there when you need it, and you can be flexible with your hours. Your family will see more of you (in theory) and a proportion of your household costs (communications, heat/light/power/water) become tax-deductable.

IN MY EXPERIENCE PRODUCTIVITY ACTUALLY INCREASES

NANOMETER CHALLENGES

MARCO CASALE-ROSSI FROM SYNOPSYS CONSIDERS THE CHALLENGES THAT ARE SLOWING DOWN THE RUSH TO NANOMETER TECHNOLOGIES

The very first integrated circuit has recently celebrated its 50th birthday. It was a flip-flop made from four transistors and five resistors, manufactured at about 100 microns. Now we are witnessing the birth of the latest integrated circuits – dual core processors consisting of approximately 400 million transistors, manufactured at 32 nanometers. This is the result of what we sometimes call scaling. More often we refer to it as Moore's Law.

When Dr Gordon E. Moore, a co-founder of Fairchild Semiconductor and Intel, first put forward his conjecture, he took an engineer's viewpoint. He predicted that every new technology node would enable manufacturing – in the same silicon area and, therefore, at the same cost – twice as many transistors as the last; as a "by product" these transistors run 1.4x faster at each new node and the power density is the same. Moore's Law has proven to be valid for over 40 years and has been instrumental in growing the value of the semiconductor industry from nothing in 1959, to the massive $270bn that was forecast for 2010.

The Nanometer Rush

The 32/28nm technology node is about the 23rd generation of integrated circuits. It has taken 50 years to get this far, but the last 10 years or so has become a rush: the "nanometer rush".

Figure 1 shows the technical progress of the first dual core processors shipped at the last four technology nodes: 90, 65, 45 and 32 nanometers.

After Moore's Law, from one technology node to the next, integration capacity should double, frequency (clock rate) should improve by 40% and power should remain constant. What Figure 1 actually shows is that over the last four technology nodes integration capacity has increased, but is below the historical 2x level, while both frequency and power have remained about the same.

The nanometer rush, which has being going on exponentially for decades, is slowing down. Although Moore's Law is alive and well, the main problem is power.

Figure 1: Dual core processors from 90 to 32nm [Source: M.Bohr, Intel IDF 2008. www.intel.com 2010]

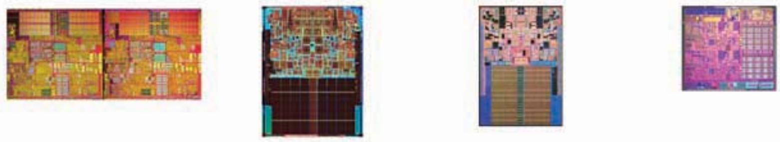

90 nanometers	65 nanometers	45 nanometers	32 nanometers
206 mm²	143 mm²	107 mm²	81 mm²
230M transistors	291M transistors	410M transistors	382M transistors
3.2GHz	3.0GHz	3.4GHz	3.5GHz
TDP = 130W	TDP = 80W	TDP = 80W	TDP = 73W
Moore's Law	Moore's Law	Moore's Law	Moore's Law
90 = 1X	**90/65 ⇨ 1.82X**	**65/45 ⇨ 1.88X**	**45/32 ⇨ 1.23X**
Clock Rate	Clock Rate	Clock Rate	Clock Rate
90 = 1X	**90/65 ⇨ 0.94X**	**65/45 ⇨ 1.13X**	**45/32 ⇨ 1.03X**
TDP @ 3.0GHz	TDP @ 3.0GHz	TDP @ 3.0GHz	TDP @ 3.0GHz
90 = 1X	**90/45 ⇨ 0.61X**	**65/45 ⇨ 1X**	**45/32 ⇨ 0.91X**

We can either achieve higher performance or lower power, but not both. Transistors get smaller but, because of higher power densities they struggle to deliver any further performance improvement.

Design – and EDA which helps engineers to automate electronic design – are increasingly important in dealing with nanometer challenges: they unleash the value – higher performance, lower power, sometimes both – of nanometer process technology.

Top Ten 32nm Challenges

The industry must tackle the following ten specific challenges in moving to 32/28nm technology:

1. Complexity

The designs shown in Figure 2 are an example of Moore's Law at work: over six technology nodes, integration capacity should increase by 25 = 32x. In this case, the absolute number of transistors has increased by 32x, but the die size has shrunk by 20%; therefore, the integration capacity has increased by 39x. Amazingly, power has remained the same.

> Although Moore's Law is alive and well, the main problem is power; we can either achieve higher performance or lower power, but not both.

It is very easy to distinguish between the six blocks in the 1996 chip. However, because the 2006 chip interleaves logical and physical hierarchy, it is much harder to tell which block does what. The 2006 chip is not only more challenging because of the technology node, the real challenge is the sheer complexity of the device. There are hundreds of millions of transistors, thousands of macros and multiple levels of hierarchy, which are pushing computer hardware and EDA software to the limit.

2. Lithography

We are still using 193nm optical lithography to manufacture 32/28nm chips. So, our ruler is currently about one order of magnitude bigger than the objects that we are trying to draw.

At 32/28nm, double patterning may be required, at least for some layers. Double patterning – and moving forward multi-patterning – has a profound impact on the rules of our game. The path from GDSII to masks and silicon becomes more complex and costly. Designs and designers also need to account for double patterning in the layout, the routing rules, the design rule checks (DRC) and so on.

Besides double patterning, another lithography challenge has arisen. The semiconductor industry introduced strained silicon some time ago to improve transistor performance: compressive strain at 90nm and tensile strain at 65nm.

Unfortunately, the strain propagates itself within a certain area – across a radius of about 2um – which strains things that were not meant to be strained. Such effects, due to the proximity of different types of transistors, have become visible, both in variations of timing and static power. These variations, typically 5% at 45/40nm and in the 10% range at 32/28nm, must be accounted for.

3. Static Power and 4. Power Density

There are two reasons why power density keeps doubling from one technology node to the next:

(i) We have already reaped the benefit from using high-k dielectric and metal gates, and static power is still increasing, both in relative terms compared to dynamic power and in absolute terms. Design and EDA are increasingly being relied upon in order to keep static power below the 50% threshold of total power consumption. The problem is further worsened by the progressive integration of multiple, heterogeneous functions – baseband, connectivity, multimedia, etc. – onto a single process technology platform: e.g. LP or HP, exclusive or, with differences in static power of up to 100x from one another.

(ii) The benefits of using copper interconnect were realized when the industry introduced it a few technology nodes ago. Today, approximately 50% of dynamic power consumption is due to interconnects' RC. Designers predict that this percentage will increase. Global interconnect length doesn't scale with smaller transistors and local wires. Chip size remains relatively constant because chip function continues to increase. RC delay is increasing exponentially. At 32/28nm, RC delay in 1mm global wire at minimum pitch is 25x higher than intrinsic delay of a 2-input NAND fanout of 5.

Figure 2: Moore's Law at work between 1996 and 2006
[Source: M. Taliercio, STMicroelectronics 2001; L. Bosson, STMicroelectronics 2006]

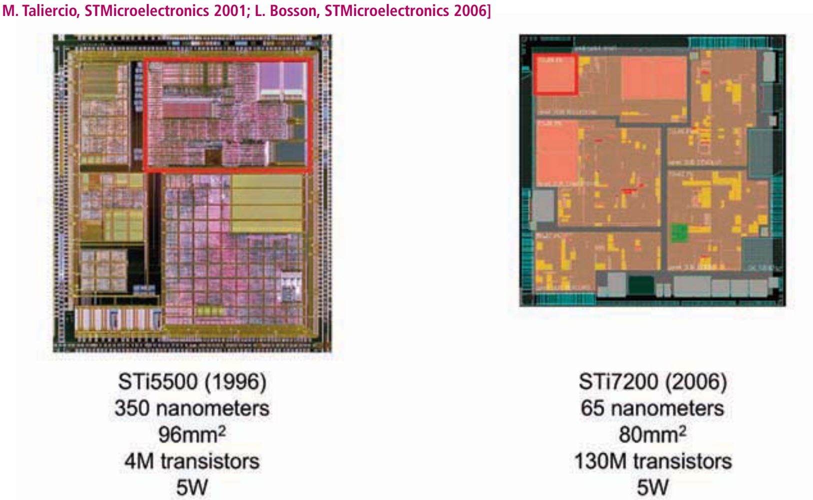

Figure 3: Atomic uncertainties lead to increasing power and timing variability
[Source: C. Kim, University of Minnesota, DAC 2007]

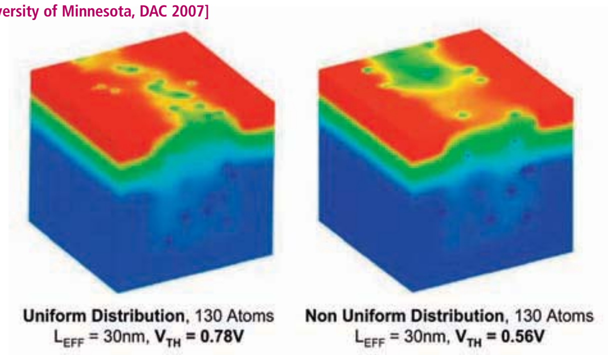

5. Reliability

Reliability issues are on the rise. Many chips have multiple functional modes to shut down and power up major functional blocks as part of the power management scheme. This approach has implications for the power and ground grid. As designers strive to achieve even lower power levels, they will have to deal with even more complex scenarios during physical implementation.

In modern electronic systems there is significant potential for interference between applications. Consider a car: it uses a variety of electronics operating at frequencies from 100MHz to 2.5GHz in the power train, car body, safety, communications, multimedia and navigation sub-systems. Reliability is an issue that now reaches far beyond simple IR-drop analysis. Designers must analyze and fix EM emissions during the design process, not after.

6. Physical Verification

The number of design rules doubles with each new technology node, while the complexity of those rules increases at a rate of 5x. The net result is that each new design rules manual contains three times as many pages as the last, and that each rule requires 3x to 4x more DRC steps. This has a big impact on process development and design and layout complexity and, of course, impacts verification accuracy and runtime.

Any attempt to exhaustively describe the design rules for the router dramatically worsens its runtime. On the other hand, postponing design rule checks until after routing requires huge numbers of errors to be fixed outside of the design environment, with uncertain effects on timing and, possibly, power.

7. Variability and Uncertainty

Nanometer transistors now consist of very small numbers of atoms. The atomic uncertainty, both in dimensions (or how many) and nature (of what kind), leads to increasing timing and power variability. Figure 3 shows the differences between two transistors that arise because of a variation in the distribution of just 130 atoms of dopants, which translates to a 30% variation in VTH.

Figure 4: Computational requirements exceed Moore's Law
[Source: CA Malachowsky, co-founder, NVIDIA. EDPS 2009]

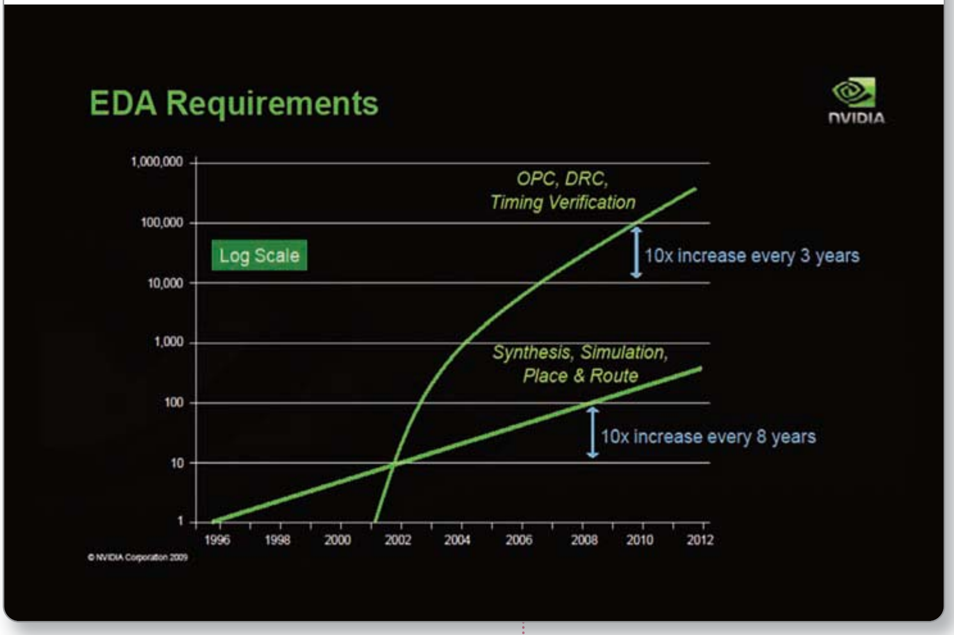

Random variations, such as random dopant fluctuation and edge roughness, are on the rise. These make optimisation more complex and signoff more difficult to achieve. VDD, VTH and LEFF variations now account for 85% of timing and 80% of power variability. New models, current-based, and techniques are required for optimization and signoff.

8. Test

The number of defects, and therefore faults, is increasing. So too is the number of test patterns required to guarantee coverage and quality. The compression ratio needed to maintain the status quo is getting very high.

During a given process technology node development, we strive to model what we see and what we want to use in the design flow: SPICE models, design rules and so on. However, the information available during this phase is abstract because we use process qualification vehicles and not real applications or products. It is limited because we are bound by both cost and time considerations and it is unstable because the process technology is still maturing.

Over the life of the process technology node, especially during the volume production phase of its lifespan, thanks to metrology and test we accumulate priceless information which we use today only in a very limited manner.

The industry is used to DFT, DFM and test for manufacture (TFM). Perhaps test for design (TFD) is next. We should analyse the data that we accumulate through test and metrology in order to generate new, better models for use in the design flow. These models may be very different from the initial ones, since they reflect the mature, stable phase of the process technology lifespan, and therefore may lead to better designs, with higher manufacturing yield.

9. Computing

Unfortunately the CPU runtime required to perform a certain task increases at a pace that exceeds even Moore's Law, see Figure 4.

There are two categories of EDA tools:

(i) Design implementation and functional verification, whose computational complexity has risen by approximately 100x as we have scaled from 350 to 32nm.

(ii) Signoff and mask data preparation, whose computational complexity has risen by approximately 100,000x.

While we can explain the first number using Moore's Law (seven technology nodes, 27 = 128 or about 100x), we cannot explain the second. It is clearly a drawback of nanometer process technology.

10. "Bandwidth"

Besides the specific nanometer technology challenges we have described above, there are also more "macro" challenges. For example, multi- and many-core architectures require the bandwidth to increase at the same pace as Moore's Law. 2D technology is running out of steam.

Even assuming an infinite scaling, interconnect is a big challenge moving forward and an impediment to better performance and lower power. Multicore devices require communications bandwidth to increase in line with Moore's Law. A state-of-the-art multicore device requires a bus bandwidth of 100GBps. An off-chip memory solution can deliver a bandwidth of up to 100GBps. A 2D planar multi-chip package solution can support around 100 to 200GBps, so multicore is already pushing the limits of 2D packaging. Many-core will need several hundreds of GBps. It is apparent that "more of Moore" requires "more than Moore".

> Reliability analysis and physical verification are becoming an integral part of design implementation flow

Complexity Remains Number 1 Issue

The sheer complexity of design remains the biggest challenge at nanometer process nodes.

While waiting for EUV, or higher refraction index fluids, we have to live with what we've got, i.e. 193nm immersion (water) lithography. EDA can help a lot in making sure that GDSII is as lithography-friendly as possible and that the life of the mask data preparation tools is reasonably acceptable. Stronger methodologies and more automated flows are helping designers tackle chip power problems. EDA can also help engineers to model and compute variability in order to reduce uncertainty.

Over the last 50 years we have successfully implemented design for manufacturability and design for testability; now the challenge is to exploit the wealth of information that test generates, in order to improve design and manufacturability, i.e. yield.

Reliability analysis and physical verification are becoming an integral part of design implementation flow. This will ensure that problems are dealt with as they arise and not as an afterthought to design.

Multi-everything is an increasingly important part of our lives. We can still dream about processors running at 10GHz or more, but the reality is 3 to 4GHz cores, software tools and Tcl scripts, which we must adapt to the multicore, many-core era.

Last, but not least, 3D IC can keep Moore's Law alive and well longer than the current 2D technologies. 3D technology is coming, possibly sooner than one would expect. ●

JAVIER ORENSANZ FROM ARM EXPLORES THE MOST WIDELY USED PERFORMANCE AND POWER PROFILING METHODOLOGIES AND THEIR APPLICATION TO THE DIFFERENT STAGES IN THE DESIGN OF NEW LINUX AND ANDROID PRODUCTS

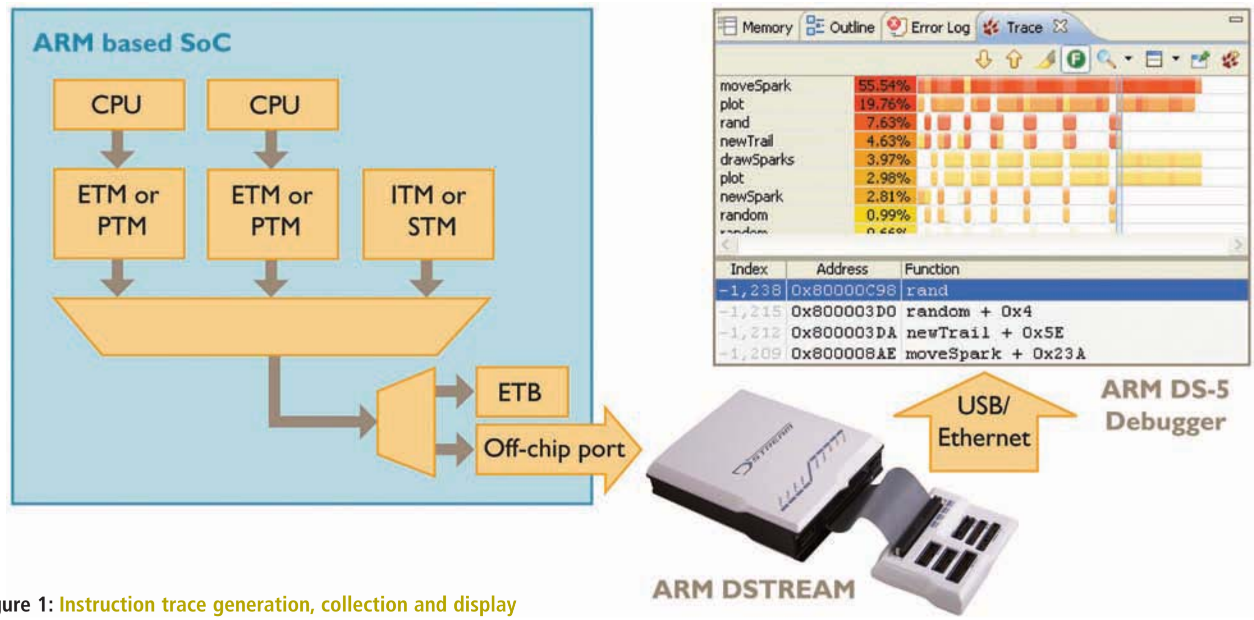

Figure 1: Instruction trace generation, collection and display

PERFORMANCE ANALYSIS ON ARM EMBEDDED LINUX AND ANDROID SYSTEMS

In the highly competitive market for smartphones, tablets and mobile Internet devices, the success of new products depends strongly on high performance, responsive software and long battery life.

In the PC era it was acceptable to achieve high performance by clocking the hardware at faster frequencies. However, this does not work in a world in which users expect to always stay connected. The only way to deliver high performance while keeping a long battery life is to make the product more efficient.

On the hardware side the need for efficiency has pushed the use of lower silicon geometries and SoC integration. On the software side performance analysis needs to become an integral part of the design flow.

Processor Instruction Trace

Most Linux-capable ARM processor-based chipsets include either a CoreSight Embedded Trace Macrocell (ETM) or a Program Trace Macrocell (PTM).

The ETM and PTM generate a compressed trace of every instruction executed by the processor, which is stored on an on-chip Embedded Trace Buffer (ETB) or an external trace port analyzer. Software debuggers can import this trace to reconstruct a list of instructions and create a profiling report. For example, DS-5 Debugger can collect 4GB of instruction trace via the ARM DSTREAM target connection unit and display a time-based function heat map.

Instruction trace is potentially very useful for performance analysis, as it is 100% non-intrusive and provides information at the finest possible granularity. For instance, with instruction trace you can measure accurately the time lag between two instructions. Unfortunately, trace has some practical limitations.

The first limitation is commercial. The number of processors on a single SoC is growing and they are clocked at increasingly high frequencies, which results in higher bandwidth requirements on the CoreSight trace system and wider, more expensive, off-chip trace ports. The only sustainable solution for systems running at full speed is to trace to an internal buffer, which limits the capture to less than 1ms. This is not enough to generate profiling data for a full software task such as a phone call.

The second limitation is practical. Linux and Android are complex multi-layered systems, and it is difficult to find events of interest in an instruction trace stream. Trace search utilities help in this area, but navigating 4GB of compressed data is still very time-consuming.

The third limitation is technical. The debugger needs to know which application is running on the target and at which address it is loaded in order to decompress the trace stream. Today's devices do not have the infrastructure to synchronize the trace stream with kernel context-switch information, which means that it is not possible to capture and decompress non-intrusively a full trace stream through context switches.

Sample-Based Profiling

For performance analysis over long periods of time sample-based analysis offers a very good compromise of low intrusiveness, low price and accuracy. A popular Linux sample-based profiling tool is Oprofile.

Sample-based tools make use of a timer interrupt to stop the processor at regular intervals and capture the current value of the program counter in order to generate profiling reports. For example, Oprofile can use this information to display the processor time spent on each process, thread, function or line of source code. This enables developers to easily spot hot areas of code.

At a slightly higher level of intrusiveness, sample-based profilers can also unwind the call stack at every sample to generate a call-path report. This report shows how much time the processor has spent on each call path, enabling different optimizations such as manual function inlining.

Sample-based profilers do not require a JTAG debug probe or a trace port analyzer and are therefore much lower cost than instruction trace-based profilers. On the downside they cause a target slow-down of between 5% and 10%, depending on how much information is captured on every sample.

> *Sample-based tools make use of a timer interrupt to stop the processor at regular intervals and capture the current value of the program counter in order to generate profiling reports*

It is important to note that sample-based profilers do not deliver "perfect data" but "statistically relevant data", as the profiler works on samples instead of on every single instruction. Because of this, profiling data for hot functions is very accurate, but profiling data for the rest of the code is not accurate. This is not normally an issue, as developers are mostly interested in the hot code.

A final limitation of sample-based profilers is related to the analysis of short, critical sequences of code. The profiler will tell you how much processor time is spent on that code. However, only instruction trace can provide the detail on the sequence in which instructions are executed and how much time each instruction requires.

Logging and Kernel Traces

Logging or annotation is a traditional way to analyze the performance of a system. In its simplest form, logging relies on the developer adding print statements in different places in the code, each with a timestamp. The resulting log file shows how long each piece of code took to execute.

This methodology is simple and cheap. Its major drawback is that in order to measure a different part of the code you need to instrument it and rebuild it. Depending on the size of the application this can be very time consuming. For example, many companies only rebuild their software stacks overnight.

The Linux kernel provides the infrastructure for a more advanced form of logging called "tracing". Tracing is used to automatically record a high number of system-level events such as IRQs, system calls, scheduling and event application-specific events. Lately, the kernel has been extended to also provide access to the processor's performance counters, which contain hardware-related information such as cache usage or number of instructions executed by the processor.

Kernel trace enables you to analyze performance in two ways. First, you can use it to check whether some events are happening more often than expected. For example, it can be used to detect that an application is making the same system call several times when only one is required. Secondly, it can be used to measure the latency between two events and compare it with your expectations or previous runs.

Since kernel trace is implemented in a fairly non-intrusive way, it is very widely used by the Linux community, using tools such as perf, ftrace or LTTng (**http://lttng.org**). A new Linux development will enable events to be "printed" to a CoreSight Instrumentation Trace Macrocell (ITM) or System Trace Macrocell (STM) in order to reduce intrusiveness further and provide a better synchronization of events with instruction trace.

Combining Sampling with Kernel Trace

Open source tools such as perf and commercial tools such as the ARM Streamline performance analyzer combine the functionality of a sample-based profiler with kernel trace data and processor performance counters, providing high-level visibility of how applications make use of the kernel and system-level resources.

For example, Streamline can display processor and kernel counters over time, synchronized to threads, processes and the samples collected, all in a single timeline view. For example, this information can be used to quickly spot which application is thrashing the cache memories or creating a burst in network usage.

Instrumentation-Based Profiling

Instrumentation completes the pictures of performance analysis methodologies. Instrumented software can log every function – or potentially every instruction – entry and exit to generate profiling or code coverage reports. This is achieved by instrumenting, or automatically modifying, the software itself.

The advantage of instrumentation over sample-based profiling is that it gives information about every function call instead of only a sample of them. Its disadvantage is that it is very intrusive, and

Figure 2: Streamline timeline view

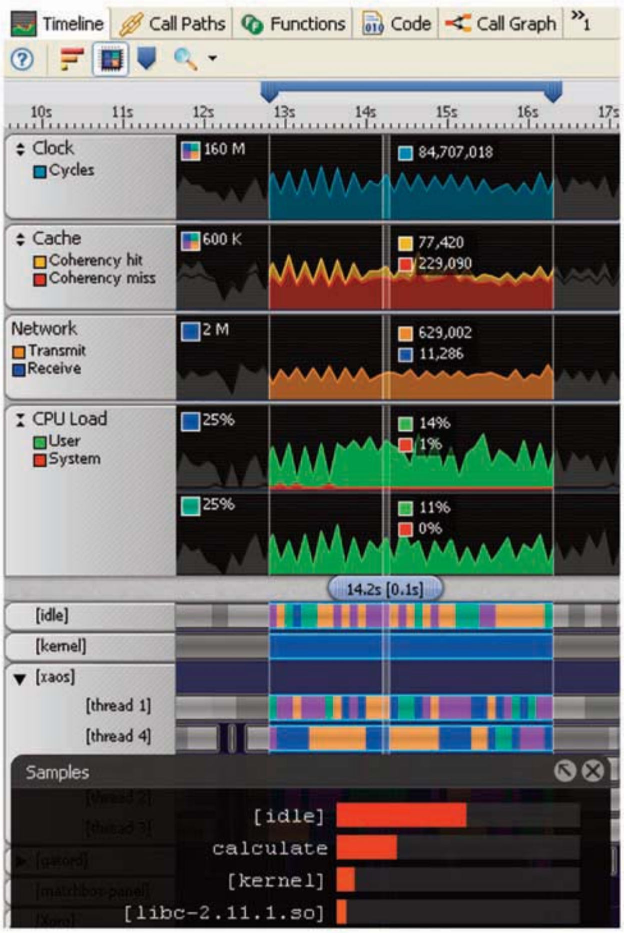

FIND OUT MORE ABOUT ANDROID NDK FOR ARM

FOR ANDROID NATIVE DEVELOPMENT ON ARM PROCESSOR-BASED SYSTEMS

Android provides the Native Development Kit (NDK) (http://developer.android.com/sdk/ndk/index.html). ARM offers DS-5 as its professional software tool-chain for both Linux and Android native development.

may cause substantial slow-down.

The Android TraceView software uses instrumentation to generate time-based logs and profiling reports for Android Java applications. A major issue of TraceView is that in order to trace the execution of each line of Java code it needs to disable the DalVik JIT machine. The resulting interpreted code can be about 30 times slower than the original just-in-time compiled code.

Using the Right Tool for the Job

All of the techniques described so far may apply to all stages of a typical software design cycle. However, some are more appropriate than others at each stage.

Instruction trace is mostly useful for kernel and driver development, but has limited use for Linux application and Android native development, and virtually no use for Android Java application development.

Performance improvements in kernel space are often in time-critical code handling the interaction between kernel, threads and peripherals. Improving this code requires the high accuracy and granularity, and low intrusiveness of instruction trace.

Secondly, kernel developers have enough control of the whole system to do something about it. For example, they can slow down the processors to transmit trace over a narrow trace port, or they can hand-craft the complete software stack for a fast peripheral. However, as you move into application space, developers do not need the accuracy and granularity of instruction trace, as the performance increase achieved by software tweaks can easily be lost by random kernel and driver behaviour totally outside of his control.

In the application space, engineering efficiency and system visibility are much more useful than perfect profiling information. The developer needs to find quickly which bits of code to optimize, and measure accurately the time between events, but can accept a 5% slow-down in the code.

System visibility is extremely important in both kernel and application space, as it enables developers to quickly find and remove any issues such as misuse of cache memories, processors and peripherals not being turned off, inefficient access to the file system or deadlocks between threads or applications. Solving a system-related issue has the potential to increase the total performance of the system ten times more than spending days or weeks writing optimal code for an application in isolation. Because of this, analysis tools combining sample-based profiling and kernel trace will continue to dominate Linux performance analysis, especially at application level.

Instrumentation-based profiling is the weakest performance analysis technique because of its high level of intrusiveness. Optimizing Android Java applications has better chances of success by using manual logging than open-source tools.

High-Performance Android Systems

Most Android applications are developed at Java level in order to achieve platform portability. Unfortunately, the performance of the Java code has a random component, as it is affected by the JIT compiler. This makes both performance analysis and optimization difficult.

In any case, the only way to guarantee that an Android application will be fast and power-efficient is to write it – or at least parts of it – in native C/C++ code. Research shows that native applications run between five and 20 times faster than equivalent Java applications; in fact, most popular Android apps for gaming, video or audio are written in C/C++. ●

	Low Cost	Low Intrusiveness	Accuracy	Granularity	System Visibility
Logging	•••	•••	•••••	•	••
Kernel trace	•••••	••••	•••••	•••	•••
Instruction trace	•	•••••	•••••	•••••	•
Sample-based	••••	•••	•••	••	••••
Instrumentation	•••••	•	•••••	••••	•

Table 1: Comparison of methodologies

www.stewart-of-reading.co.uk

Check out our website, 1,000's of items in stock.

HP8563E SPECTRUM ANALYSER
9KHZ - 26.5GHZ Synthesised........ £POA

HP33120A FUNCTION GENERATOR
100 MicroHZ – 15MHZ
Unused Boxed £595
Used, No Moulding, No Handle £395

MARCONI 2955 RADIO COMMUNICATION TEST SET
ONLY £625
Also available Marconi 2955A & 2955B

ENI 3200L RF POWER AMPLIFIER
250KHZ-150MHZ 200W 55Db

AGILENT E4402B Spectrum Analyser 100HZ – 3GHZ with Option 1DN Tracking Gen; 1 DR Narrow Res; A4H GPIB, UKB..£5800
HP 35670A FFT Dynamic Signal Analyser 2 Channel. Unused in original box...£4000
AGILENT 83752B Synthesised Sweeper 0.01-20GHZ...............................£6000
HP83711B Synthesised 1-20GHZ with Opt IEI Attenuator........................£5000
AGILENT/HP E4431B Signal Generator 250KHZ-2GHZ Digital Modulation...£2750
MARCONI 2024 Signal Generator 9KHZ-2.4GHZ Opt 04..........................£1250
MARCONI/IFR 2030 Signal Generator 10KHZ-1.35 GHZ £995
MARCONI 2022E Synthesised AM/FM Signal Generator 10KHZ-1.01GHZ ...£500
HP8566A Spectrum Analyser 100HZ-22GHZ....................................£1950
HP8568A Spectrum Analyser 100HZ-1500MHZ..................................£1250
AVCOM PSA-37D Spectrum Analyser 1MHZ-4.2GHZ.................................£-
IFR 1200S Service Communication Monitor......................................£1500
HP6624A Power Supply 0-20V 0-2A Twice, 0-7V 0-5A; 0-50V 0.8A Special price................................350
AVO/MEGGAR FT6/12 AC/DC breakdown tester...................................£-
MARCONI/IFR/AEROFLEX 2025 Signal Gen 9KHZ—2.51GHZ Opt 04 High Stab Opt 11 High Power etc As New.......£2500
SOLARTRON 1250 Frequency Response Analyser 10uHZ-65KHZ.................£995
HP3324A Synthesised Function Generator 21MHZ..........................£500
HP41800A Active Probe 5HZ-500MHZ £750
ANRITSU MS2601A Spectrum Analyser 10KHZ-2.2GHZ 50ohm..................£750

AGILENT E4421B 250KHZ-3GHZ Signal Generator £2500

HP53131A Universal Counter Opt 001
Unused Boxed 3GHZ..................£850
Unused Boxed 225MHZ...............£595
Used 225MHZ...........................£495
HP8569B Spectrum Analyser 0.01-22GHZ...................................£1000
HP54616C Oscilloscope Dual Trace 500MHZ 2GS/S Colour..............£1250
QUART LOCK 10A-R Rubidium Frequency Standard.................£1000
PENDULUM CNT90 Timer/Counter /Analyser 20GHZ.....................£1950
ADVANTEST R3465 Spectrum Analyser 9KHZ-8GHZ......................£-
HP Programmable Attenuators £300 each
33320H DC-18GHZ 11db
33321G DC-18GHZ 70db
Many others available
AGILENT E3610A Power Supply 0-8v 0-3A/0-15v 0-2A Unused
AGILENT E3611A Power Supply 0-20V 0-1.5A/0-35V 0-0.85V Unused
CIRRUS CRL254 Sound Level Meter with Calibrator............................£95
CEL328 Digital Sound Level Meter with CEL284/2 Acoustical Calibrator..........
CEL 269 Digital Sound Level Meter with CEL282 Acoustical Calibrator

FLUKE SCOPEMETERS
99B Series II 2Ch 100MHZ 5GS/G from £325
97 2Ch 50MHZ 25MS/S from £225

STEWART of READING

17A King Street, Mortimer,
Near Reading RG7 3RS
Telephone: 0118 933 1111
Fax: 0118 933 2375
9am – 5pm Monday - Friday

Used Equipment – GUARANTEED
Prices plus Carriage and VAT

Please check availability before ordering or CALLING IN

MARCONI 2945 RADIO COMMUNICATION TEST SET with....
Opt 01 – 600 ohm Matching Unit
Opt 03 – High Stability OCXO
Opt 06 – Memory Card Drive with Real Time Clock **Opt 08** – SSB Demodulator
Opt 21 Demodulation Filters
Opt 22 POCSAG Decode
Only £3,000

HP 8569B
Spectrum Analyser 0.01-22GHZ
£995

HP6269B Power Supply
0-40V 0-50A £400

AMPLIFIER RESEARCH
POWER AMPLIFIER 1000LAM8 £POA

Designed in the UK, Made in the UK.

Tel. 01298 70012
www.peakelec.co.uk
sales@peakelec.co.uk

West Road House
West Road
Buxton
Derbyshire
SK17 6HF

PEAK® electronic design ltd

Atlas LCR - Passive Component Analyser

The Atlas LCR (Model LCR40) is now supplied with our new premium quality 2mm plugs and sockets to allow for greater testing flexibility. Supplied with 2mm compatible hook probes as standard, others available as an option.

Test inductors (from 1uH to 10H), capacitors (1pF-10,000uF) and resistors (1Ω to 2MΩ). Auto-range and auto component selection.

Basic accuracy of 1.5%. Battery and user guide included.

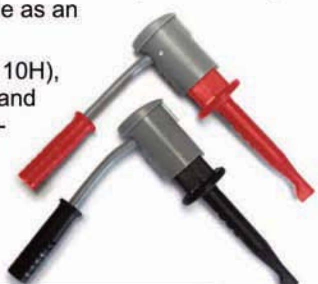

Optional Probes

£67.50 + VAT
£81 inc UK VAT
Please add £3 for UK P&P

USB 2.0 PORT SHARING FOR HIGH FUNCTIONALITY SMARTPHONES

OSCAR FREITAS FROM FAIRCHILD SEMICONDUCTOR IN THIS ARTICLE FOCUSES ON SOME OF THE SOLUTIONS THAT CAN BE APPLIED TO A DESIGN OF A SMARTPHONE WHERE MANY PROCESSING ICS NEED TO SHARE A SINGLE USB PORT AND COMMUNICATE AT HIGH SPEED

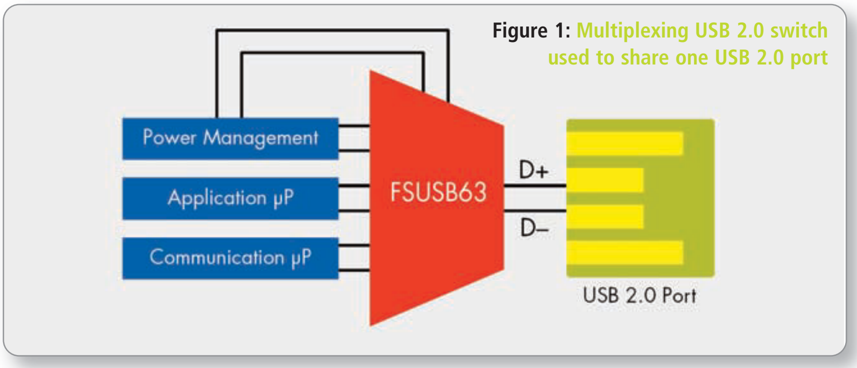

Figure 1: Multiplexing USB 2.0 switch used to share one USB 2.0 port

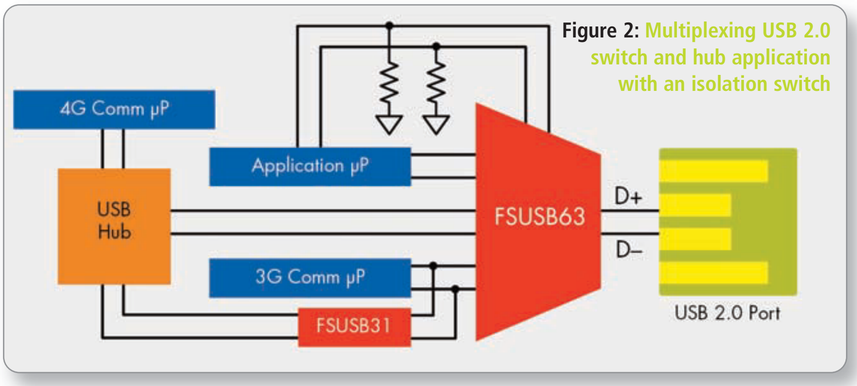

Figure 2: Multiplexing USB 2.0 switch and hub application with an isolation switch

Mobile devices need many processing integrated circuits (IC) to handle the diverse functionality demands of their users. A typical smartphone has a communication processor, an application processor and a power management IC, all needing to share a single USB port and to communicate at high speed USB data rates of 480Mbps. This article focuses on some of the solutions that can be applied to this problem by comparing and contrasting different solutions ranging from a USB hub to a simple analogue switch.

Designing Smartphones

One way to design a smartphone would be to let the power management device within the phone control the single USB 2.0 port on the phone. This is accomplished by using a 3:1 multiplexing USB switch which diverts the USB 2.0 port to itself. By default, this also diverts to the application processor used for most multi-media functionality like MP3 playing or video processing. It also diverts to the communication processor used for radio communication for either data access or for phone calls (see Figure 1).

This architecture has the advantage of allowing the handset to be in sleep when the functionality is not used. Additionally, the power management unit can wake up the relevant processor when the USB 2.0 port activity is detected or when either processor needs to use the USB port. Upon a USB plug's first insertion, the power management IC can also interrogate the USB lines to determine whether a dedicated USB charger or a charging host port has been attached so that the battery can be charged from the VBUS signal directly. The physical layer (PHY) within the communication and application processors utilizes the full 480Mbps high-speed data bandwidth of the USB switch when communicating with a USB host like a PC.

A recent trend for 4G phones requires two processing elements within handsets to concurrently access the USB port. For this application, a hub may be a better option. The dilemma is that USB hub connections are usually very capacitive and have much higher power but the need to concurrently access the USB port is infrequent.

One way to design these handsets is shown in Figure 2, where the 3G communication processor needs to access the port separately without the 4G processor. With an isolation switch (FSUSB31) and very short PCB trace stubs, the capacitance can be minimized such that a high-speed transmit USB data eye, in the path from the 3G communication processor to the USB host, has plenty of margin to the USB specification. In this example, the application processor controls the USB 3:1 multiplexing switch and it can have pull-up and pull-down resistors on its control lines to default to the application processor when the latter is powered-down in standby mode.

The functionality above is made possible by switches that take very little power when powered-up and almost zero power when disabled. In the 4G handset application shown in Figure 2, the FSUSB63 will always be powered-up and consumes just a few microamperes of current. Meanwhile, in the application shown in Figure 1, when the power management IC can shut off the FSUSB63 in standby mode, the current consumption drops to under a microampere.

High-speed and low-power consumption are typically opposing features in most integrated circuits. Often, the solution to achieve both is to lower the voltage and use fine geometry processes, however, the USB specification calls for high voltage signals. The design of a robust handset where D+ and D– signals can survive a short circuit to the 5V VBUS signal prevents low voltage solutions. However, more recently, low power, charge pumped USB switches have been providing the answer and meeting the demanding USB transmit eye diagrams as shown in Figure 3. In fact, most of these designs can run over twice the high speed USB data rates (> 1Gbps) without any changes made to either the architecture or the chosen devices.

USB Switching

Often, switching from one processor to another via a USB switch is done by first disconnecting all paths, then allowing sufficient time to elapse – ensuring the USB host port controller can recognize a disconnection and switch to the other path; thereby allowing the host to reset and re-enumerate the new USB device.

This is done in software on the handset and requires the attention of software engineers who write software based on the constraints of the hardware chosen. However, newer USB switches allow for greater use of portable device software – making this situation obsolete. This is because any change in the selection control signals of these USB switches automatically starts the process of breaking the connection for a predetermined time. This is governed by the USB 2.0 high-speed specifications for a USB host to recognize the disconnection and then makes the new connection henceforth internally.

As mentioned earlier, power conservation is paramount in most portable devices and so processors' supply voltages are often scaled down to 1.2V or lower. Therefore, when interfacing with higher voltage devices whose supply is tied directly to the battery voltage, the current drawn from the battery by inputs that are at lower voltages is considerable.

One way to eliminate the current consumption is by using voltage translators. Since USB switching is an inherently slow process, these USB switches draw negligible current in such circumstances by setting the input thresholds based on the lowest supply voltage. Also, input buffers are designed to conserve power at the worst case voltage difference which makes the system designer's job painless.

Printed circuit board (PCB) space and cost are always considerations in ever smaller portable devices. Switches are now designed into tiny packages with 0.4mm pin pitch which take up so little PCB area that they far outweigh the alternatives of big proprietary connectors with specialized dongles for USB connectivity. Additionally, the cost of the switch is far lower than a USB hub. Sometimes there is a need to cascade USB switches together in applications where a side loading controller using Media Transfer Protocol allows movies to be downloaded from the host PC to the SD memory card without burdening the application processor within the smartphone. In that case, the capacitance in the OFF state presented by a USB switch is critical to making sure high speed USB traffic is feasible for the fastest possible download time. The latest USB switches' OFF capacitance is so low (2pF typical) that it allows multiple switches to be tacked on at the final stages of the design of a handset, in order to add new features to the smartphone.

Focusing on the USB Port Design

To minimize the number of connectors on a smartphone, sometimes the USB port is designed to be attached to an analogue audio headset in addition to a high speed and/or full speed USB. Special USB switches handle this need with the headset microphone signal routed through the VBUS signal of the USB connector, while the D+ and D– signals are routed to the right and left speaker signals of the headset respectively.

An example USB switch (FSA800), determines, upon initial plug-in, whether the USB port is a USB charger via the USB battery charging specification algorithm. This allows a processor to control the switch path selection based on the ID pin signal state of the USB connector. Other USB switches utilize the ID pin for detection and automatic switch configuration, and make possible a myriad of diverse and sophisticated accessories. These can include a wide range of accessories such as very specific factory testing cables that leverage existing factory test jigs for maximum test cost benefit, full featured music playing headsets that have all the remote control buttons on the headset for MP3 playback, and FM radio listening or making/receiving calls with click-and-pop free, seamless switching between uses.

High-speed USB switches have enabled enormous functionality on today's smartphones while limiting the thickness of the handset to just one tiny, micro-USB, high-speed port connector. This USB port is routed to an application processor, a communication processor, a USB hub when concurrent USB communication with multiple devices is needed, an audio driver and a whole host of other, lower speed functions, if desired. This functionality is possible because of the advanced circuitry inherent in the latest USB switches without sacrificing power, PCB die area and cost. All this while providing USB charger detection and, thus, embracing the global push towards a universal USB charger for all portable devices. ●

Figure 3: High-speed USB 2.0 transmit eye diagram

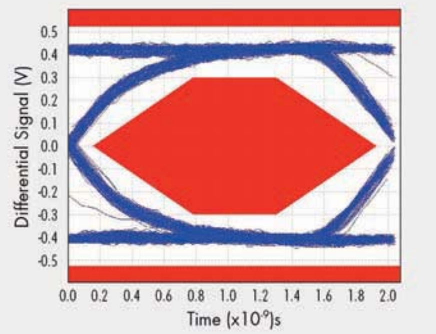

Figure 4: Audio headset accessory utilizing single USB port

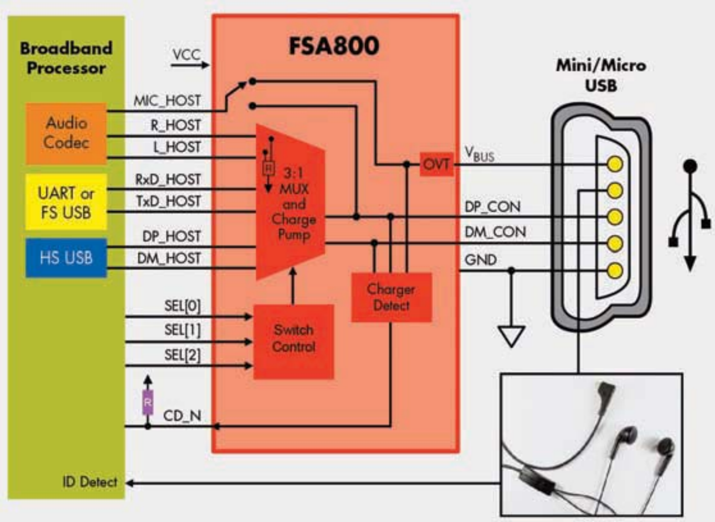

HIGH PERFORMANCE, INTEGRATED POWER ICS FOR MOBILE MICROPROCESSORS – NOT JUST FOR PORTABLES

SAM NORK, DIRECTOR AT THE BOSTON DESIGN CENTER, **JEFF MARVIN**, MANAGER AT THE BURLINGTON DESIGN CENTER AND **STEVE KNOTH**, SENIOR PRODUCT MARKETING ENGINEER AT THE POWER PRODUCTS GROUP AT LINEAR TECHNOLOGY CORPORATION EXPLAIN THE BENEFITS OF A NEW GENERATION OF PMICS FOR PORTABLE AND NON-PORTABLE DESIGNS

There is a long and growing list of power efficient microprocessors offered by Freescale, Intel, ARM and others designed to provide low power consumption and high performance processing for a wide range of wireless, embedded and networking applications. The original intent of these products was to enable OEMs to develop smaller, more cost-effective portable handheld devices with long battery life, while at the same time offering enhanced processing performance to run feature-rich multimedia applications.

Recently, demand for this same combination of high power efficiency and processing performance has spread to non-portable applications as well. Automotive infotainment systems and other embedded applications demand similar levels of power efficiency and processing horsepower as the latest high-end portable devices. In all cases, however, a highly specialized, high-performance power management companion IC is necessary to properly control and monitor the microprocessor's power system and ensure that all of the efficiency benefits possible with these processors can be realized, regardless of the application.

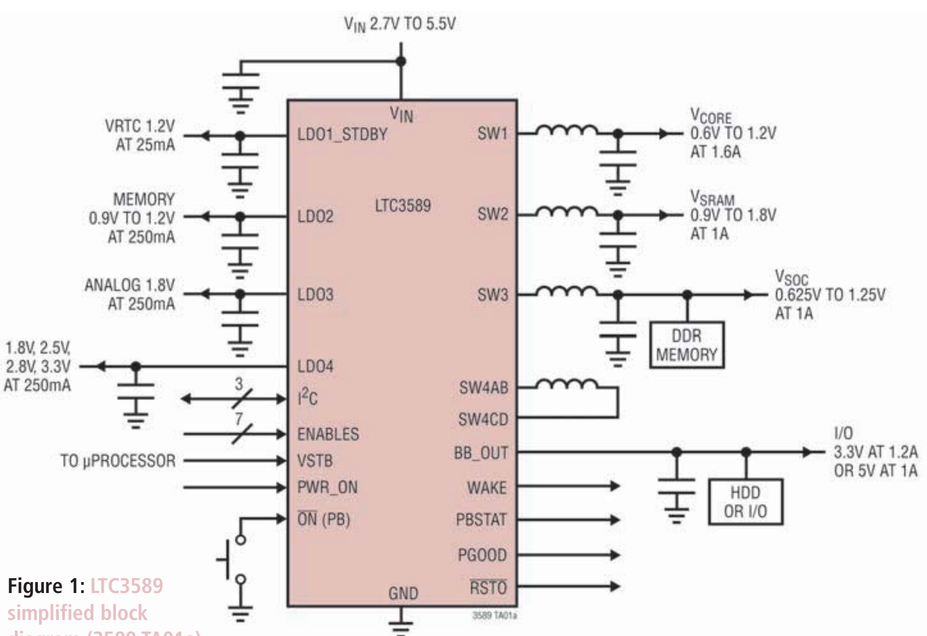

Figure 1: LTC3589 simplified block diagram (3589 TA01a)

Achieving higher processing power without increasing system power consumption requires lower voltage operation at ever increasing currents. Both portable and embedded systems include a variety of components optimized for operation at different voltages due to either the needs of the application or the line widths of the processing technology. The end result is that systems employing the latest "portable" processors require a large number of high current, low voltage rails – typically at or below 1.8V. In addition to numerous low voltage rails, many of these applications also require 3V or 3.3V rails for powering large portable hard disk drives, memory, I/O supplies for external logic circuitry and so on.

In automotive or embedded applications, all supply voltages that directly interface to the processor can be generated by high-efficiency buck DC-DCs or LDOs, depending on the current requirements.

In the case of portable applications, the main power source is typically a large single cell Li-Ion/Polymer battery which may have a voltage above or below the 3.3V system supply in the product.

Figure 2: LTC3589 LDO regulator application circuit

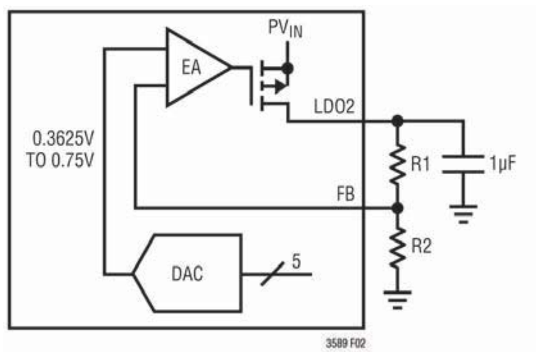

Figure 3: LTC3589 step-down switching regulator application circuit

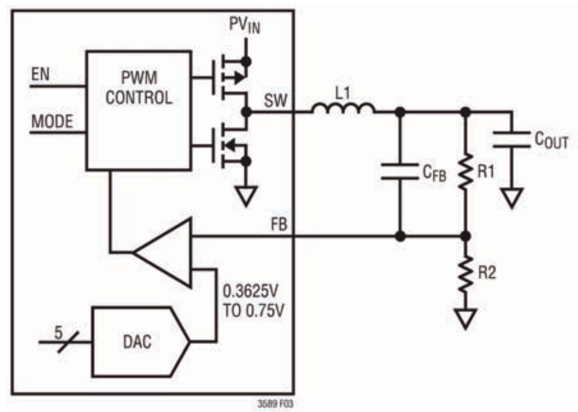

Applications such as these (e.g. handy terminals, bar code scanners, RFID readers, etc) require a buck-boost supply to generate the 3.3V rail. Additional complexity of these "portable" processor systems, whether they are battery-powered or not, include the need to sequence all supplies on and off in a specific order and the ability to be adjusted up and down dynamically depending on the processing needs of the system.

For the system designer, a single integrated solution that addresses all of the microprocessor and associated application power supply needs is a huge advantage. Handling these needs across a wide range of applications requires a highly flexible, programmable and efficient multioutput power supply solution.

Design Challenges

Most of today's modern feature-rich electronic systems still require voltage rails in the +3V range, for example to power I/O or a peripheral rail in an automotive infotainment system. Integrating synchronous buck-boost switching capability into the power management IC (PMIC) allows 3.3V regulation across the entire input voltage range, 2.7V to 5.5V, with high efficiency, resulting in increased operating margin. However, achieving high efficiency with a buck-boost design is much more challenging than a simple step-down DC-DC converter, particularly if low noise and good load step transient response are required.

The 12V car battery, the starting point for many automotive supplies, is far from the quiet, stable supply required by these systems. In addition to noise, this 12V "supply" can be subjected to reverse battery conditions or load dumps where the voltage can range or spike anywhere from -36V to 80V. The ideal supply for these systems must protect both itself and the applications circuits from these demanding electrical conditions while providing stable, low noise output voltages.

> For the system designer, a single integrated solution that addresses all of the microprocessor and associated application power supply needs is a huge advantage

Thermal conditions in automotive environments are equally challenging; PMIC junction temperatures may approach 125°C even for 85°C ambient, requiring supply stability over a -40°C to +125°C junction temperature range with robust overtemperature protection. Due to these harsh conditions, typically this system/battery voltage is pre-regulated to 3.3V or 5V prior to supplying the PMIC. In many instances, these intermediate supplies can be disrupted due to cold crank and severe noise transients. A buck-boost supply is advantageous here as well to ensure that the critical 3.3V rail associated with the processor system does not trigger a power-on reset.

Reducing Heat, Optimizing System Efficiency

Many industry-standard PMICs come with a variety of linear regulators onboard. However, linear regulators, if not managed properly with sufficient copper trace routing, heat sinks, or well-designed input/output voltage and output current levels, can generate localized thermal "hot spots" on the PC board itself. Alternatively, a switching regulator provides a more efficient way to step down voltages when the difference between input and output voltage is high and/or if the output current is large.

PMIC usage is commonplace in today's feature-rich devices with low-voltage microprocessors onboard. As a result, implementing switch-mode based power supplies for the majority of voltage rails is increasingly necessary. LDOs, however, provide low noise outputs and great power supply rejection ratio (PSRR) performance, so tradeoffs must be assessed. In many cases, the correct IC partitioning includes both DC/DC and linear regulators.

Virtually all applications today are sensitive to heat in the system. As the processing performance and associated operating currents increase, it is increasingly important to use switching regulators in place of LDOs. This is particularly true in highly integrated power supplies since single ICs are limited in their ability to dissipate power. Furthermore, achieving optimal power dissipation requires many of the core processing rails to dynamically adjust, depending on the processing operations being performed. Higher supply voltages are necessary to achieve higher clock rate

Figure 4: LTC3589 dynamic voltage scaling transition

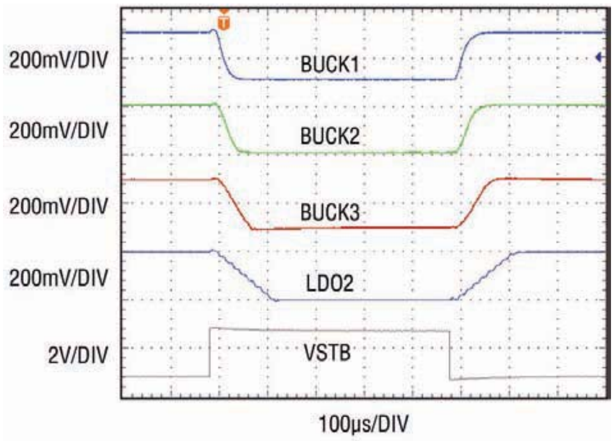

operation. Similarly, very low voltages are adequate for less processing-intensive modes of operation. Since the corresponding supply currents tend to track the input supply voltage, it is desirable to operate the processor at the lowest supply voltage possible. Dynamically adjusting the processor voltage supplies requires a serial port, such as I2C, to communicate the changes. Virtually all of today's high-end portable processors support this functionality, however, taking advantage of it requires an equally flexible and programmable power solution.

Summarizing, the main challenges for the system designer include:

- Integration of buck-boost regulators;
- Balancing power dissipation with the high level of integration of multiple switching regulators and LDOs;
- Voltage transients and temperature extremes of automotive systems;
- Integrating dynamic I2C control;
- Solution size and footprint.

A Simple Solution

In the past, PMICs have not had sufficient power to handle these modern systems and microprocessors. Any solution to satisfy the power management IC design constraints outlined above must combine a high-level of integration, including high-current switching regulators and LDOs, dynamic I2C control of key parameters with hard-to-do functional blocks such as buck-boost regulators. Further, a device with high switching frequency reduces the size of external components, and ceramic capacitors reduce output ripple. Such an IC must also be suitable for the rigorous automotive environment, although the input voltage is typically pre-regulated from the system or battery voltage.

For example, the LTC3589 is a complete power management solution for ARM-based processors and advanced portable microprocessor systems. The device contains three synchronous step-down DC/DC converters for core, memory and system on chip (SoC) rails, a synchronous buck-boost regulator for I/O at 2.5V to 5V, and three 250mA LDO regulators for low noise analogue supplies.

An I2C serial port is used to control regulator enables, output voltage levels, dynamic voltage scaling and slew rate, operating modes and status reporting. Regulator start-up is sequenced by connecting regulator outputs to enable pins in the desired order or via the I2C port. System power-on, power-off and reset functions are controlled by a pushbutton interface, pin inputs, or I2C interface. Voltage monitors and active discharge circuits guarantee a clean power-down before the next enable sequence plus selected regulators can be exempt from pushbutton control for supplies, such as memory, when the system must be kept alive during a standby mode. The LTC3589 supports i.MX, PXA and OMAP processors with eight independent rails at appropriate power levels, with dynamic control and sequencing. Other features include interface signals such as the VSTB pin that toggles between programmed run and standby output voltages on up to four rails simultaneously.

> In the past, PMICs have not had sufficient power to handle these modern systems and microprocessors

High Level of Integration

The LTC3589 is a complete power management solution for portable microprocessors and peripheral devices. It generates a total of eight voltage rails for supplying power to the processor core, SDRAM, system memory, PC cards, always-on real time clock and hard disk drive (HDD) functions. Supplying the voltage rails are an always-on low quiescent current 25mA LDO, one 1.6A and two 1A step-down regulators, a 1.2A buck-boost regulator and three 250mA low dropout linear regulators. Supporting the multiple regulators is a highly configurable power sequencing capability, dynamic voltage slewing DAC output voltage control, a pushbutton interface controller, regulator control via an I2C interface, with extensive status reporting and interrupt output.

Its internally-compensated, constant frequency current-mode step-down switching regulators provide up to 1A, 1A and 1.6A. Step-down regulator switching frequencies of 2.25MHz or 1.125MHz are independently selected, including phasing, for each step-down regulator using the I2C command registers. The power-on default frequency is 2.25MHz and includes edge rate adjustments for reduced EMI. Each of the step-down regulators have dynamically slewing DAC input references and external feedback pins to set the output voltage range. The step-down regulators' three operating modes – pulse-skipping, Burst Mode operation, or forced continuous – are set using the I2C interface. In pulse-skipping mode the regulator will support 100% duty cycle. Burst Mode operation is advantageous for best efficiency at low output loads. Forced continuous mode minimizes output voltage ripple at light loads along with optimum dynamic slew control between voltage output set points.

The single inductor, 4-switch buck-boost DC/DC voltage mode converter generates a user-programmable output voltage rail from 2.5V to 5V. Utilizing a proprietary switching algorithm, the buck-boost converter maintains high efficiency and low-noise operation with input voltages that are above, below or equal to the required output rail. The buck-boost error amplifier uses a fixed 0.8V reference and the output voltage is set by an external resistor divider. Burst Mode operation is enabled through the I2C control registers. No external compensation components are required for the buck-boost converter.

Dynamic Rail Control

The device has the I2C control features required by high-end portable applications processors, Dynamic Voltage Scaling and selectable voltage slew settings. To enable the IC's slewing DAC reference operation, the three LTC3589 step-down switching regulators and linear regulator LDO2 have programmable DAC reference inputs. Each DAC is programmable from 0.3625V to 0.75V in 12.5mV steps:

$$V_{OUT} = \left(1 + \frac{R1}{R2}\right) \bullet (0.3625 + BxDTVx \bullet 0.0125)(V)$$

R1 and R2 make up the feedback resistor divider for setting the output voltage of the regulators, see Figures 2 and 3 for details. 0.3625 is the minimum value of the 5-bit DAC reference into the error amplifier. 0.0125V is the DAC LSB step-size. BxDTVx is the binary code (0 to 31 decimal) stored in the I2C register.

The DAC references can be commanded to independently slew between two voltages at one of four selectable slew rates. Each DAC has two independent output voltage registers, voltage register select, slew rate and start controls. The regulators do not need to be enabled to change the DAC outputs.

Figure 4 shows buck regulators 1, 2 and 3, and LDO2 slewing between 0.8V and 1.2V at the four possible slew rates; the slew is initiated by the VSTB pin (grey). The levels are 8 DAC codes apart.

The versatile I2C serial port is used to control regulator enables, output voltage levels, operating modes and status reporting. The I2C serial port on the LTC3589 contains 13 command registers for controlling each of the regulators, one read-only register for monitoring each regulator's power good status, one read-only register for reading the cause of an IRQ event and one clear IRQ command register. The LTC3589 I2C port supports random addressing of any register, and registers may be written in any order using multiple START sequences. All registers can be read back to verify software and hardware integrity.

New Generation of PMICs

By replacing discrete power IC components or traditional large overly integrated PMICs (i.e. with audio, codecs, touch screen interfaces, etc), a system designer can use a new generation of compact PMICs that integrates key power management functions for improved performance with smaller and simpler solutions.

High performance mobile processors typically have a unique set of power supply requirements, including multiple high current and low noise rails, programmable sequencing and dynamic I2C adjustment.

These high-end processors were originally developed for handheld applications but are now being implemented in non-portable and embedded systems such as automotive infotainment. ●


ELECTRONICS WORLD
THE ESSENTIAL ELECTRONICS ENGINEERING MAGAZINE
Volume 117 • Issue 1897
January 2011 • £4.60
www.electronicsworld.co.uk
EXPLOSIVE END OF ELECTRIC CHARGE AS WE KNOW IT
NEW ON THE ROAD
NEW USB DESIGN PROJECTS
THEORY AND R&D
SPECIAL REPORT
TECHNOLOGY MOBILE HEALTH ADVANCES
AUDIO REVOLUTIONARY 3D FILTER
ALSO IN THIS ISSUE: THE TROUBLE WITH RF: "NO ANTENNA CABLE?"
ELECTRONICS WORLD
THE ESSENTIAL ELECTRONICS ENGINEERING MAGAZINE
Volume 117 • Issue 1898
February 2011 • £4.60
www.electronicsworld.co.uk
FOCUS ENGAGE IN THE NEW COMMS ACT
THEORY WHERE DID EINSTEIN GO WRONG?
ANALYST THE SIX MYTHS OF CHIP MAKING
LT3743
ELECTRONICS WORLD
THE ESSENTIAL ELECTRONICS ENGINEERING MAGAZINE
Volume 117 • Issue 1899
March 2011 • £4.60
www.electronicsworld.co.uk
SPECIAL REPORT: DESIGN WITH MCUS
POWER SUPPLIES CLAMP SIZING
TREND OPTICAL TRENDS FOR 2011
DESIGNSPARK
SEE WHAT YOU MIGHT SPARK
Download your free copy at www.designspark.com/pcb
EMBRACING THE NEW PARADIGM
DESIGNSPARK BY RS COMPONENTS
TECHNOLOGY ELECTRON ORBITS IN MULTILAYER GRAPHENE DISCOVERED
AUDIO TUTORIAL AUDIO TESTING WITH STANDARDS IN MIND
PRODUCTS NEW CHIPSETS, MODULES AND SCOPES
ALSO IN THIS ISSUE: USB DESIGN PROJECTS • ON THE ROAD • AUDIO COLUMN

WRITING SOFTWARE FOR THE VINCULUM-II

JOHN HYDE FROM FUTURE TECHNOLOGY DEVICES INTERNATIONAL (FTDI) WILL IMPLEMENT A SERIES OF USB PROJECTS OVER THE COURSE OF THE FOLLOWING EDITIONS OF *ELECTRONICS WORLD*

In the previous articles in this series we looked at developing simple USB-based designs using slave-based converter ICs and the Vinculum-I dual host peripheral controller. An introduction to the user-programmable Vinculum-II dual host/slave controller will enable a wider range of USB problems to be solved at lower cost. This article looks at the RTOS framework offered with the Vinculum-II which provides a huge head start for even the most sophisticated of applications.

The FTDI Vinculum-II has a 16-bit Harvard architecture CPU that controls three major buses, a 32-bit data memory accessing 16KB of RAM, a 16-bit program memory accessing 256KB of Flash, and an 8-bit peripheral bus with a wide collection of integrated devices as shown in Figure 1.

The implementation is 'hardware-rich' with silicon-assist built into every peripheral device and this means that more processing resources can be devoted to user developed applications. Rather than provide a lengthy register map with each bit explained in immense detail, FTDI provides a device driver for each of these integrated peripherals.

Figure 1: The Vinculum-II is a hardware-rich implementation

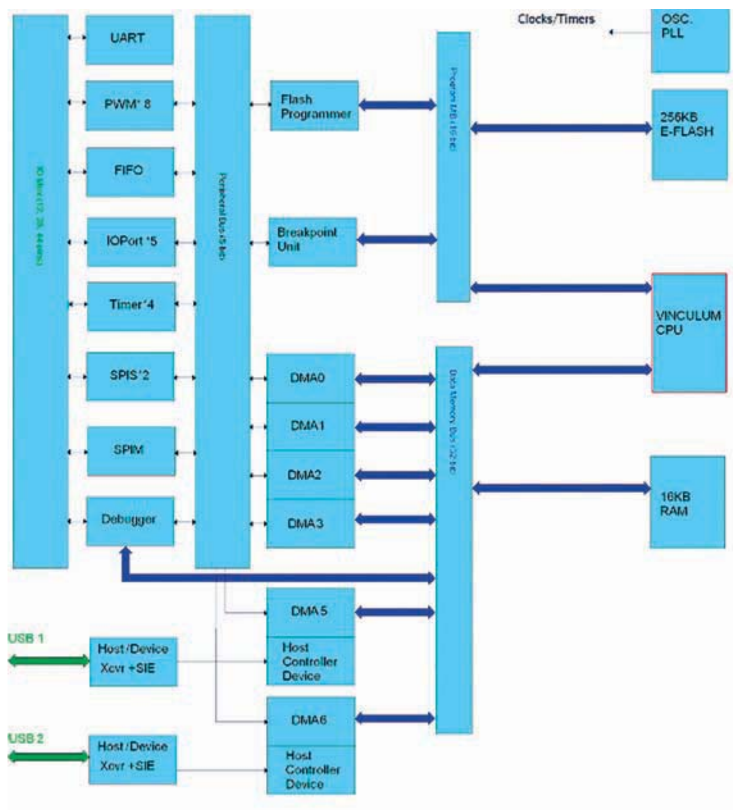

Writing Software for the Vinculum-II

Writing USB host controller software can be a daunting task since there are many timing constraints that must be accommodated. When the concurrent operation of two host/slave controllers is combined with high-speed operation of serial ports, parallel ports and timers the programming challenge becomes excessive.

To make USB host controllers easy to use, the Vinculum-II includes a real-time operating system (RTOS) that owns the hardware; this manages the low-level hardware interactions and user-level tasks act more like 'traffic cops' guiding the overall operation of the application. You don't have to deal with hardware registers – this design freedom allows greater focus on the application.

Multitasking RTOS

You have to learn some new words and concepts to be successful with a multitasking RTOS. This will take some effort so let me first explain the benefit of becoming familiar with these new terms. You probably write your code using flow charts or state machines. Flow charts are good for describing sequential processes, while state machines are good if there are small numbers of possible states with well-defined transition rules. However, both are poor at describing more complex systems with several interdependent parts.

Multitasking, on the other hand, is a good fit for such systems – you define a task to handle each part then define how the parts interact. A significant weakness of the sequential and state machine approaches is that they are inflexible. A good programmer can initially create a workable solution using these approaches but as requirements change and enhancements are demanded the workable design invariably turns into spaghetti code that is difficult to debug and, even worse, to maintain. The multitasking RTOS approach forces code that is structured so that it can grow and change easily. Changes are implemented by adding, deleting or changing some tasks while leaving other tasks unchanged. Since your code is compartmentalized into tasks, propagation of changes through the code is minimized. This will also reduce testing efforts. So, you have some hard work now to save time and effort later – this is a good deal.

The first paradigm shift you need to make is to partition your program into a set of smaller tasks – each will do one job and it will do it very well. You must also be comfortable with data structures since an RTOS will use a lot of them. Note too that task now has a specific meaning, it consists of a collection of code bytes that is the program, a

Figure 2: A program has several tasks that interact

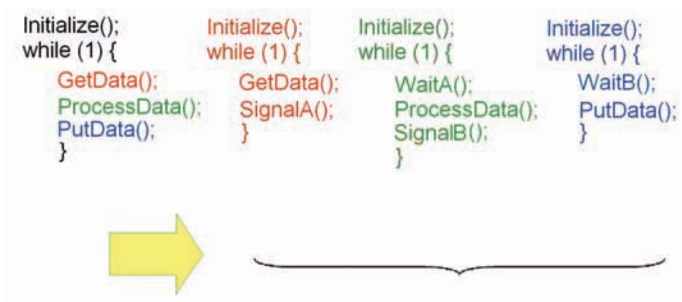

Figure 3: Tasks continuously move though this state diagram

collection of variables that are data bytes on the stack and a data structure, also on the stack, called the task context.

A 'thread' is a data structure used to describe a task and its operational status. If you have two, or more, identical peripherals (Vinculum-II has several duplicate units), you can define two threads each with the same code object but with different stack and context objects. Once your application is divided into several tasks you will define how these tasks interact. The primary inter-task communications mechanism is a 'semaphore', which is another system-defined data structure. Several operations are defined for a semaphore such as Initialize, Signal and Wait. A task that creates data will signal when it has data while a task that consumes data will wait until a semaphore is set. Figure 2 shows a simple embedded program split into multiple tasks.

Each task is written as if it has sole ownership of the CPU. You could now allocate the coding of each task to different programmers with different areas of expertise. Also if a better data processing algorithm is discovered then only one task has to be changed; you need not be concerned about the impacts to the input and output processes since they now operate independent of the processing task. Are you beginning to see some of the benefits of this "divide-and-conquer" approach?

When you divide your program into multiple tasks you will decide that some are more important than others and you can assign these a higher priority. Figure 3 shows the classical multitasking RTOS task state diagram – specific details on the Vinculum-II implementation are covered later. As tasks are created they are placed on the ReadyToRun list where the RTOS determines the highest priority task and makes this the Running task; execution of this task continues until it is blocked for some reason (waiting for a resource, such as a semaphore or a timer) when it is placed on the Waiting list; the RTOS then places the highest priority task on the ReadyToRun list as the Running Task; and so the process continues.

Vinculum-II Software Architecture

Figure 4 shows a block diagram of the layered Vinculum-II software architecture. Following a RESET the software environment for Vinculum-II must be set up; the steps taken during this initialization are part of the Kernel Services module of Figure 4 and are shown in more detail in **Figure 5**. All Vinculum-II programs implement these steps but with different data and, once initialized, the run-time diagram shown in Figure 3 describes the operation of your program.

Kernel Services

Here we are looking deeper into the kernel services initialization steps:

StartVOS: this function call initializes all of the internal data structures and sets up the operational parameters of the kernel. The Vinculum-II Operating System (VOS) needs to know how many device drivers will be used so that it can organize and set aside memory for data buffers. System timing parameters are also set using StartVOS.

ConfigureIOMux: Vinculum-II is available in three package sizes (32, 48 and 64-pin) and this function call sets up the mapping of peripheral Input and Output functions with the physical pins on the package. The number of available IO pins varies with package size and you cannot get every peripheral signal connected to the outside world on the smallest package. Note too that you should be careful not to map away the debugger pin.

InitializeDrivers: Before a driver can be used its driver must be initialized. The driver for most peripherals is small but the USB Host controller driver, for example, will set up several threads to manage the root hub, optional downstream hubs and transaction list processing.

CreateResources: A Vinculum-II program will consist of multiple independent tasks that interact with each other. During initialization the threads for each of these tasks will be created along with the semaphores, mutexes and, perhaps, shared buffers needed for intertask communications. Each thread has a context (object), and handles for shared objects (such as semaphores) may be provided as data within this context. Each thread has its own stack that is initialized with a known pattern so that VOS can track memory usage.

A 'thread' is a data structure used to describe a task and its operational status

StartScheduler: Once all of the program objects have been initialized we start the real-time operating system. The VOS scheduler uses a round-robin, priority-based, pre-emptive algorithm to run the highest priority task. It also tracks statistics such as thread CPU usage and this enables applications to be profiled and tuned if necessary.

Additional Device Drivers

File System Driver: Layered on top of the BOMS/MSC driver and the SD card driver is a FAT file system driver. It supports devices with FAT12, FAT16 or FAT32 structures and includes everything needed to simply open, read or write and then close data files. The API supports random length file reads and writes but this does cause the driver to run read modify-write cycles on the physical device. I would recommend doing your own sector buffering.

Device Manager: This provides a consistent and standard interface to the underlying on-chip peripheral device drivers and any added device drivers. The API includes Open, Close, Read, Write and IOCtl (IO Control) functions. All devices are accessed using these standard API functions so communicating over the UART is the same as

Figure 4: Vinculum-II layered software architecture

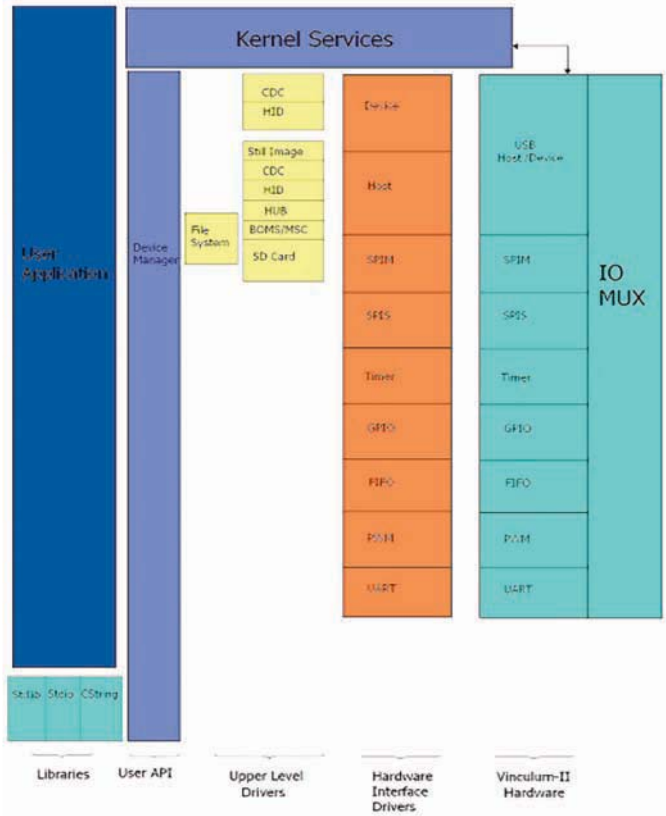

Figure 5: Software initialization steps

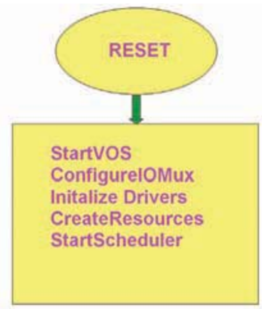

communicating over SPI, or the USB Host for that matter. This will standardize and simplify your application code and make it easier to change your hardware to match what your marketing team has sold.

Any differences between peripherals, such as setting the baud rate of the UART, are handled by the IOCtl API function. The read and write functions are used to stream data to and from devices and four DMA channels are available for user applications. Each host controller also has a DMA channel which the driver uses to move data into and out of any specified user data buffers. These read and write requests can be any length since the file system driver handles all USB packet size issues.

Using the information compiled in this article, we are now in a position to develop our own Vinculum-II based application programs. Next month we will go through a series of different examples using the Vinculum-II tool chain. ●

When Lightning STRIKES

HUW MUNCER OF TE CONNECTIVITY IS ALWAYS ON THE ROAD. AS A SALES MANAGER, HE SEES WHAT ENGINEERS ASK FOR AND NEED, BUT ALSO HOW AN ENGINEER'S JOB HAS CHANGED OVER THE YEARS

When I took over the beat of a salesman for Africa little did I realise that part of my role would be governed by a force of nature – lightning. Africa, in general – and Southern Africa in particular – is one of the world's hot spots for lightning storms, making protection against damage from lighting strikes of paramount concern for wireless telecoms.

While Africa is a notably vulnerable region, lightning presents a problem in any telecom market.

Arbitrary and random, lightning strikes with millions of volts and tens of thousands of amps. It's estimated that the temperature of a lightning strike is three times hotter than the sun's surface. Tests have found that even indirect strikes can cause surges on lines with horizontal arcing in excess of 40m.

I recently heard a story from a customer in Johannesburg who was testing a piece of equipment that had been hit more than 260 times in one year. Everyone I talk to in Africa has a lightning story, be it about network equipment or a TV in a home; but for wireless telecom carriers the threat posed by lightning presents an untenable situation.

Although, in general, lightning protection in telecoms is achieved using overvoltage protection components, telecoms require both overvoltage and overcurrent protection. The primary reason overcurrent protection is placed in a telecom circuit is to protect against interaction with AC power lines. Such protection is not an 'optional', although specific regulatory requirements depend on applicable standards.

The challenge for telecoms is selecting overvoltage protection devices that do not fail during a lightning event since telecoms must remain fully functional after the event. A lightning strike that hits close to a wireless carrier's network can induce high-voltage spikes in the system wiring. These spikes can damage sensitive electronic equipment at either end of the network. Therefore, they must be shunted to ground.

It is a common misconception among engineers I meet with that protecting against a lightning strike must be complicated. There seems to be a lack of awareness that a simple, effective device is readily available to confront this problem, namely the Gas Discharge Tube, or GDT.

GDTs help protect against overvoltage events such as lightning strikes. The main advantage of using them to help protect against damage from lightning is their ability to handle significantly more energy than any other type of component, including semiconductors and other passives. Talking amongst fellow component salesman I've found that knowledge about this technology is misunderstood within the telecom world – many of the design engineers we meet with are vaguely aware of GDTs but don't completely understand how they work.

Put very simply, GDTs contain sealed ceramic tubes filled with gas that works as an insulator since it doesn't conduct current. In the event of a lightning strike, voltage builds up across the GDT. While being hit by thousands of volts, the gas inside the tube breaks down (ionizes) – in an almost "magical" way – electricity transforms the gas into plasma. The plasma then conducts the lightning down to ground, therefore protecting the telecom equipment from potential damage.

As I've witnessed over the years, lighting presents a very real concern for telecom customers. GDTs can provide a much-needed solution in markets beset by lightning problems, not only in Africa but in any area where wireless telecoms require protection against lighting strikes. Although they can deliver this simple, essential protection, GDTs continue to remain a bit of a mystery to most engineers. To my mind, GDTs represent the "hidden gem" of lightning protection.

The underused potential of this technology for telecom protection illustrates how many solutions may go undiscovered simply because engineers aren't asking the right questions. For instance, when researching a problem on the web they can only find the right answer if they know what to look for. This shows that the components salesman, like me, still has an important role to play. With our experience in the field, we know how components used in one way might actually solve another design issue. After all, a customer's problem can be a boon for both the component salesperson and the design engineer when they work together to provide the means to solve it. ●

> *Many of the design engineers we meet with are vaguely aware of Gas Discharge Tubes but don't completely understand how they work*

Enabling CONNECTIVITY

IN THIS THREE-PART SERIES **MIKE HALL** OF MICROSOFT DISCUSSES THE CURRENT TRENDS IN HARDWARE, SOFTWARE AND EMERGING TECHNOLOGY THAT AFFECT THE EMBEDDED DEVICE DEVELOPER

Recently, I was asked to present at London's Future World Symposium, a conference focused on connected, smart and electronically enabled devices and experiences. My presentation, entitled "*Enabling Connectivity*", is also the topic of this column.

It would be very easy to focus my presentation purely on the topic of connectivity, but connectivity is really only one part of having a smart, connected, electronically enabled device. In forthcoming posts I will address other issues such as software trends, the user experience and cloud computing infrastructure; in this post I will focus on connectivity.

The connectivity story wouldn't be complete without a short hop into the Tardis to look at where the computing and connectivity story started. Interestingly, there's a blog post written by Andrew Coates called 'Standing on the Shoulders of Giants' that compares the evolution of software with changes to plumbing technology, from caveman times to modern day. This in many respects mirrors the transition from mainframe computing to smart, connected devices.

Mainframes to Devices

It's hard to imagine that just 50 years ago the cutting edge of computing technology was the mainframe computer – a small number of people queued up with punch cards to get access to limited computing resources in fixed locations.

The minicomputer was somewhat similar to the mainframe; it's true that more people had access to computing power then, but still in fixed locations with limited networking capabilities.

By the 1980s, access to computing technology was becoming very widespread thanks to desktop PCs. However, connectivity was still limited to the enterprise space or limited dial-up options. During the 1990s we really started to see widespread connectivity between people, computers and the Web. Computing as an experience was still tied to a physical location – a home or work office.

The ability to break out of a fixed computing location came with the introduction of Wi-Fi and laptop/netbook computing. Wi-Fi is simply an extension of an existing LAN infrastructure, though, which means users are constantly jumping from one island of connectivity to another, and in many cases are locked out of being able to connect by needing to pay for access.

The computing devices of the past were human-machine interaction devices. We are now seeing the next iteration of growth: both devices and connectivity are being driven by device-to-device (machine-to-machine or M2M) communication as opposed to human-device-cloud communication.

Demand for computational power and connectivity is increasingly driven by M2M communication. These are not general-purpose computing devices, but are fixed-function computing devices that are increasingly remote, mobile and movable.

Connectivity Challenges

Embedded device developers face a range of choices when figuring out which connectivity modules and operating system platforms to develop against. For M2M solution developers, this creates huge development and integration challenges if a device is to connect to the enterprise or cloud. Also, given the specific and often exacting requirements stipulated by cellular operators, this often results in developers pretty much building a custom solution for each geographical region and dealing with the integration issues it creates in the back end.

> *Embedded device developers face a range of choices when figuring out which connectivity modules and operating system platforms to develop against*

When you add on top of this the sheer time-to-market delay involved in custom development, and then the often lengthier delay caused by the testing and certification procedures that cellular operators need to go through to ensure that they can at least anticipate the behaviour of a device on their network, it's not a surprise that a lot of OEMs and enterprise customers have historically given up because the whole process was too hard, too long and too expensive.

The good news is that in the last couple of years there has been a huge upsurge in interest in the connected devices space. Operators are increasingly looking to M2M as a major driver of new revenue streams, particularly with the advent of 4G. In addition, there's more willingness and flexibility to develop new commercial structures that are needed to enable the development of solutions and operators. M2M service providers are also investing in services and billing layers to enhance their ability to manage connected devices. There's still much to do to make connectivity another component that solution developers can add to their toolbox, and even more to do to realize the full potential of connected device solutions with this basic foundational element in place.

Let's assume that we've solved the basic issue of getting connected. Our sparkling new embedded device now has a data pipe to the cloud and can communicate with other devices and cloud-hosted Web services.

There are other issues within the device that now need to be addressed. Some of these relate to silicon trends, specifically the move from single-core to multicore and how our code is going to take advantage of available processing power. There are also software trends to consider when building smart and connected devices, such as the move from low-level assembler through C/C++ to higher-level languages. In addition, the user experience both for the device shell and applications should be taken into consideration, together with cloud-computing capabilities. ●

ABOUT MIKE HALL

MIKE HALL IS A PRINCIPAL SOFTWARE ARCHITECT in the Windows Embedded Business at Microsoft, working with Windows Embedded Compact and Windows Embedded Standard.

WHAT THE READERS SAY

HAVING THE COURAGE TO PRINT

Referring to the articles '*The end of electric charge and electric current as we know them*' part one and two in your January and February 2011 editions, I would like to congratulate you on your courage in publishing this un-refereed article, the conclusions of which would demand a major paradigm change.

I find it most refreshing to be able to read about a theory in an authoritative journal such as '*Electronics World*' that has the courage to present, and call 'explosive', such a model that clearly highlights major issues in the current mainstream established view.

Now that this has been presented, one wonders if there are any 'established experts' who would also have the courage to comment on this article.

Ian Montgomery

EW Does a Great Service to Engineering

I write to urge you to continue publishing articles by Ivor Catt as his insights are of great importance to science and engineering. Many journals will not publish his papers for the simple reason that they upset too many apple carts. The issues raised by Catt need to be more widely known and discussed objectively and this can only be achieved by publication. *Electronics World* does a great service to science and engineering by publishing Catt's articles.

Steve Crothers

Providing Insight

I met Ivor in the 1970s and he and our colleague David Walton published many articles in Wireless World, various other magazines and we also published a book Digital Electronic Design. During those years we recognized the work of JA Fleming, Oliver Heaviside and others who had recognized that electric current as we know it is a secondary effect and the more fundamental energy current was the basic primitive.

Ivor has continued with this work for the past 30 years and, as we gain more insight into the theoretical flaws of basic EM theory, we can construct more accurate models which will be better suited to the challenges of electronics in the foreseeable future.

More effective modelling and understanding of electrical energy can provide insight into creating improved designs for transmitting, consuming and processing this valuable energy source.

Malcolm F. Davidson
USA

RATIONAL DEBATE

It is alarming that Ivor's clearly argued assertions on Electric Charge and Current are not debated rationally but this can be understood in that many 'academics' would be embarrassed by a change to their established understanding.

In July 1961 I joined Ferranti computer department as a new graduate development engineer. I was assigned to Ivor for the first six months of my industrial baptism.

This was extremely valuable as Ivor would always examine concepts from their basic elements.

It was a conversation at that time which lead Ivor to develop his seminal paper on metastability, published some years later.

John Raymond Dore

Failing to Notice

I first met Ivor in the early 1970s when he was working at Marconi_Elloitt avionics in Borehamwood. It was very shortly after this that I read Ivor's crosstalk paper, which became, together with some of his other writings, the starting point for our collaboration in fast digital systems.

It is truly fascinating that we both failed to notice that the superposition of modes of transmission in a 3-wire system, as experimentally demonstrated and documented in that paper, is in itself a demonstration that charge and current are the result of the energy current flow and not its cause. As Ivor summarizes in his articles: "One reason why it is illegal is that the electric currents in the right hand conductor are in opposite directions for the two modes, and classical theory says there cannot be two electric currents in opposite directions along a single conductor. However, two electromagnetic waves (or light rays) can be in the same point in space, for instance when we shine a torch at another lighted torch pointing in the opposite direction, or when we send two pulses from left and right through each other down a coaxial cable. Similarly, the Even and Odd Mode TEM Waves in our photos can coexist, but not their associated electric charges and currents."

I trust that the publication of these articles will stimulate a fruitful dialogue which will lead others to identify the discrepancies between the Maxwellian electromagnetic model and observed reality.

David Walton

IF YOU WOULD LIKE TO COMMENT
on this subject on any other that you have read on in Electronics World magazine, please write to the Editor at Svetlana.josifovska@stjohnpatrick.com

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

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

Figure 1: The general overview of the proposed system

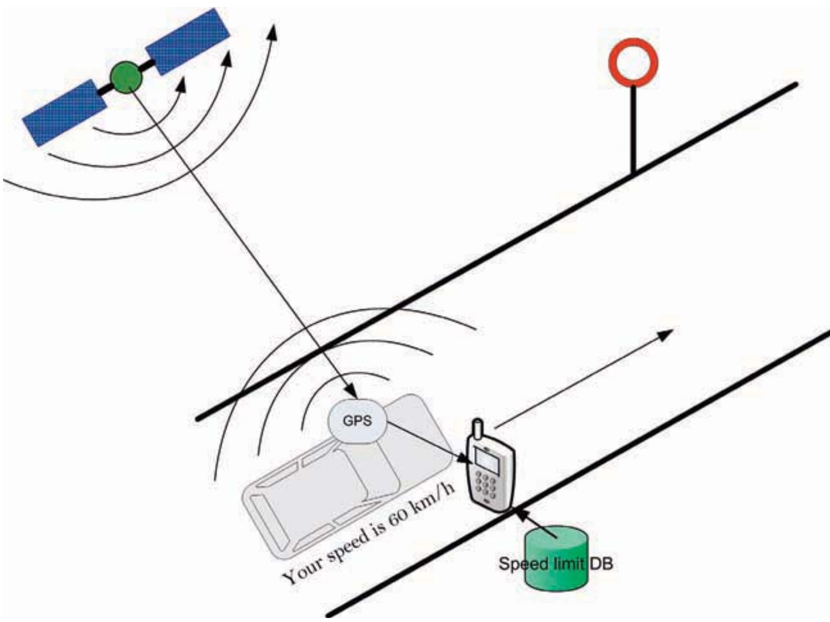

GPS-BASED DRIVER WARNING SYSTEM FOR SPEED CAMERAS WITH MOBILE PHONES

ERBUĞ ÇELEBI FROM THE DEPARTMENT OF COMPUTER ENGINEERING AT THE CYPRUS INTERNATIONAL UNIVERSITY IN NORTH CYPRUS PRESENTS THE DESIGN AND DEVELOPMENT OF A GPS-BASED SYSTEM FOR MOBILE PHONES THAT GIVES INFORMATION AND ALERTS OF SPEED CAMERAS ON THE ROAD AHEAD

It is indisputable that mobile phones with various mobile applications have become a big part of our lives. The main reason of this rapid evolution is that the mobile phones are embedded with many technologies such as WiFi, GPRS, GPS, FM transmitters and Bluetooth among others. By introducing these technologies mobile phones have been turned into smart phones enabling them to operate as small portable computers.

The aim of this article is to introduce and embed another useful feature to the mobile phones. This new feature will allow and analyse information gathered via a mobile phone and GPS receiver, and inform the driver whether there is a fixed speed camera ahead. Furthermore, this system is giving a warning alarm to the driver with sound and visual effects, if the vehicle's speed exceeds the speed limit. The design and development steps of such systems are presented in this article.

Use of a GPS-Based System

Each year many people are being injured or lose their lives in car accidents all over the world. One of the main reasons cited for this is the high speed of vehicles in a speed limited zone. So, an intelligent system which is aware of the location of speed limits and has information about its location can be very useful in order to warn the driver to control the vehicle more carefully. Also, nowadays fixed speed cameras are installed in many places and drivers who are not familiar with the road may mistakenly pass the camera with higher speed than the speed limit, which can lead to a fine. The locations of fixed speed cameras are determined from the statistics of traffic accidents, so they are located on the most dangerous roads and junctions. An autonomous system that warns the driver of fixed speed cameras on the road will decrease the accident risks.

The exact information of the location and speed can be achieved by using a global positioning system (GPS). GPS is a US space-based global navigation satellite system, providing reliable positioning, navigation and timing services to worldwide users on a continuous basis; it uses four or more GPS satellites.

Currently there are some secure driving tools in the market that inform of the location of cameras or speed limit zones. RoadPilot mobile is one such system. It includes a real-time alert facility enabling motorists to notify other members that belong to the RoadPilot user community if specific mobile fixed speed camera sites are active, i.e. if there is a police speed-checking vehicle present.

When not in the range of a safety camera, the handset screen displays an accurate GPS speed reading. The alarm function is directional, which means it does not give alerts to cameras on the opposite side of the road and, also, features auto-ranging and speed sensitivity controls in order to prevent alerts when travelling at less than the prescribed speed limit.

Road Angel Mobile (RAM) is a similar system, based on a subscription basis so that a direct connection with a data centre is established. The transmitted GPS information will be partially analysed by the mobile phone and the rest by the computers in the centre. But, since it is a managed service, with a direct link to a centre, it is an expensive one.

TrafCam is a software package which turns the mobile phone into a GPS speed camera warning. The problem with this system is that it isn't easy to retrieve and update stored information.

However, in this article, we are introducing a real-time, low cost, GPS-based secure driving system. The system gathers the latitude, longitude and speed of the moving vehicle via the GPS device. This data is transmitted to the mobile phone via a Bluetooth connection.

The coordinates, speed limit and name of specific locations (speed-limit points) are also saved in the database, available on the mobile phone, and can be retrieved and updated by connecting the mobile phone to the server. The obtained information through the GPS receiver is continuously compared to the available speed-limit points by the means of Vincenty distance, which will be explained in more detail later on. If the distance is within a predefined threshold and the speed is higher than the maximum speed limit, a warning alarm will be activated in order to inform the driver to reduce his/her speed. The system has been successfully tested in a practical scenario. Also, such a system can be embedded in vehicles, which can be more practical in terms of making the adaptive cruise control smarter.

Figure 2: Block diagram of the proposed system

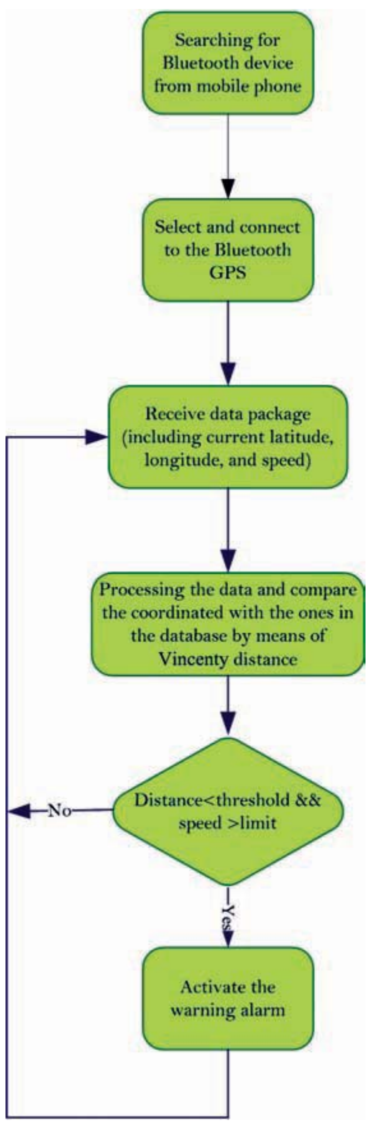

System Design

Mobile phones are widely used in many applications including mobile payment, information gathering, navigation, etc. Here we are proposing the use of a mobile phone as equipment to detect fixed speed camera on the roads. In order to design an autonomous speed control based on GPS, the following parts are required:

- A GPS receiver which supports Bluetooth communication;
- A mobile phone with the Symbian 60 operating system;
- Float library and BLUElet API;
- An IDE with J2ME support;
- Java wireless toolkit 2.2.

The system can also be implemented in a mobile phone with a built-in GPS receiver, however, in order to increase its accuracy, an individual GPS receiver will need to be attached to the dashboard of the car. There is no support for floating point numbers and operations in mobile applications that are compiled for CLDC 1.0. However CLDC 1.1 supports floating points. This application is compiled to support CLDC 1.0 for backward compatibility, so there is need for a floating point library. We have used henson.midp. Float is a library used to obtain rapid arithmetic calculation and parsing for floating numbers. BLUElet is also an API for Bluetooth developers to simplify the coding of Bluetooth applications. An overview of the general steps of the proposed system is described in Figure 2.

Implementation

There are two major ways of developing software for mobile devices which has the Symbian operating system. The first is to use C++, which is the native language for Symbian, and the second is to use Java. In order to decide which of these two languages is more productive, one should consider some issues such as performance, portability and system usage of the languages. Although both languages are object-oriented, Java approaches are better than C++, which makes programming easier. It is easier to make errors when coding in C++ than in Java because Java is a higher level language than C++ and has automatic garbage collection. Also one of the advantages of using Java is that it can be used on devices that have other operating systems rather than Symbian.

The GUI Implementation

In this section, we discuss the design of the visual components in other word Graphical User Interfaces (GUI), functions and accessibility of the application. The application is a subclass of MIDlet.

Main Window

The main window is implemented by MainMenu class that consist of visual elements and listeners. Listener events are handled by CommandListener. This window contains a menu and the StringItem in order to inform the user. User events are generated by clicking the menu items that are handled by the CommandListener.

```
public class MainMenu extends
Form implements CommandListener{
...
}
```

Illustration 1: Class MainMenu

Since the menu extends form, it means that MainMenu is a form, and implementing CommandListener means that in MainMenu the CommandListener function will be used.

As shown in Illustration 2 there are five commands implemented on the main menu which are defined by cStart, cOptions, cUpdate, cAbout, cExit variables. Those commands are handled by commandAction() method.

```
public MainMenu(Displayable parent){
 ...
  addCommand(cStart = new
Command("Start", Command.SCREEN, 0));
  addCommand(cOptions = new
Command("Options", Command.SCREEN,
1));
  addCommand(cUpdate = new
Command("Update", Command.SCREEN,
1));
  addCommand(cAbout = new
Command("About", Command.SCREEN, 1));
  addCommand(cExit = new
Command("Exit", Command.EXIT, 0));
  setCommandListener(this);
 ...
}
```

Illustration 2: Command implementation

Bluetooth Discovery

Here, an external GPS device is required to gather the GPS data of the user's location. This external GPS device communicates with the mobile phone via Bluetooth. In order to communicate with a Bluetooth-supported device, the first step is to discover the devices around the mobile phone. Once the devices are discovered, users may select one and connect to it.

Thanks to a BLUElet API we can discover and connect to a Bluetooth device. Once the Midlet application is started, it searches for the devices around and allows the user to select one. When the user selects a device, a *url* of this device is built by this API again. Illustration 3 shows how to connect the selected device over this *url*.

```
StreamConnection con =
(StreamConnection) Connector.open( url );
```

Illustration 3: Connecting to a Bluetooth device

Receiving data over the established connection is shown in Illustration 4.

```
DataInputStream in =
con.openDataInputStream();
byte[] bytes = new byte[100];
int length=in.read(bytes);
String s = new String(bytes,0,length);
```

Illustration 4: Receiving data over Bluetooth

System Components

This section will introduce the core elements of the system that obtains the geographical information and decides whether we are approaching a fixed speed camera and alert us to it.

Figure 3: RoyalTek RLT 2100 GPS receiver

> Nmea is a combined electrical and data specification for communication between marine electronic devices such as echo sounder, sonar, anemometer and other types of instruments

Here, mobile phones are used in order to warn the drivers if they exceed the speed limit whilst driving. The mobile phone must be placed on the dashboard of the car with a phone holder. As mentioned earlier, mobile phones which have built-in GPS receivers may not give accurate data, due to their weak signal reception. Instead, in our system an external GPS receiver that can be placed on the dashboard is used.

For the purposes of this article, we are using a RoyalTek RBT 2100 GPS receiver (Figure 3) that has Bluetooth support to communicate with another device such as a mobile phone. Once the connection is established with a Bluetooth device, we can receive the GPS data from it to obtain the coordinates of our location.

GPS devices send data with the Nmea protocol. Nmea is a combined electrical and data specification for communication between marine electronic devices such as echo sounder, sonar, anemometer, gyrocompass, autopilot, GPS receivers and many other types of instruments.

Obtaining the Geographical Position

We consider the packages (strings) that start with $GPGGA to get the fix GPS data like longitude and latitude. The structure of a $GPGGA package is shown in Table 1.

Every sentence received from the GPS device contains the '*' character in the third last character, hence a check for this character is used to obtain the end of the package and checksum. All sentences received from the GPS device start with the '$' character and data fields are separated with the ',' character. The data structure of each GPGGA package is also shown in Table 1.

From this sentence we will need latitude and longitude values to evaluate the location. Those values can be obtained from the 2nd, 3rd, 4th and 5th data fields, where these values corresponds to latitude, direction of latitude (N or S), longitude and direction of longitude (W or E) respectively. The longitude and latitude values are in degree minute.minute format. Hence latitude is in ddmm.mmmm format and longitude is in dddmm.mmmm. However we will need the corresponding decimal values in order to evaluate the distance between the fixed speed-cameras. We

Field Name	Example Data	Description
Sentence Identifier	$GPGGA	Global Positioning System Fix Data
Time	070834	07:08:34 Z
Latitude	4064.8963, N	40d 64.8963' N
Longitude	08071.6838, W	80d 71.6838' W
Fix Quality: - 0 = Invalid - 1 = GPS fix - 2 = DGPS fix	1	Data is from a GPS fix
Number of Satellites	08	8 Satellites are in view
Horizontal Dilution of Precision (HDOP)	1.3	Relative accuracy of horizontal position
Altitude	140.2, M	140.2 meters above mean sea level
Height of geoid above WGS84 ellipsoid	-34.0, M	-34.0 meters
Time since last DGPS update	Blank	No last update
DGPS reference station id	Blank	No station id
Checksum	*75	Used by program to check for transmission errors

Table 1: Global Positioning System fix data

can make this conversion with the following equations:

lon = ddd + mm.mmmm/60

lat = dd + mm.mmmm/60

In Illustration 5 we demonstrate the required Java code to process the Nmea data and to obtain the latitude value in decimal format. This code checks the position of the '*' character and process the Nmea if its position is greater than 0. This control is followed by a set of string processing operations of the Nmea data if it starts with the "$GPGGA" token.

Obtaining the Speed

The speed of the car is as important as its location, and its speed may be evaluated in several ways. For example, two points (start and end) may be recorded within a time interval to calculate the speed. However, we have to consider that the distance on the globe cannot be calculated with a simple formula such as Euclidean Distance, because the path is not a straight line, instead it is a circular path.

Instead of those calculations, GPS receivers provide information about the speed of the vehicle. The speed can be obtained from the Nmea data with the $GPVTG packages. Information about the course over ground can be evaluated with $GPVTG data, where the details of this package are explained in Table 2. The speed of the car can easily be calculated from that.

Calculating the Distance

A distance measurement is required in order to specify whether the vehicle is approaching a speed camera or not. For this requirement, a regular measurement of the distance between the car and the closest fixed camera is done. Apart from this, the driver must be informed when the fixed speed cameras are, say 100 meters, ahead. So, we will need a way to calculate the distance in meters.

The distance between two points on a globe is not the same as the distance on the Earth's surface. The great-circle distance, or orthodormic distance, is the shortest distance between any two points on the surface of a sphere. There are many equations to calculate the great circle distance on the globe, but the most accurate one is the Vincenty Formula, shown in Equation 1. In this equation, $\varnothing_s,\Delta_s$; $\varnothing_f,\Delta_f$ are the geographical latitude and longitude of two points (base "start point" and the destination "fore point") respectively, and $\Delta\lambda$ indicates the difference of longitude values. The $\varnothing_s,\Delta_s$; $\varnothing_f,\Delta_f$ values must first be converted to radians and r should be the value of the Earth radius, which is 6371.01km.

Equation 1:

$$\Delta d = r.arctan\left(\frac{\sqrt{(cos\varnothing_f.sin\Delta\lambda)^2+(cos\varnothing_s sin\varnothing_f - sin\varnothing_s.cos\varnothing_f.cos\Delta\lambda)^2}}{sin\varnothing_s.sin\varnothing_f+cos\varnothing_s.cos\varnothing_f.cos\Delta\lambda}\right)$$

Field Name	Example Data	Description
Sentence ID	$GPVTG	
Course	89.68	Course in degrees
Reference	89.68	T = True heading
Course		Course in degrees
Reference	89.68	M = Magnetic heading
Speed	0.00	Horizontal speed
Units	N	N = Knots
Speed	0.00	Horizontal speed
Units	K	K = KM/h
Checksum	*5F	

Table 2: $GPVTG package: Course over ground

Illustration 5: Evaluating the latitude from Nmea data

```
int checkSum = nmea.indexOf('*');
if (checkSum == -1)
 return;
String [] nmeaStr = StringProcess.getArray(
 nmearaw.substring(0, checkSum), ",");

if (nmeaString[0].equals("$GPGGA")) {
 if (nmeaStr[2].length() > 8 && nmeaStr[3].length() == 1) {
 degree = Integer.parseInt(nmeaStr[2].substring(0, 2));
 minute = Integer.parseInt(nmeaStr[2].substring(2, 4));
 fraction = Integer.parseInt(
             nmeaStr[2].substring(5, 9).concat("0000").substring(0, 4));
 latitude = new Float(degree).Add( new Float(minute).Div(60)).Add(
 new Float(fraction).Div(600000));
 if (nmeaStr[3].charAt(0) == 'S')
 latitude = latitude.Neg();
 }
}
```

Illustration 6: Execution of SoundPlayer

```
public void run() {

 if ( busy )
 return;

 try
 {
 busy = true;
 InputStream is = getClass().getResourceAsStream("AlarmSounds/Polaris.mp3");
 Player player = Manager.createPlayer(is, "audio/mpeg");

 player.realize();
 player.prefetch();
 player.start();

 Thread.sleep(3000);

 player.stop();
 player.close();

 busy = false;
 close();
 }
 catch(Exception e)
 {}
 }
```

The Alert Tones

Once the system decides that the driver is close to the camera, it gives an alert. In this study we prefer to play an alert tone to warn the driver, and we play the alert in a thread as, while the alert is sounding, the system continues to take the GPS data and measures the distances. Illustration 6 is showing how to execute the SoundPlayer in a thread, where Polaris.mp3 is chosen as the alert.

The Entire System

In this article we have presented a system that can be used by drivers in their cars to be informed when they are approaching fixed speed-cameras. The locations of fixed cameras are important, as they are installed on the most dangerous roads. Although there are traffic signs to inform the drivers of speed cameras, some can be missed for variety of reasons.

In this article, a low cost and fast system is introduced that can be used for better navigation whilst driving. GPS was used to obtain the required coordinates and the speed of the vehicle. The proposed software is installed in a mobile phone so that it will give the required warnings to reduce the speed once the vehicle enters a speed-limited zone.

As a future work, an online map could be integrated into the system to display the location of the car and where the speed cameras are on a map. Another improvement could be to add all of the speed limits on the roads instead of just showing where the cameras are; a warning system for humps that are used as the speed traps could also be added. ●

WHAT THE READERS SAY

HAVING THE COURAGE TO PRINT

Referring to the articles '*The end of electric charge and electric current as we know them*' part one and two in your January and February 2011 editions, I would like to congratulate you on your courage in publishing this un-refereed article, the conclusions of which would demand a major paradigm change.

I find it most refreshing to be able to read about a theory in an authoritative journal such as '*Electronics World*' that has the courage to present, and call 'explosive', such a model that clearly highlights major issues in the current mainstream established view.

Now that this has been presented, one wonders if there are any 'established experts' who would also have the courage to comment on this article.

Ian Montgomery

EW Does a Great Service to Engineering

I write to urge you to continue publishing articles by Ivor Catt as his insights are of great importance to science and engineering. Many journals will not publish his papers for the simple reason that they upset too many apple carts. The issues raised by Catt need to be more widely known and discussed objectively and this can only be achieved by publication. *Electronics World* does a great service to science and engineering by publishing Catt's articles.

Steve Crothers

Providing Insight

I met Ivor in the 1970s and he and our colleague David Walton published many articles in Wireless World, various other magazines and we also published a book Digital Electronic Design. During those years we recognized the work of JA Fleming, Oliver Heaviside and others who had recognized that electric current as we know it is a secondary effect and the more fundamental energy current was the basic primitive.

Ivor has continued with this work for the past 30 years and, as we gain more insight into the theoretical flaws of basic EM theory, we can construct more accurate models which will be better suited to the challenges of electronics in the foreseeable future.

More effective modelling and understanding of electrical energy can provide insight into creating improved designs for transmitting, consuming and processing this valuable energy source.

Malcolm F. Davidson
USA

RATIONAL DEBATE

It is alarming that Ivor's clearly argued assertions on Electric Charge and Current are not debated rationally but this can be understood in that many 'academics' would be embarrassed by a change to their established understanding.

In July 1961 I joined Ferranti computer department as a new graduate development engineer. I was assigned to Ivor for the first six months of my industrial baptism.

This was extremely valuable as Ivor would always examine concepts from their basic elements.

It was a conversation at that time which lead Ivor to develop his seminal paper on metastability, published some years later.

John Raymond Dore

Failing to Notice

I first met Ivor in the early 1970s when he was working at Marconi_Elloitt avionics in Borehamwood. It was very shortly after this that I read Ivor's crosstalk paper, which became, together with some of his other writings, the starting point for our collaboration in fast digital systems.

It is truly fascinating that we both failed to notice that the superposition of modes of transmission in a 3-wire system, as experimentally demonstrated and documented in that paper, is in itself a demonstration that charge and current are the result of the energy current flow and not its cause. As Ivor summarizes in his articles: "One reason why it is illegal is that the electric currents in the right hand conductor are in opposite directions for the two modes, and classical theory says there cannot be two electric currents in opposite directions along a single conductor. However, two electromagnetic waves (or light rays) can be in the same point in space, for instance when we shine a torch at another lighted torch pointing in the opposite direction, or when we send two pulses from left and right through each other down a coaxial cable. Similarly, the Even and Odd Mode TEM Waves in our photos can coexist, but not their associated electric charges and currents."

I trust that the publication of these articles will stimulate a fruitful dialogue which will lead others to identify the discrepancies between the Maxwellian electromagnetic model and observed reality.

David Walton

IF YOU WOULD LIKE TO COMMENT
on this subject on any other that you have read on in Electronics World magazine, please write to the Editor at Svetlana.josifovska@stjohnpatrick.com

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

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

BATTERY CONDITIONER EXTENDS THE LIFE OF LI-ION BATTERIES

BY **GEORGE H. BARBEHENN**, SENIOR APPLICATIONS ENGINEER, LINEAR TECHNOLOGY

Li-Ion batteries naturally age, with an expected lifetime of about three years. But, that life can be cut very short – to under a year – if the batteries are mishandled.

It turns out that the batteries are typically abused in applications where intelligent conditioning would otherwise significantly extend the battery lifetime.

Modern Li-Ion batteries are constructed of a graphite negative terminal, cobalt, manganese or iron phosphate positive terminal and an electrolyte that transports the lithium ions.

The electrolyte may be a gel, a polymer (Li-Ion/Polymer batteries) or a hybrid of a gel and a polymer. In practice, no suitable polymer has been found that transports lithium ions effectively at room temperature. Most 'pouch' Li-Ion/Polymer batteries are, in fact, hybrid batteries containing a combination of polymer and gel electrolytes.

The charge process involves lithium ions moving out of the negative terminal material, through the electrolyte and into the positive terminal material. Discharging is the reverse process. Both terminals either release or absorb lithium ions, depending on whether the battery is being charged or discharged.

Figure 1: **Yearly capacity loss vs temperature and state of charge for Li-Ion batteries**

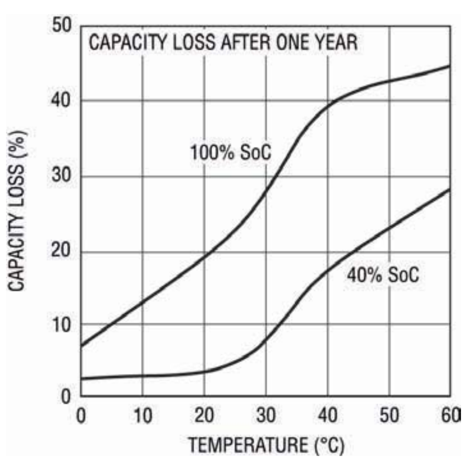

The lithium ions do not bond with the terminals, but rather enter the terminals much like water enters a sponge; this process is called 'intercalation'. So, as is often the case with charge-based devices such as electrolytic capacitors, the resulting charge storage is a function of both the materials used and the physical structure of the material. In the case of the electrolytic capacitor, the foil is etched to increase its surface area. In the case of the Li-Ion battery the terminals must have a sponge-like physical make-up to accept the lithium ions.

The choice of positive terminal material (cobalt, manganese or iron phosphate) determines the capacity, safety and aging properties of the battery. In particular, cobalt provides superior capacity and aging characteristics, but it is relatively unsafe compared to the other materials. Metallic lithium is flammable and the cobalt positive terminal tends to form metallic lithium during the discharge process. If several safety measures fail or are defeated, the resulting metallic lithium can fuel a "vent with flame" event.

Consequently, most modern Li-Ion batteries use a manganese or iron phosphate-based positive terminal. The price for increased safety is slightly reduced capacity and increased aging.

> *Battery aging results in an increase of the battery series resistance (BSR) and reduced capacity*

Aging is caused by corrosion, usually oxidation, of the positive terminal by the electrolyte. This reduces both the effectiveness of the electrolyte in lithium-ion transport and the sponge-like lithium-ion absorption capability of the positive terminal. Battery aging results in an increase of the battery series resistance (BSR) and reduced capacity, as the positive terminal is progressively less able to absorb lithium ions.

The aging process begins from the moment the battery is manufactured and cannot be stopped. However, battery handling plays an important role in how quickly aging progresses.

Conditions That Affect the Aging Process

The corrosion of the positive terminal is a chemical process and this chemical process has an activation energy probability distribution function (PDF). The activation energy can come from heat or the terminal voltage. The more activation energy available from these two sources the greater the chemical reaction rate and the faster the aging.

Li-Ion batteries that are used in the automotive environment must last 10 to 15 years. So, suppliers of automotive Li-Ion batteries do not recommend charging the batteries above 3.8V. This does not allow the use of the full capacity of the battery, but is low enough on the activation energy PDF to keep corrosion to a minimum. The iron phosphate positive terminal has a shallower discharge curve, thus retaining more capacity at 3.8V.

Battery manufacturers typically store batteries at 15°C (59°F) and a 40% State of Charge (SoC) to minimize aging. Ideally, storage would take place at 4% or 5% SoC, but it must never reach 0%, or the battery may be damaged. Typically, a battery pack protection IC prevents a battery from reaching 0% SoC. But pack protection cannot prevent self-discharge and the pack protection IC itself consumes some current.

Although Li-Ion batteries have less self-discharge than most other secondary batteries, the storage time is somewhat open-ended. So, 40% SoC represents a compromise between minimizing aging and preventing damage while in storage (see Figure 1).

In portable applications, the reduction in capacity from such a reduced SoC strategy is viewed negatively in marketing specifications. But it is sufficient to detect the combination of high ambient heat and high battery SoC to implement an algorithm that minimizes aging while ensuring maximum capacity availability to the user.

Battery Conditioner Avoids Conditions that Accelerate Aging

Linear Technology has the LTC4099 battery charger and power manager contains an I2C controlled battery conditioner that maximizes battery

Figure 2: **The LTC4099 with I2C controlled battery conditioner**

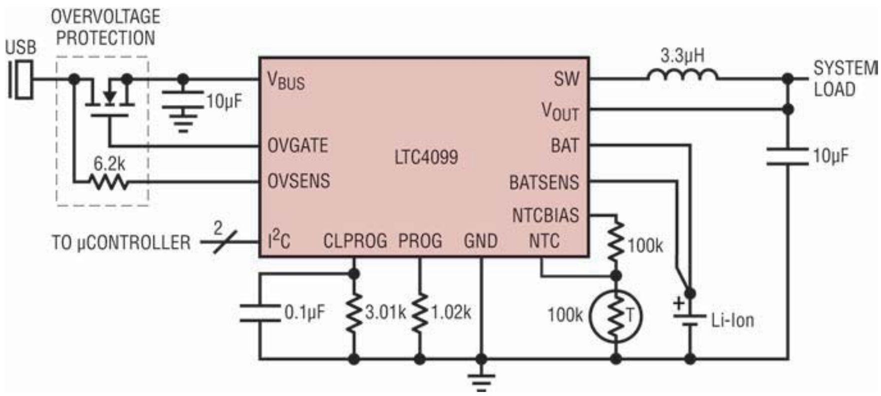

Figure 3: **Battery discharge current vs voltage for the LTC4099 battery conditioning function**

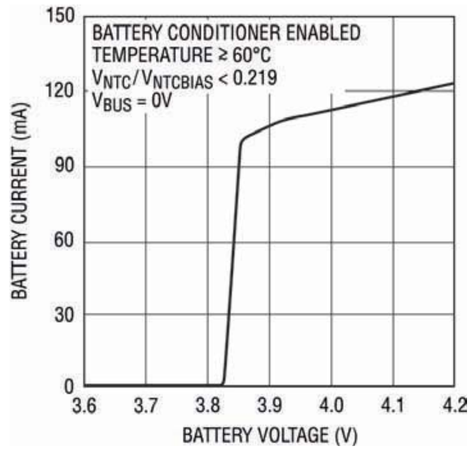

operating life, while also optimizing battery run time and charging speed (see Figure 2).

It has a built-in battery conditioner that can be enabled or disabled (default) via the I2C interface. If the battery conditioner is enabled and the LTC4099 detects that the battery temperature is higher than ~60°C, it gently discharges the battery to minimize the effects of aging.

The LTC4099 NTC temperature measurement is always on and available to monitor the battery temperature. This circuit is a micropower circuit, drawing only 50nA while still providing full functionality.

The amount of current used to discharge the battery follows the curve shown in Figure 3, reaching zero when the battery terminal voltage is ~3.85V. If the temperature of the battery pack drops below ~40°C and a source of energy is available, the LTC4099 once again charges the battery. Thus, the battery is protected from the worst case battery aging conditions.

Although the aging of Li-Ion batteries cannot be stopped, the LTC4099's battery conditioner ensures maximum battery life by preventing the battery-killing conditions of simultaneous high voltage and high temperature. Further, the micropower, always-on NTC monitoring circuit ensures that the battery is protected from life-threatening conditions at all times. ●

PRODUCT OF THE MONTH

APACER
CONSOLIDATES ITS INDUSTRIAL PATA SOLID STATE DRIVES (SSD) PORTFOLIO

Apacer has consolidated its offerings of PATA Solutions by launching its new 2.5/1.8-inch high-speed and capacity AFD (ATA-Flash-Drives) SSDs.

The SLC (Single-Layer-Cell) Flash AFD 255 comes in 2.5-inch 44pin connectors at 32 and 64GB, while the SLC 1.8-inch AFD 185 supports ZIF connectors and offers a maximum capacity of 32GB.

With sequential read/write speed of up to 125/110MB/sec respectively, both can enhance computer performance and meet user's needs for highly stable and reliable storage solutions. Furthermore, the SLC PATA SSD series support wide temperature range -40/+85degC. Also introduced with MLC (Multi-Level-Cell) flash, the AFD 255-M and AFD 185-M have maximum capacities of 256GB and 128GB, and sequential read/write speed up to 120/90MB/s, respectively.

Both new series feature built-in S.M.A.R.T. technology and intelligent power failure recovery function to safeguard data and enhance product stability.

http://eu.apacer.com/business/industrial-ssd/pata/

IDT Expands Wireless Infrastructure Portfolio

Integrated Device Technology (IDT) announced a digitally-controlled IF VGA with very high SNR for improved Quality of Service (QoS) for receiver systems located in cellular base stations and other wireless infrastructure equipment.

The new device expands IDT's product offerings for base station equipment in the growing 4G wireless infrastructure market, which today includes high performance timing, RapidIO solutions, and other devices from IDT's rich communications portfolio.

The IDT F1200 digitally-controlled IF VGA's noise figure is less than 3.0dB and a broad 23dB control range with 7-bit resolution, resulting in better QoS for cellular base stations. The extremely low distortion provides increased spurious-free dynamic range (SFDR), making the device flexible for use in myriad applications. In addition, the excellent gain accuracy makes it easy to digitally compensate, while the 200 ohm differential input and output impedances allow it to integrate seamlessly into the RF signal path, minimizing design effort and risk.

www.idt.com

COMPLETE ENERGY MEASUREMENT SOC FOR AC/DC POWER SUPPLIES

Maxim Integrated Products has announced the launch of a new Teridian/Maxim energy-measurement system on chip (SoC) – the 78M6613. This is the industry's first SoC energy-measurement solution for AC/DC power supplies that brings a higher level of management and control to servers and other equipment in data centers.

The new Teridian/Maxim chip comes as Environmental Protection Agency (EPA)/ENERGY STAR mandates toward greater efficiency and accurate energy measurement in data centers begin to bear down on enterprises around the globe. The 78M6613 enables the capture and reporting of real-time energy data, which provides data center managers with the ability to quantify where energy is needed, being used and, more importantly, being stranded. Visibility into this critical data is an absolute requirement in order to manage and control energy usage in power-starved data centers that cannot keep up with growing demands placed upon them.

The SoC offers an embedded analog front-end and compute engine, small footprint and embedded firmware.

www.maxim-ic.com

Automotive-Grade Programmable Stepper-Motor Driver ICs

The A3981 and A4980 from Allegro MicroSystems Europe are two new automotive-grade programmable stepper-motor driver ICs that are designed to operate bipolar stepper motors in full, half, quarter and sixteenth-step modes at up to 28V and ±750mA in automotive, consumer and industrial applications.

Both devices offer a single-chip solution with internal digital timing circuitry and a built-in translator for ease of operation. The A3981 can be controlled by simple step and direction inputs, or through an SPI-compatible serial port that also can be used to program many of the integrated features and to read diagnostic information.

The A4980 is identical to the A3981 except for a low-voltage operation feature which will allow the IC to run down to a load supply voltage of 3.3V (compared with 7V on the A3981). Such a low voltage operation may be necessary in some automotive applications where the battery voltage during cold cranking can drop to a critically low level.

www.allegromicro.com

RUGGED PUSH BUTTON SWITCH WITH DUST AND WATER PROTECTION

The latest high quality customisable switch to be introduced by Foremost Electronics, the Essex based specialist distributor of electromechanical components, is the Marquardt Mechatronik GmbH model 1683 pushbutton switch suitable for use in appliances requiring protection class II.

Suitable for use in harsh industrial environments, including high dust and water areas, the 1683 pushbutton switch is protected to IP65 and IP67 levels and can control devices of up to 4,000W with its powerful switching system for currents up to 16A. The highly reliable double lifting contact system has a proven high life endurance of over 1,000,000 operations and has a number of different connection options including quick connect, solder and PCB terminals.

This attractive and ergonomically designed switch can accommodate single and double pole contact options in the same compact package, with on/off and momentary actions. The switch cap is available in a range of colours, may be printed with customer specific legends.

www.4most.co.uk

MKJ TRINITY MINI CIRCULAR CONNECTORS ARE AVAILABLE ON 3D WEB MODELLING SITE

Connector manufacturer and supplier, ITT Interconnect Solutions, has announced the immediate availability of its Trinity MKJ series mini-circular connectors as downloadable 3D models on its interactive 3D web modelling site. The 3D modelling function was launched in the second half of last year and provides an invaluable resource for design engineers to try out different connector solutions in prototype designs early in the design process. Engineers request a downloadable 3D CAD drawing which they can then import into their own CAD systems, saving time and effort.

Now with Trinity mini-circular interconnects available as CAD models, engineers have even greater design flexibility. This broad product range offers a unique combination of design, functionality and flexibility and suits numerous applications in multiple markets, including industrial, medical, aerospace and military. MKJ0, MKJ1 and MKJ4 series connectors are up to 71% lighter and up to 52% smaller than equivalent military standard environmental interconnects, yet provide similar mechanical and electrical performance.

www.ittcannon.com

Industrial C90 Circular Connector Family with 360º EMC braid protection

Harwin has announced that users of its C90 industrial range of circular connectors can now benefit from 360º EMC screening via a new series of screening braid adapters.

Fully RoHS-compliant and available from stock in many standard sizes, the new adapters are manufactured in black zinc nickel plated aluminium alloy and enable easy connection of RFI and EMI screening braid to Harwin's popular C90 connector series.

A choice of termination is available. The first option is a constant force spring system, allowing termination without special tooling. Alternatively, a Tie-Dex Micro Band system may be used which is easily terminated using industry-standard application tooling. Both termination styles provide full 360 degree attachment of the braid ensuring optimum screening performance.

Designed for use in harsh industrial environments, agricultural and construction equipment, these adapters, together with the appropriate braids, provide proven EMC performance. Devices are in stock now for immediate delivery.

www.harwin.com

COMPACT 120W MIL-SPEC DC-DC CONVERTER IN 12V AND 15V OUTPUT VERSIONS

Roband Electronics has launched a compact military specification DC-DC converter in 12V and 15V output versions, providing a power output of 120W from a package measuring just 76.3 x 38.2 x 10.2mm.

The new RO-MIL-2212 is a high-efficiency, high power-density, single-rail unit comprising an isolated, high-frequency, pulse width modulated push-pull converter. While it is specified to operate from a nominal 28V supply, this versatile power supply will accept a wide input voltage range – from 16-40V (12V output version) and 14.5-40V (15V output version). Options include an 80V input surge capability and a variant that will operate from an input as low as 9V. External synchronisation is also available as an option.

The RO-MIL-2212 operates over the military temperature range without de-rating. It is unconditionally stable and does not require additional external components for correct operation. A remote sense facility is provided to compensate for any voltage drop at the load due to the length of the power supply leads.

www.roband.co.uk

New Battery Disconnection Unit Makes Design-In Easier and Faster

The Battery Disconnection Unit (BDU) is of compact size and serves as an interface between the high-voltage battery and the electric drive. The BDU contains two main relays and a current sensor as well as precharge and diagnostic electronics. It is designed for a continuous current of 300A. Over one minute, the BDU can even master currents of up to 600A, thus allowing it to be used in very powerful drives. In case of failure the EV relays disconnect currents of up to 2500A without problem, thus enhancing the safety of the new car generation.

With the relays from the EV relay family, Panasonic offers a solution in a wide range of sizes and switching capacities. Millions of EV relays are already designed into all kinds of vehicles and make the EV relay the leader in the automobile market segment.

www.panasonic-electric-works.co.uk

NEW ISSUE OF TECH AUDIO PUBLICATION LINEAR AUDIO: VOLUME 1

Linear Audio Publishers of Hoensbroek, The Netherlands, have announced the publication of the second issue of their semi-annual tech audio 'bookzine' Linear Audio. The printed publication's goal is to provide high quality, in-depth technical audio design and construction projects to professionals and serious hobbyists.

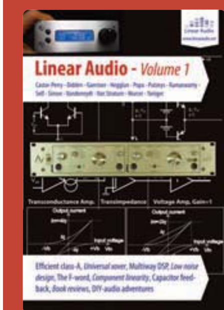

Linear Audio Volume 1 packs 180 pages of technical audio articles, and includes a high-efficiency class A power amp, low noise phono and mic preamp design, DSP-based and analog cross-overs, a review of feedback, passive component linearity measurements, book reviews and more.

Volume 1 authors include Ian Hegglun, Douglas Self, Kendall Castor-Perry, Bruno Putzeys, Scott Wurcer and others. Linear Audio Volume 1 can be ordered online.All article abstracts and author bio's can be read at

www.linearaudio.net

Optimized Embedded Single-Board Computer with AMD Geode LX800

Kontron announced the new 2.5-inch Pico-ITX SBC KTLX800/pITX, featuring the AMD Geode LX800 processor that combines maximum cost and power efficiency with interfaces tailored to the needs of industrial automation and test and measurement. With long-term availability and an average power consumption of only 5-7W the Kontron KTLX800/pITX is an ideal foundation for extra-small and cost-sensitive industrial applications.

The Kontron 2.5-inch Pico-ITX SBC KTLX800/pITX offers a wide choice of interfaces, including VGA, plus LVDS with JILI30, two serial interfaces, Fast Ethernet, SATA, four USB 2.0 and 16 GPIO.

This extremely compact SBC (100 x 72mm) enables efficient development of cost-optimized small form-factor embedded designs and ultra-mobile devices such as those used in industrial automation and test and measurement.

Ensuring long-term availability, the Kontron 2.5-inch Pico-ITX SBC KTLX800/pITX comes with up to 1GB of DDR RAM system memory as supported by the AMD Geode LX800 500MHz processor which uses an average of just 0.9W.

www.kontron.com/sbc/

HAMEG INTRODUCES NEW ENTRY LEVEL DSO/MSO

Hameg Instruments has complemented its HMO Series of DSO/MSO Oscilloscopes with entry-level bandwidths of between 70 and 200MHz.

The HMO3524/3522/2524, introduced some two years ago, excel due to the Mixed-Signal Oscilloscope upgrade capability (even post-sale) and the hardware-accelerated trigger and decode functionality, supporting I2C, SPI, UART/RS-232 protocols. With the new HMO derivatives, this enhanced feature-set will be now available to customers with budgets starting at €1,150.

The new instruments will be available in 2- and 4 channel versions with 70, 100, 150 and 200MHz bandwidths. All models will provide 1GSa/s sampling rate and 1Mpt of memory per channel, which can be combined to 2GSa/s and 2Mpts respectively.

All new HMOs are MSO-ready when used with the logic probe (option HO3508) in order to analyze 8 digital channels. With the existing option HOO10, two independent serial protocols I2C, SPI and/or UART/RS-232 can be analyzed, hardware accelerated and time-synchronous. Both analogue and digital channels may be used with this option.

www.hameg.com/embedded

STMICROELECTRONICS UNVEILS NEXT-GENERATION INEMO ENGINE

STMicroelectronics (ST) unveiled the first details of its new iNEMO Engine.

An industry first, the iNEMO's new advanced filtering and predictive software engine integrates the outputs from a 3-axis accelerometer, 3-axis gyroscope and 3-axis magnetometer; fusing these sensors' data through sophisticated algorithms, the iNEMO delivers dramatically more accurate and reliable sensor performance, suitable for next-generation smart consumer devices that use a host of new motion-based applications.

For many current applications, such as freefall detection, screen rotation or pedometers, a single MEMS accelerometer meets their system requirements. However, a new class of advanced applications is emerging, for example: location-based services, enhanced motion-based gaming, pedestrian dead-reckoning for indoor and multi-floor navigation, robotics or human-body tracking. These applications require multiple MEMS sensors, together with advanced software, to achieve better overall system performance in terms of accuracy, resolution, stability and response time.

The iNEMO Engine is available now in a two-chip sensor solution (LSM303DLHx e-compass and the L3G4200D gyroscope).

www.st.com/mems

Rapid Electronics Helps Guitar Heroes with New Stomp Box Range

Rapid Electronics has enlarged its instrument enclosure range with the 1590 'stomp' series of die-cast aluminium boxes, as well as introducing two other new ranges of high quality enclosures from Hammond Manufacturing.

The Hammond 1590 die-cast aluminium enclosures are designed to be the housings of choice for stomp box or guitar effects pedal manufacturers. The 1590 is a rugged, easy to machine die-cast aluminium enclosure able to cope with the demanding on-stage environment in which they will be used.

Available in two sizes, 112 x 60 x 27mm and 119 x 94 x 34mm, they are finished in a smooth gloss polyester powder paint, which does not chip after machining and provides a good surface for labels and silk screening. A lap joint seals the units to IP54, and the painted finish is only applied to the external surfaces, maintaining RFI integrity.

www.rapidonline.com

WORLD'S SMALLEST WAVEFORM MONITOR FOR BROADCASTING APPLICATIONS

Tektronix unveiled the compact WFM5200 waveform monitor and WVR5200 rasterizer to meet the needs of video content production, post production and broadcast for affordable 3G-SDI monitoring capability that can be tailored to specific applications.

The world's smallest full-featured waveform rasterizer and waveform monitor can monitor up to four 3G-SDI signals simultaneously. The WFM/WVR5200 products are well suited for use in OB vans, studio control rooms, and post production houses doing editing, special effects and colour correction. With software that can be tailored to specific applications, the new instruments meet the growing trend toward "fit for purpose" products that can easily be integrated into an existing system architecture.

While the new WFM/WVR5200 products come in a compact package, they deliver the advanced capabilities required in modern production, post production and broadcast applications.

www.tektronix.com

LIGHT-SPEED DEVELOPMENTS IN PHOTONIC CRYSTAL TECHNOLOGY

European scientists and engineers are working together on the COPERNICUS project, to develop cutting-edge photonic crystal technology that has the potential to make electronic devices much faster, smaller and more efficient.

Photonic technology uses light – instead of electric current – to send and receive signals at extremely high speeds. The technology has a huge number of applications, from telecommunications, medicine and manufacturing to aviation, construction, consumer equipment and many other areas.

Photonic devices will address the pressing need for low-power, ultra-high bandwidth data links in server farms, optical storage networks and on-board Internet/entertainment systems, where demand is driving the data bandwidth and technology integration level rapidly upwards. Next generation telecom systems will also benefit from these devices.

The COPERNICUS project brings together eight European academic and industrial partners with high profiles in photonics, nanotechnology, modelling and developing new technologies for telecommunications and aerospace. It has received funding of nearly €3m from the European Commission's Information Society Technologies Programme.

PROFESSOR DR DOGAN IBRAHIM, Near East University in Nicosia, Cyprus:

We are coming to the end of conventional electric current based signal processing as far as the speed and low power operation are concerned. The photonic crystals and nanoscale materials seem to be the way forward for ultra-fast and low-power signal processing applications. With the development of all-optical gates, filters, photodetectors and other nanophotonic devices, it is clear that the size of the future electronic devices will be reduced considerably with an increase in the speed, perhaps well beyond 100GHz, using all-optical processing. We should then be able to develop signal processing elements, data links and optical storage devices with extremely low-power and ultra-high bandwidth. I am very happy about the research and development in this field.

BURKHARD VOGEL, Managing Director, Germany:

In the case of a financial support of the research and development of the usage of light in electronic components and as the main transfer medium of information, the EU moves towards the right direction. The worldwide competition is very strong and, compared with the sales amount in recent years, the funding of only EUR 3m seems to be rather small.

However, we should be lucky with an EU administration that got enlightened by light as one of the major future business drivers in the electronics field.

Independent of the EU decision, I guess, the management of light (sunlight in the energy creation field too) will be one of the main tasks that could solve most of the current problems of energy consumption and processing of information.

BARRY MCKEOWN, RF and Microwave Engineer in the Defence Industry, and Director of Datod Ltd, UK:

The physical development of 3D Photonic Crystals is currently at the forefront of the advance in physical chemistry for the control of the flow of light for the realisation of optical computers. The initially highly focused and time-limited project key aim of developing compact demultiplexing receivers requires unspecified breakthrough technologies to emerge: a big ask. Opened end research programs are out of the question here. The well balanced composition of the collaboration should enable dissemination of the emerging technologies through the on-line tutorials and UK involvement should bode well for the future of an advanced UK material science industry.

MAURIZIO DI PAOLO EMILIO, Telecommunications Engineer, INFN – Laboratori Nazionali del Gran Sasso, Italy:

Just as the transistor and microelectronics transformed communications and human society in the 20th century, "light" and microphotonics will revolutionise the way we communicate in the 21st century. Photonics technology can already be found in nearly every home, office and manufacturing plant; technological innovations in areas as varied as entertainment systems, displays, data storage, telecommunications systems and Internet connectivity among others.

Computers have enhanced human life to a great extent. The speed of conventional computers is achieved by miniaturizing electronic components to a very small micron-size scale so that those electrons need to travel only very short distances within a very short time. Our daily lives demand solutions to increasingly sophisticated and complex problems, which requires more speed and better performance of computers. Optical interconnections and optical integrated circuits will provide a way out of these limitations to computational speed and complexity inherent in conventional electronics. Optical interconnections and optical integrated circuits have several advantageous over their electronic counterparts. They are immune to electromagnetic interference and free from electrical short circuits.

HAFIDH MECHERGUI, Associate Professor in Electrical Engineering at the University of Tunisia:

The European collaborative research projects allow several scientific teams and organizations to work together in promising fields. Indeed, the co-operation on the COPERNICUS project contributes enormously to make European space of research attractive for the best crop of scientists and to become a reference post for them.

I think that the COPERNICUS project must be encouraged as it will make possible the research teams of the European Union to show creation and innovation by undertaking work on common challenges. This will allow the development of advanced technologies for our century.

Integrate Touch Sensing Quickly and Easily

With Microchip's Range of Low Power, Low Cost Solutions

Microchip's mTouch™ Sensing Solutions allow designers to integrate touch sensing with application code in a single microcontroller, reducing total system cost.

Microchip offers a broad portfolio of low power, low cost & flexible solutions for keys/sliders and touch screen controllers. Get to market faster using our easy GUI-based tools, free source code and low-cost development tools.

Capacitive Touch Keys and Sliders

- Extend battery life with eXtreme Low Power MCUs
  - Proximity sensing in less than 1 µA
- High noise immunity and low emissions
- Broad portfolio of MCUs lowers system cost
  - 8, 16 & 32-bit PIC® MCUs for Capacitive Touch
  - Integrated USB, Graphics, LCD, IrDA, CAN
  - No external components
- With **Metal Over Cap technology** you can:
  - Use polished or brushed metal surfaces including stainless steel and aluminium
  - Sense through gloves
  - Create waterproof designs
  - Deploy Braille-friendly interfaces

Touch Screen Controllers

- Fully processed touch coordinates
- **Projected Capacitive technology**
  - Multi-touch enabling gestures
  - Low cost MCU implementation
  - Wide operating voltage: 1.8-5.5V
  - Low operating current 1.5 mA at 5V typical
- **Analog Resistive technology**
  - Lowest system cost, easy integration
  - Universal 4, 5 & 8-wire solution with on-chip calibration
  - I^2C™, SPI, UART or USB interfaces
  - Low power "touch to wake-up" feature

GET STARTED IN 3 EASY STEPS

- Learn more at www.microchip.com/mtouch
- Download App Notes & royalty-free source code
- Order a development tool

Enhanced mTouch Capacitive Evaluation Kit - DM183026-2 (For keys & sliders)

Projected Capacitive Development Kit - DM160211

Analog Resistive Touch Screen Development Kit - DV102011

Intelligent Electronics start with Microchip