ELECTRONICS WORLD

THE ESSENTIAL ELECTRONIC ENGINEERING MAGAZINE

LCDs
THE WORKINGS
MCUs
MEDICAL DESIGN

MILITARY SKILLS

SPILL INTO THE CIVILIAN MARKET

TECHNOLOGY
DUAL
NANOIMPRINT
LITHOGRAPHY
SYSTEM

MIPI
HIGH-SPEED
SERIAL VIDEO
INTERFACE

PRODUCTS
STANDARD
AND PRESS-FIT
CONNECTORS

ALSO IN THIS ISSUE: TIPS 'N' TRICKS NEW COMPETITION @ LETTERS

FEATURES

REGULARS

FUNDING INEQUALITY

As editors we are regularly inundated with information: which company brought out which new product, who cut a deal with whom, about the latest technological advances and so on. Among this information are announcements made by universities, government bodies, small and large firms receiving grants for certain projects and R&D programmes.

There's no shortage of such announcements and they come from all over the world, but one thing is clear, in the UK this funding seems to be a lot smaller – regardless of the size of the project or future importance of the technology – than anywhere else.

Case in point: Only recently I was writing stories involving two different projects, with similar size consortia but different technologies; one was in Germany, the other in the UK.

In Germany, it was announced, there's a project focusing on the development and manufacture of photovoltaic cells modules which involves the government and four industrial partners. The programme's goal: to manufacture solar cells in a cost-effective manner, increase the number of areas of application and make Germany a key player in this field.

Size of funding: £250m.

Here in the UK, on the other hand, we have a consortium formed also of four industrial partners but a university too. The government is sponsoring the programme. The programme's goal: to develop a route to low-cost LEDs for solid state lighting and make the UK a key player in this field.

Size of funding: £3m.

That's nearly hundred times less than what the German government is giving away. Who do you think will make a better go of it?!

THIS STORY OF FUNDING INEQUALITY HAS BEEN REPEATED MANY TIMES OVER THE PAST FEW YEARS

This story of funding inequality has been repeated many times over the past few years. There's no shortage of interesting technological fields, and our academics and firms like to get involved. After all, there are a lot of expertise and capabilities in the UK. So yes, the UK wants to be seen as doing something in new technologies, and the government wants to make this country competitive in emerging markets. However, is its commitment a serious one or just a token gesture? How much research is likely to be done on a say, £1m to £3m grant? I'd like to hear from those involved in such projects – maybe it's me who's missing a trick.

Svetlana Josifovska
Editor

IMEC develops improved FinFETs

Stage delay of NAND and NOR gates in dependence of fan-in, slowest input and SEM picture of NAND3 gate

IMEC has fine-tuned a manufacturing process which will enable reproducible FinFETs with fin widths of down to 5nm. The academics used 193nm immersion lithography and dry etching.

"Although the performance benefits of FinFETs have been recognised for many years, several bottlenecks had to be overcome to bring FinFET technology to manufacturing. These advances [at IMEC] have reduced the gap for FinFETs to become a manufacturing technology," said Luc van den Hove from IMEC.

When scaling CMOS circuits towards the 32nm node and beyond, the use of FinFETs may solve problems associated with short channel effects and leakage. The use of these ultra-thin bodied devices eliminates the need for channel doping and hence the problems of doping fluctuations – due to only a few doping atoms – in the nanoscale planar devices. This results in closer tolerance parameters together with reduced junction leakage. IMEC showed experimentally that the performance of FinFET circuits is better than that of bulk CMOS and meets the needs of future digital libraries.

FinFETs' parameters have been considerably improved by a dielectric passivation process based on the use of fluorine in the metal-hafnium based gate-stack during gate etching. This reduces both NBTI and PBTI (negative and positive bias temperature instability).

Apart from the extremely small inverter delay, FinFETs provide an excellent stacked device performance too, permitting higher stack heights and thus enabling the same number of functions to be implemented with fewer logic gates and smaller silicon area.

GM system alerts drivers

General Motors (GM) has launched a new in-vehicle system which alerts drivers to an impending accident.

The 'vehicle to vehicle' (V2V) technology uses GPS data to plot the precise location of a car and relays that information over a wireless network to other vehicles within its range, triggering a warning if they get too close to each other. If on a collision course, the system sounds an alarm, alerting the driver to do a corrective action.

Car communication can take place at distances of up to 250m,

including around corners, which will improve visibility and help eliminate 'blind spots'.

At present the V2V device can be retrofitted to vehicles at a cost of £125. Andrew Marshall, of Vauxhall, a brand owned by GM, said: *"We have deliberately based this technology on inexpensive, proven components, giving it the potential to become standard equipment in many vehicles."*

GM plans to install it in new cars as of 2012.

Scavenging device good for powering pacemakers

Scientists at the University of Southampton have created a tiny generator powered by vibrations, which will also help power medical implants including pacemakers.

The generator has been developed to power devices where replacing batteries is very difficult. It was initially planned for powering wireless sensors on equipment in manufacturing plants, but other uses are also found for it, such as powering road and rail bridges sensors that will monitor the condition of these structures.

The generator, in the centre of the chip, powers an accelerometer

Similar energy-scavenging devices already exist, however according to the team leader, Dr Steve Beeby, this generator is up to 10 times more efficient.

Its size is a miniscule cubic centimeter. It uses vibrations around itself to make magnets on a cantilever in the device move, which in turn generates power. As for pacemakers, the beating of the human heart would be strong enough to keep the magnets inside the device shake. The produced power is microwatts in capacity, sufficient to fuel sensors or indeed pacemakers.

"The big advantage of wireless sensor systems is that by removing wires and batteries there is the potential for embedding sensors in previously inaccessible locations," said Beeby.

Work on the project was funded by the EU as part of a £10m Vibration Energy Scavenging (Vibes) project that focuses on using environmental vibrations to generate power.

Dual nanoimprint lithography system will cover many applications

A nanoimprint lithography system developed by Obducat AB of Sweden promises to enable high volume production of advanced micro and nanostructures for a wide field of applications.

The twin 'Sindre' imprint soft press technology is a fully automated system and employs compressed gas to ensure a uniform distribution of pressure across the whole imprint surface, thus producing a homogeneous residual layer of under 20nm with variations of less than 5nm. The technology is scalable up to any imprint size area, opening a path for the nanostructuring of large surface devices, such as flat panel displays and high density interconnects on printed boards.

The way it works is the first imprint step replicates the master stamp into an intermediate polymer stamp (IPS) using a polymer liner material. This ensures optimum stamp lifetime to reduce the costs associated with producing the intensive master and provides high throughput on large areas while guaranteeing thin homogeneous residual layers.

The IPS produced at the imprint head is transported to a second transfer unit (STU) imprint head, where the pattern is printed onto the final target substrate before the IPS is discarded.

Any loose particles on the master stamp are self-cleaned in the IPS process, leaving a clean master stamp for the next IPS to be produced. The STU process is carried out by printing the IPS once into a pre-heated, polymer-layer spin, coated onto the target substrate. Spin coating ensures a homogeneous residual layer thickness.

The IPS/STU system can be used to print on uneven surfaces, such as epitaxially grown semiconductors that are much used for producing LEDs. It adapts to the curvature, while any major growth defects such as spikes on the substrate are no problem.

Some of the first new products derived from the technology will be high brightness LEDs for projection devices and a new generation of hard drives employing discrete track recording (DTR). Initially, the expected wafer throughput is 30-90 wafers per hour, but higher values are anticipated after the system is optimised for specific processes.

SEM micrographs showing an IPS (left) and the STU imprint obtained on a hard disk substrate. This is a typical DTR hard disk drive pattern that will enter the market in the near future. It was generated with an Obducat electron beam recorder system.

IN BRIEF

● The engineering and technology sector in the UK is facing a growing recruitment crisis and there is little confidence that the problem will improve in the short or medium term, warned the IET (Institution of Engineering and Technology).
The organisation's annual skills survey of 500 companies highlighted that some 52% of businesses expect difficulties in recruiting adequate, suitably qualified, engineers over the next four years, up from 40% in 2006. As a result, they are turning to countries such as India, China and South Africa to plug the skills gap with 48% of companies recruiting from overseas in the last 12 months to cover specific skills shortages.

● A new report by IDTechEx entitled *"Active RFID and Sensor Networks 2007-2017"* states that the market for active RFID and sensor networks is growing faster than other RFID technology sectors.
Active RFID, where a battery drives the tag, is responsible for an increasing percentage of the money spent in the RFID market. It will rise from 12.7% of the total RFID market this year to 26.3% in 2017, a leap in value to over $7bn.
The primary factors creating this growth will be Real Time Locating Systems (RTLS) and ubiquitous RFID sensor systems – mainly disposable, including ones in the form of Smart Active Labels (SALs).

● A new report released last month forecasts a huge increase in personalised home learning with technological advances propelling e-learning to be the principle education method in the subject of maths by 2020.
E-learning has many years to go to reach its potential, which is to tailor learning to provide one-to-one education for all. Ron Van Der Meer, the report's author, said: *"The future is e-learning, but not the kind of e-learning as it is done today, which is still in the first motorcar design stage when they took the horse away from the cart and stuck a motor on top."*

Progress on GaN wafer production will give better devices

A novel process for producing gallium nitride (GaN). Inside a vacuum chamber, a wafer of glass or other substrate material is heated. Species of nitrogen are produced using a plasma and introduced into the main chamber, where they react with a vapour of trimethyl gallium (TMGa) to form a film of GaN on the wafer

A spin-off from the Macquane University in Australia promises to commercialise innovative technology to produce GaN more economically. BluGlass from Sydney claims that its remote plasma CVD (RPCVD) process for the lower than usual temperature of 700^0C deposition of GaN on glass substrates has significant advantages over current commercial processes that form GaN by conventional MOCVD at temperatures of typically 1000^0C on more costly heat-resistant materials, such as sapphire and silicon carbide.

The RPCVD process is expected to be scaleable beyond the size limits for the present sapphire and silicon carbide substrates. It will thus offer substantial improvements in production efficiency over current processes, resulting in lower wafer and finished device costs.

BluGlass workers are now optimising the process to improve device performance, while scaling up the current 2-inch wafer (which is the standard size for GaN) to 4-inch and later to 6-inch, with 8-inch or even larger wafers eventually being used.

GaN is vital for the production of many high performance electronic devices, especially high brightness LEDs that use 80% less energy and last up to one hundred times longer than conventional filament bulbs. It is also used to make blue and UV laser diodes, which are expected to be applied to the next generation of optical storage devices that may succeed conventional DVDs.

The total world market for GaN devices reached some $3.8bn in 2005, but could reach $100bn per annum if cheaper LEDs become widely accepted for general lighting products in homes and work places.

● The German government and private industry plan to invest up to €360m in research on organic photovoltaics. BASF and Bosch are founding members of the technology initiative of Germany's Federal Ministry of Education and Research (BMBF). The two firms, supported the government and other industrial partners, plan to develop new technology which will make the manufacture of solar cells cost-effective and increase the number of areas of application.

Last year, photovoltaic modules had a global market volume of €8bn. The segment is expected to grow by more than 20% annually until 2020. The BMBF will provide €60m for research to develop this market, whilst the industrial partners which include Merck and Schott, will add the rest €300m.

● Two academics, Dr Prashan Premaratne and Quang Nguyen, devised seven hand gestures which will control a TV or VCR and banish the remote control as we know it today.

The system comprises a web camera, gesture processing unit, hardware interface for the control unit and a universal remote control which is built into the control unit. The webcam captures the hand gestures and the software converts this into a signal which operates the remote controller. The series of commands devised by the researchers included switching the equipment on, turning the volume up and down, changing channels, play and stop – all in different lighting conditions and from different distances.

The solution is claimed to be accurate and power-efficient as it uses low processing power.

● A consortium comprising Filtronic, Forge Europa, QinetiQ, Thomas Swan and Cambridge University, has won a three-year contract under the DTI technology programme, worth just under £3m, to develop a route to low-cost LEDs for solid state lighting.

The project will demonstrate high-quality LEDs on 150mm silicon substrates to enable large scale production of LED lamps for solid-state white lighting. It combines expertise in III-V semiconductor manufacture, GaN growth capabilities and expertise in LED packaging. The project is expected to achieve significant progress on the solid-state lighting roadmap and provide a route for the UK to enter this major future market.

GOING DIGITAL IS NOT GREEN, SAYS A REPORT

Running costs of different TV types

New digital technology consumer systems are raking up the electricity bills for households and likely to need the output of 14 more power stations in the UK alone by 2020 to power them, says a new report from the Energy Saving Trust.

The worst offenders are large screen digital TVs and digital radio receivers, specifies the report, entitled *"The Ampere Strikes Back"*.

At present household appliances consume around a third of the average home's electricity. However, this is expected to rise to nearly half of household's electricity expenditure as homeowners buy more high-tech systems which are seen as less efficient or more energy-intensive (due to higher screen sizes, for example) than old technology.

"The old-fashioned, bulky cathode ray tube TV on average consumed about 100 watts of electricity when it was switched on," said Paula Owen, author of the report. *"What we are seeing now is a trend for much bigger flat-screened TVs. On average, we are seeing a three-fold increase in the energy needed to power these TVs."*

Whilst the average consumption of traditional analogue radio for example is only 2W, digital radios consume 8W, and if the radio is being listened to via a digital TV set or the set-top box, this rises to a staggering 100W. On standby alone televisions consume some 1.4% of all domestic electricity.

At present, UK consumers spend around £12bn per year on consumer electronics; some systems have been improved with technology when it comes to energy consumption, but others have not.

"Pretty much in every other sector, such as fridges and washing machines, we find that as the technology moves on the products get more and more efficient. Consumer electronics [seems] does not work like that," said Owen.

Top Five Tips

Board Layout for QDR-II SRAM

Good board layout techniques are important on any high-speed design, in particular QDR-II SRAM designs have some specific requirements. High-speed memory interfaces require stringent control in order to ensure the uncorrupted transfer of data. Below are a number of considerations, which need to be taken into account when laying out the design.

01. Trace lengths on buses must match to each other and their associated clocks to within 0.5 inch. There are two groups of buses and trace lengths must be matched to others in that group.

02. All traces are uncoupled. In particular, "K" and "KN" are not a differential pair.

03. Power rails, such as "VTT", must be planes, not traces.

04. Care must be taken to keep reference voltages, such as the QDR-II SRAM's "VREF", noise-free.

05. Simultaneous Switching Outputs (SSO) are a real concern in QDR-II SRAM designs, since there are usually can be no design constraints on how many signals can switch simultaneously.

This month's top tips were supplied by Sid Mohanty, Strategic Marketing Manager at Lattice Semiconductor.

If you want to send us your top five or ten tips on any engineering and design subject, please write to the Editor: svetlana.josifovska@stjohnpatrick.com

PHYSICAL SYNTHESIS FACES UP TO DESIGN CHALLENGES

GARY MEYERS, PRESIDENT AND CEO OF SYNPLICITY, CAN'T WAIT FOR THE MOMENT WHEN THE EDA INDUSTRY TRANSITIONS FROM LOGIC SYNTHESIS TO PHYSICAL SYNTHESIS, A PLATFORM WHICH, HE SAYS, WILL TACKLE THE RISING COMPLEXITY AND DEEP SUB-MICRON ISSUES

When I started in the FPGA business, FPGAs contained only tens of logic cells and designs were easily captured with graphical schematic capture software. Benefiting from Moore's Law trends, FPGAs grew bigger over time and by the early 1990s, FPGAs had grown to thousands of logic cells and capturing designs with schematics was becoming tedious.

Device complexity and capability were out-running the capabilities of the design entry tools. Fortunately, there was a solution waiting in the wings: logic synthesis. Logic synthesis had been commercially developed a few years before, mostly aimed at the ASIC market, and combined the efficiency of RTL abstraction with a concise text-based syntax for design description. In time, these benefits, coupled with logic synthesis engines optimised for FPGAs, caused a fundamental shift in the way FPGAs were designed. By the late 1990s, logic synthesis had nearly replaced schematic capture.

Logic synthesis has served us well. Because of this technology, today's designers routinely complete designs that consist of tens or even hundreds of thousands of logic cells. We are however starting to see some strains on the technology as it is tasked with ever larger design challenges. For example, with larger devices and longer runtimes, timing closure has become a difficult problem. The designer may know that he can reach his required speed, but getting there requires multiple time-consuming design flow iterations to converge. This condition has become common enough that various tool strategies for coping with it have reached the market, such as floorplanning or post-synthesis optimisations. But these solutions have been only partially successful in solving the timing closure problem.

> **We are starting to see some strains on the technology as it is tasked with ever larger design challenges**

As it turns out, current design challenges go beyond design size and complexity. In this deep-submicron era there are new issues to be reckoned with. Power has now become a problem in many FPGA designs. Soon, other issues encountered in deep-submicron SoC design, such as on-chip clock variation, will start to resist hardware-centric solutions or cap performance gains well below physical device limits. FPGA vendors are doing an excellent job of handling these effects so that they are invisible to FPGA users but, from here on, the task is going to get more difficult with each generation of technology.

Just as in the early 1990s, we now need a transition to a new design technology that can provide the basis for coping with both the increased design complexity and the threatening intrusion of deep submicron issues into the realm of digital FPGA design. What we need is a transition from logic synthesis to physical synthesis. Because physical synthesis simultaneously combines logic optimisation with placement and routing, the technology has the information needed to tackle problems ranging from timing closure to power optimisation. Physical synthesis provides the needed platform to tackle both rising complexity and deep sub-micron issues. Like with previous transitions this one is only a matter of time. ■

MILITARY EXPERTISE SPILLS INTO
CIVILIAN MARKETS

By Joel Bainerman

The export of Israeli military products has always been a major component of the Israeli industrial landscape. Today, Israel is one of the top five defense exporters in the world, behind only the United States, Russia and Britain with 2006 sales of $4.4bn.

The fact that Israel is a major arms developer is nothing new. What is new is that the major market penetration Israeli firms – once solely focused on the worldwide military market – are making in civilian security and aviation markets.

Israeli exports of homeland security equipment reached $300m last year, up 22% per year since 2002.

As a result of the push into additional markets, Israeli start-up companies in the sector are attracting international venture capital.

For instance, Scent Detection Technologies (SDT) recently completed a Series A $5m round of financing from Menlo Park based, veteran, venture capital fund, Sequoia Capital. The company has developed a low-cost trace detection sensor and mobile handheld "sniffing" devices. Its first product is the Mini-Nose, a portable, highly sensitive, non-invasive device capable of trace and particle detection aimed at markets such as aviation security, military, police, border-control, mass transit, prisons and ports.

Biometrics security company, IDesia, has raised $3m in investment capital from France-based venture capital firm, Partech International. The company's biometric technology can identify the unique natural electro Bio-Dynamic Signature (BDS) of individuals. The electronic signature is based on the electric signals of humans, which are automatically and continuously given off as our beats, brain whirrs and breathing functions. No two people have the same signature.

The core basic principle of IDesia's technology is that a user will touch the tiny BDS sensor for anything up to 8 seconds, depending on the levels of accuracy required, and the sensor will then identify their electronic signature. The technology is ideal for the consumer electronics market because the sensors are smaller and more durable than other methods, they are easy to use, require little power and are much, much cheaper than alternatives; all of which are critical factors for acceptance in the mass market.

The boom in exports to civilian security markets is also fuelling basic research in Israel in these areas. Recently, a team of Israeli scientists created a molecule that can function as an ultra-miniaturised version of a keypad locking mechanism. Researchers at the Weizmann Institute of Science in Rehovot have developed the first "molecular keypad lock" that operates similarly to the electronic locks on today's ATM machines, but uses a molecule in place of electrical circuits. The result is a lock that offers much greater security than simple molecular logic gates.

The molecule, synthesised in the lab of organic chemist Professor Abraham Shanzer, is composed of two smaller linked units – fluorescent probes – separated by a molecular chain to which iron can bind. One of these probes can shine bright fluorescent blue and the other fluorescent green, but only if the surrounding conditions are right. These conditions are the keypad inputs; rather than the electric pulses of an electronic keypad, they consist of iron ions, acids, bases and ultraviolet light.

The challenge the scientists faced in creating the functional 12-key prototype keypad lock was in generating sequences that can be distinguished one from another. On a calculator, entering the sequence 2+3+4 will yield the same result as 3+4+2, but a keypad lock set to one password

ODF Optronics's Eye Ball R1 in action

(234) will not open for the other (342). "The numbers are not relevant, the order is," says Shanzer.

The Israeli scientists found that by controlling the opening rate of the logic gate within the reaction time-frame, they were able to produce different, distinguishable outputs, depending on the input order. By adding light energy, which also influences the glow of the molecules, they were able to produce a molecule-size device that lights up only when the correct chemical 'passwords' are introduced.

"Faster and more powerful molecular locks could serve as the smallest ID tags, providing the ultimate defence against forgery," explains Shanzer. "Further ahead, the technology could also be used to design smart diagnostic equipment that can detect the release of biological molecules in the body that could indicate disease or exposure to chemical or biological weapons."

ENHANCING HOMELAND SECURITY

Due to the large budgetary allocation by the US government towards the goal of safeguarding America's borders, the "homeland security market" has also been attracting many Israeli firms.

One of them, Tel Aviv based ODF Optronics, has introduced a new electro-optical device which was originally designed to enable Israeli army troops and law enforcement personnel to gather intelligence by viewing images and listening to voices in a target area, prior to an assault into a room or entering subterranean spaces, such as cellars, wells or tunnels.

Called the Eye Ball R1, the palm-sized, camera-laden device was designed to be thrown into environments deemed too risky for personnel to enter, like collapsed buildings or for counter-terrorism operations. Capable of producing a 360-degree view, it works by dispatching images and sound wirelessly to a remote receiver. Weighing about half a kilo, the Eye Ball R1 is encased in a rugged rubber and polyurethane housing. That allows it to be thrown through windows or bounced off walls. When it comes to a rest, the ball

automatically stabilises itself and begins to wirelessly transmit footage and sound, up to 200 meters away.

This information is gathered by a remote control device, called the Personal Display Unit (PDU), which includes a colour video screen. Its omni-directional camera rotates at 4rpm until a target is identified and gives the operator a 55-degree vertical and 41-degree horizontal field of view, as well as invisible near-infrared spectrum night vision capability for low-light deployment/night operations.

The main users are the military and law enforcement units as well as first responders - firefighters, emergency medical technicians search-and-rescue teams, and correctional facilities.

Another unique anti-terrorist product Israel has put on the market is a rifle that incorporates the most sophisticated night vision technology around. Developed by Haifa-based Corner Shot, the rifle uses a unique application of night vision technology to enable the shooter to shot at an angle and around corners, allowing police and soldiers to shoot from behind cover. The soldier/policeman can also observe the conditions around a corner using a video camera mounted on the front

part of the barrel, which transmits the picture to a display installed next to the rifle's butt. The camera can be exchanged with night vision equipment and can transmit pictures of the battlefield to a command centre. In this way the soldier can stay behind cover and still watch what is happening, providing firepower if necessary.

Elbit Systems, one of Israel's largest defence firms, has unveiled a robot for urban warfare and counterterrorism. Called "Viper", it is capable of shooting bursts of automatic fire and throwing fragmentation, smoke and stun grenades.

The Viper, which is about a meter long and weighs approximately two kilos, is powered by a special electrical engine and operated by remote control, governed by a program implanted in its 'brain' in advance. It is capable of climbing stairs, getting past obstacles and, at the same time, checks what is going on around it by means of a system of sensors. It is equipped with a special nine-millimetre calibre Uzi machine gun, on which a laser pointer has been installed. The Viper is carried to the battlefield by a soldier on his back in a special carrier. When it is necessary to infiltrate a building safely where, for

IDesia's biometric technology

example, armed terrorists are hiding, the soldier lowers it to the ground, turns it on and from that moment controls it from a distance.

In Jerusalem, Camero Technologies developed a system using FCC compliant micro-power Ultra Wide-Band (UWB) signals to see through any type of wall, including steel-enforced concrete walls. UWB technology eliminates unwanted RF "chatter" to produce three-dimensional pictures of the space behind a wall from a distance of up to 20 meters, with extremely high resolution. The device calculates the distance and orientation of people, furniture and weapons on the other side – up to 300 feet into a building, in real time. The image is displayed to the user in an intuitive way and can be presented in a semi real-time mode (i.e. "jumpy" video) or in a 'stills' mode (utilising longer exposure time).

CIVILIAN AVIATION IS NEXT

With its established expertise in avionics for military use, Israeli firms have adapted this know-how to the civilian aviation market, mostly in the security portion of that market. Recently, The Israel Civil Aviation Authority (ICAA) approved a flight guard commercial airliner protection system built by Elta Systems Group for installation on Boeing 767 passenger jets.

Flight countermeasure dispenser system launches flares at instructions from a radar-based missile approach warning system (MAWS). The system has two main components: Elta's autonomous Pulse Doppler Radar system, a series of six miniaturised pulse-Doppler sensors to give all round coverage to detect the launch of surface-to-air missiles; and a counter-measure dispensing system that jams and diverts heat-seeking missiles.

The radar missile warning system has higher reliability than infrared systems, reducing false alarm rates from one every two hours of flight time to one per year. The flares are intended to burn for a very short time to avoid any damage even if discharged in low altitude.

Another company, Steadicopter Ltd, has announced the completion of the first model of an unmanned helicopter. The helicopter performs fully autonomously, including all

the aspects of the flight: takeoff, hovering, waypoints flight and landing, while transmitting video images in real time.

The Steadicopter craft weighs just 6.5 kilos and is 1.5m in length. It can be controlled by a laptop up to a range of about five kilometres. The chopper can stay in the air for 90 minutes, navigate using pre-defined GPS coordinates and relay information back, using day and night (thermal) cameras. Flight details and routes are fed from a standard PC at the ground station before the flight, with the possibility of interfering, changing mission, orders and commands (directions, flight speed and altitudes) throughout the course of the flight.

The first product of Steadicopter is designed to fly at around few hundreds feet, with an operation range of 10km and can adjust to winds of up to 25 knots. Changes in direction, flight speed and altitude can be made throughout the flight. The helicopter can be fitted with cameras that survey areas up to 8.8 miles away and transmit real time video images.

The craft uses GPS, gyroscopes and various, other, off-the-shelf instrumentation to guide and control the helicopter. It can come in a variety of configurations that can fit military and security applications such as mine detection, stealth scouting, border patrol and police surveillance, as well as civilian applications such as high voltage line inspection, media coverage and agriculture purposes.

Urban Aerodynamics Ltd in northern Israel has developed an aviation vehicle that can fly amid skyscrapers and park inside buildings. Its purpose is to become an effective life-saving rescue system.

Called the X-Hawk, the vehicle is a rotorless Vertical-Take-Off and Landing (VTOL) vehicle. Unlike a helicopter, the X-Hawk's propellers are not extended, but incorporated into the body of the aircraft, enabling it to pull up close to the windows of tall buildings without danger of collision.

These unique characteristics make the X-Hawk suitable for rescue and law enforcement work, evacuating injured people from high-rise buildings, high speed pursuit and other daring police activities. The X-Hawk is expected to be able to

Viper Robot by Elbit

achieve a maximum speed of 200kph and to remain airborne for up to an hour and a half (like small helicopters).

Dr Raffi Yoeli, the founder of the concept and the company said: "In a regular helicopter, if you want to move left, you have to first tilt to the left. Then you have to correct the movement by tilting to the

"WITH ITS ABILITY TO GO WHERE NO HELICOPTER HAS GONE BEFORE, THE X-HAWK CAN DOCK ON THE SIDE OF A BUILDING AND RESCUE PEOPLE FROM TOP FLOORS OF A BUILDING"

right in order to straighten out. We eliminate those movements – a very important point when you're working in constricted air space like an urban environment."

The vehicle's control system is a row of vanes on the top and bottom of the panel; each of the ducts looks and works like a Venetian blind. They can either turn in unison or in a few degrees in one direction or the other.

With its ability to go where no helicopter has gone before, the X-Hawk can dock on the side of a building and, thus, be used in medical response teams to rescue people from top floors of a burning building, for example. ■

LISTEN TO YOUR RADIO MODULE

by Myk Dormer

To the non-specialist, even the simplest of radio modules can seem a strange and demanding device. It requires ripple-free power, proper mounting on a good groundplane and the provision of a carefully selected and located aerial. And even then, exacting interface requirements must be met, often the datastream must be specifically coded and bandwidth and timing restrictions worked within.

And then everything is guaranteed to work, or is it? Unfortunately, no. Even when good practice has been followed and the module is being used well within its parameters, there will occur occasions when

"SOMETIMES, EVEN WHEN EVERY GUIDELINE AND RULE HAS BEEN FOLLOWED AND THE RADIO MODULES HAVE BEEN RE-TESTED BY THE MANUFACTURER AND FOUND GOOD, THE LINK DOES NOT PERFORM"

the implemented link achieves nothing like its calculated, or promised, range. In rare cases it may fail to operate at all.

Sometimes, this can be due to simple bad practice; like using an inappropriate aerial, a noisy power supply, a badly coded datastream or, occasionally, even a defective radio. These cases are relatively easy to diagnose. But, sometimes, even when every guideline and rule has been followed and the radio modules have been re-tested by the manufacturer and found good, the link does not perform. The problem is most likely not in the module itself, but rather it is the result of an interfering signal. An interferer can originate from outside the customer's equipment (from nearby electronic equipment, or from other radio systems) or from within (from switch mode power supply oscillators and from digital clock signals). It can be a single frequency 'spur'; it can be the modulated carrier of another radio; or it can be a 'forest' of unstable, digital noise.

A typical ISM band radio receiver can have an input sensitivity better than -115dBm (about 400nV). An unwanted signal on a channel of this order of magnitude will be undetectable on a scope, or a simple spectrum analyser, but will still compromise the link range significantly.

So how do you go about identifying and solving such interference? There are a few simple diagnostic methods that can save a great deal of heartache and lost time:

- Once you've got the suspicion that things are not all they ought to be, you need to be able to reliably reproduce your findings. A repeatable test set-up will be needed in your lab if you are to make any meaningful conclusions.

  Conducted tests, where a signal generator connected to the aerial port replaces the transmit end of the link, are useful but only if interference is entering the radio by radiation through the casing or conduction through the interface.

  Otherwise, a radiated test will be needed.
- Placing a large value attenuator on the transmitter output or replacing the transmitter aerial with a dummy load will reduce the range to manageable proportions, or, if available, a signal generator feeding an aerial can replace the transmitter-end of the link.

- Monitor the receiver audio. You can tell a lot about a link by looking at the analogue AF output of the receiver on a scope or, even better, listening to it, with a small audio amplifier and speaker. The output of a simple FM receiver in the absence of a signal should be white noise. Anything else is symptomatic of problems. Warbles, tones or 'burbling' noises should suggest either digital noise from a CPU, or another transmitter in your vicinity, while the absence of noise ('quieting') – when you remove your wanted signal indicates an un modulated interfering carrier (typically a harmonic of a stable clock oscillator).
- Monitor the channel. Use an inexpensive 'scanning' receiver set to your operating channel frequency or better build yourself a monitor receiver from another, identical radio module and aerial. Add a voltmeter on the "signal strength" output and an audio amp and speaker on "AF out". Power it from batteries.
- Check that there are no unexpected signals present on your channel. Conduct this survey both in the lab and at your final installation location if possible, and be aware that external interferers are frequently intermittent.
- Eliminate your own hardware. Observe the output of your receiver when your

entire design is operating and when only the receiver itself is powered, preferably from a temporarily connected battery.

By replacing the monitor receiver aerial with a small 'search coil' (2-4 turns of 1mm diameter wire wound around a 5mm former), it can be used at shorter range to 'sniff out' an interfering device. If the signal is found to have an external origin (another radio system, or a piece of equipment over which you have no control) then you have very few options: Accept the reduced range or change frequency to a 'clear' channel.

Module manufacturers now often provide several frequency versions of even the simplest modules, in identical pin-outs, while many of the more sophisticated units are multiple-channel by design.

Occasionally, the radio design itself will be found to have a 'blocked' channel, when the channel frequency falls onto a harmonic of one of the internal reference oscillator frequencies. Your module manufacturer should be able to advise you of this, although all too few will actually admit it. If the interferer originates in your hardware, then there are certain modifications and tests that may help:

● Improve general decoupling and filtering. This ought to be general good practice, but ensure that:
a) HF decoupling (47p-1n, surface mount) caps are provided adjacent to all fast logic devices;
b) PCB tracks carrying fast digital or high frequency analogue signals are short and direct;
c) Good grounding practices (groundplanes or low-impedance earths) are used;
d) Ensure the interface connections to the radio module are clean and add filtering networks if necessary.
● Conduct an 'audit' of all the clock sources in your design. If the harmonic of a processor clock or other oscillator falls on the channel frequency, even if it is a very high order multiple (I've seen trouble from the 35th harmonic of a crystal clock!), then this is a possible culprit. Temporarily stopping the clock should be seen to remove the interferer and, if system design permits, changing to a different frequency clock can be a cure. But beware: the interferer is still

present; it's just moved and, if serious enough, could still be a general EMC problem.

● If the mechanical design constraints permit, locate the circuitry inside a shielded enclosure (metal or treated plastic) with the aerial outside. Sometimes, the addition of internal shield cans over especially 'active' circuits can be effective.
● Carefully examine external connections. Power supply or interface cables are frequently found to radiate interference generated inside a unit, so it is good practice to limit the slew rate of external signals and include filtering (RC, LC or ferrite) wherever possible.

If the interference originates from the hardware and no easy cure can be found, then there are some considerations for the (inevitable) redesign.
● Identify noisy sub-assemblies (displays, power supply modules). Conduct radio interference tests of alternative units.
● When re-draughting the PCB, p ut series 'stopper' resistors on fast logic lines (100Ω is frequently sufficient to reduce radiation without corrupting the logic waveforms), include a groundplane and filter all connections at the board edges.
● 'Sleep' as much of the digital hardware as possible. Modify the operation of the software so as little as possible that could cause noise, such as updating displays, reading EEPROMs, happens during the receiver's operating time-window. And, as always: test everything.
Good luck!
Myk

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

RANGE TESTING THE RADIO ALONE

Sometimes, it will seem that a radio link isn't quite right; range is not quite what was hoped for, but isn't unusably poor, and no absolute cause can be identified.

In such a case it is useful to conduct a range test on the radio/aerial/environment that is independent of the possibly interference producing customer hardware. This will give an estimate of what the modules themselves are capable of.

The easiest way to conduct this test is to use a proprietary encode/decoder chip pair to replace the customer hardware. The coder/decoder should be operated at approximately the same data-rate as that used in the final design, and a simple 'push button/flash light' interface.

Suitable devices for such a test would be the Holtek H12 type coder, the Radiometrix CTR44, or its evaluation kit:
www.radiometrix.co.uk/products/nbek.htm

RAPIDIO AND ETHERNET –
THE PRACTICAL TECHNICAL COMPARISONS

E thernet is currently the incumbent backplane technology across a wide range of storage, wireless, wireline, military, industrial and other embedded applications. This is due to developers continuing to move away from proprietary implementations in an effort to reduce development time while increasing performance and functionality, all at a reduced cost.

However, as data rates continue to increase, it has become apparent that many high-performance applications exceed the limits of this traditional protocol. Designing an efficient embedded backplane interconnect with excellent performance requires addressing a number of key design challenges, including header efficiency, protocol processing efficiency, effective bandwidth and quality of service, while strictly managing cost. To meet these challenges, many developers are turning to RapidIO technology as an alternative to Ethernet.

Many of the differences between Ethernet and RapidIO technologies stem from their initial design constraints. Ethernet was designed to connect a large number of endpoints with a flexible and extensible architecture, leading to the choice of a simple header and support for a single transaction type. As Ethernet was originally intended to connect computer workstations, hardware is only required to identify packet boundaries, necessitating a relatively large software stack to manage protocol processing. While this serves well in LAN and WAN applications because of the presence of powerful processors, the hardware/ software trade-off imposes a formidable performance bottleneck in high-speed embedded applications.

RapidIO technology was originally conceived as a next-generation front-side

IN THE JULY ISSUE WE STARTED A SERIES ON ETHERNET AND RAPIDIO INTERCONNECT TECHNOLOGIES: TYPES, TECHNICAL COMPARISONS AND OVERVIEW, PRACTICALITIES AND DESIGN CONSIDERATIONS. IN THIS – FOURTH – ARTICLE, GREG SHIPPEN, SYSTEM ARCHITECT FOR THE FREESCALE SEMICONDUCTOR DIGITAL SYSTEMS DIVISION AND A MEMBER OF THE RAPIDIO TRADE ASSOCIATION'S TECHNICAL WORKING GROUP AND STEERING COMMITTEE, FOCUSES ON THE TECHNICAL ASPECTS OF RAPIDIO

bus for high-speed embedded processors. The value of a front-side bus that could also function as a system-level interconnect was recognised early in the specification's development. As a

consequence, RapidIO technology was designed with a focus on embedded in-the-box and chassis control plane applications, emphasising reliability with minimal latency, limited software impact, protocol extensibility and simplified switches while achieving effective data rates from 667Mbps to 30Gbps. Protocol processing takes place in hardware and supports read/write operations, messaging, data streaming, QoS, data plane extensions and protocol encapsulation, to name a few of its capabilities.

HEADER EFFICIENCY

Ethernet's extreme flexibility is one of the main sources of its inefficiency. Use of a simple, generalised header enables Ethernet to add higher-level services as new protocol layers but even basic services require additional header fields. Ethernet's firm requirement for backwards compatibility also introduces inefficiencies such as maintaining preamble and IFG fields required for the original half-duplex shared coax PHY but for which more recently defined PHYs have less need.

As there is no opportunity to optimise the overall header, the presence of duplication across multiple headers increases parsing complexity and latency. While processing of Ethernet headers can be optimised as custom protocol stacks on top of Ethernet Layer 2, the cost of supporting custom stacks across multiple vendors or even multiple generations of the same vendor hardware can become prohibitive. As a consequence, many systems take a performance hit to utilise UDP or TCP/IP rather than communicate directly through the more efficient lower layers.

RapidIO technology was designed with optimisation of the header in mind.

Figure 1 above and below: Many embedded Ethernet applications use TCP/IP to handle packet loss because of off-the-shelf software support. While higher-layer protocols simplify application development, they also add substantial overhead – 40 bytes in the case of TCP/IP – reducing overall bandwidth efficiency

Embedded backplanes need to support significantly fewer endpoints than LANs, so a smaller address field can be utilised (i.e. only one or two bytes compared to Ethernet's MAC address of six bytes plus four bytes when IP is used). The specification provides header support for common services, such as read and write transaction types, which require additional header layers in Ethernet. Redundant fields are removed, duplicate addressing schemes avoided and fields compressed where possible.

PROTOCOL EFFICIENCY

Ethernet offers only best-effort service unless higher-layer protocols are employed to handle packet loss. Many embedded Ethernet applications use IP with TCP to handle packet loss because software support is widely available and understood.

TCP/IP, while simplifying application development, adds significantly to the overall Ethernet header, introducing 40 bytes of overhead (see **Figure 1**). UDP can also be used when this overhead is too high but then reliable delivery, if required, must be implemented in a proprietary fashion.

Overall, Ethernet is a fairly flexible standard supporting many optional layer 3+ protocols developed by different standards organisations for a broad range of applications. However, because no single, unified specification is uniformly implemented, these have resulted in a diversity of actual implementations and, as a consequence, increased Ethernet stack complexity. Even necessary technologies as TCP/IP offload engines (TOE) are plagued with many proprietary implementations, each with its own unique partitioning of processing between hardware and software and vendor-specific Ethernet stack that ties down developers (see **Figure 2**). There isn't even a standard SERDES PHY yet for 1 or 10 Gigabit Ethernet backplanes.

The RapidIO specification achieves many of its advantages because it offers a single, uniform protocol with consistent layering, managed by a single standards organisation. RapidIO technology also guarantees delivery by providing end-to-end error checking, retrying link errors, not allowing switches to drop packets and supporting virtual channels. Since RapidIO technology directly implements the protocol in hardware, headers can be processed in a straightforward and less processor-intensive manner than equivalent Ethernet implementations which utilise partial hardware offload and custom stacks, resulting in lower implementation cost, reduced overall design complexity and more stable interoperability between vendors.

EFFECTIVE BANDWIDTH

Ethernet supports a payload size from 64 to 1500 bytes (up to over 9000 with jumbo packets) and its efficiency is best with a maximum payload, although this comes at the cost of increased latency jitter. RapidIO technology transports 1 to 256 bytes, balancing large payload jitter against small payload inefficiency. RapidIO technology achieves better efficiency across the payload sizes most common to embedded backplane applications.

For control plane applications that cannot tolerate packet loss, an Ethernet fabric must be significantly over-

provisioned to avoid packet loss and limit associated latency and jitter. Given 25-35% usage for many applications, this translates to a sustainable effective throughput for layer 2 traffic of 250Mbps for 1GE and 2.5Gbps for 10GE, depending on average packet size. Note that performance is defined, not by PHY symbol rate, but rather the effective rate

"UNIQUE PARTITIONING OF PROCESSING BETWEEN HARDWARE AND SOFTWARE ENDS UP TYING DEVELOPERS TO VENDOR-SPECIFIC IMPLEMENTATIONS"

in which a protocol reliably transports data. Additionally, even with over-provisioning, end-to-end latency can still run in milliseconds since traffic must traverse multiple software stacks.

By implementing protocol processing in hardware, RapidIO technology greatly reduces effective latency in comparison to Ethernet and can deliver substantially higher fabric utilisation in complex topologies. For control plane applications, link-level error correction minimises latency jitter caused by soft errors, potentially reducing end-to-end latency below 500ns.

Throughput is also affected by overhead for operations such as reading, writing and messaging. Ethernet RDMA provides read and write operations, but as a layer 4 protocol, its high overhead is not well-suited for small control-oriented load/store operations. TCP/IP services resemble RapidIO messaging but where RapidIO messaging supports convenient 4kbyte messages and is often fully implemented in hardware, TCP/IP supports 64kbyte messages that are much more dependent upon software for processing. Additionally, the RapidIO standard defines a protocol for keeping caches coherent across the interconnect, a feat ineffective to implement in Ethernet due to low header efficiency, high latency and inconsistent levels of hardware support.

Extensive packet handling by Ethernet switches also increases overall complexity and processing load. When IP packets are routed, numerous fields must be updated and the FCS recalculated. RapidIO switches typically only update the AckID field which is not covered by the CRC and so does not force a recalculation.

QUALITY OF SERVICE

Quality of Service (QoS) is essential for many backplane applications. While Ethernet through TCP/IP can support millions of individual streams and differentiate traffic by port number and protocol fields, no universally used class of service (CoS) field exists. The RapidIO standard defines up to six flows that can be considered prioritised classes of service. Through the use of Type 9 encapsulation and virtual channels, millions of streams can be differentiated as well.

QoS is also affected by Ethernet's best-effort service, which commonly manages congestion by dropping packets, leading to latency jitter. Since flow control belongs to upper-layer protocols such as TCP, congestion cannot be promptly managed to prevent packet loss.

This lack of short-term, link-level flow control requires larger endpoint receive buffers to avoid overruns. Further exacerbating latency is the fact that error detection and recovery occurs at the system rather than link level. Thus, timeouts exist only at layer 3 and above, and are managed by offload hardware in the best case or software in the worst case, resulting in much longer timeouts and significantly increased latency jitter. Additionally, no standards exist for hardware-based recovery such as retries or timeouts, and so Ethernet drops packets for a significant period of time before failure is detected.

While there are protocols such as bidirectional forwarding detection which exchange packets to detect failures, these continuously impose bandwidth overhead dependent on the responsiveness required.

Per its spec, all RapidIO networks must provide a minimum level of prioritised service to implement logical layer ordering

and deadlock avoidance. This also improves average latency since packets marked with a higher relative priority must make forward progress since they might be responses. Optionally, switches can reorder packets from different flows and offer head-of-line blocking avoidance as well as other QoS features. With multiple flow control mechanisms, RapidIO networks allow congestion to be managed before there is a significant impact on network performance.

The RapidIO standard also defines a link-layer protocol for error recovery and various hardware-based link-to-link and end-to-end timeout mechanisms, enabling virtually all errors to recover at the link level without software intervention, substantially lowering latency jitter. Also, because RapidIO technology links carry valid traffic at all times, a broken link is promptly detected locally at the link level. As a result, failure rates, defined as undetected packet or control symbol errors, are significantly less than the hard failure rate of the devices on either end of the link, depending upon operating conditions.

COST

From a silicon standpoint, it might seem that the often touted high-volume cost

economies of Ethernet would provide Ethernet a significant advantage over RapidIO. While this is likely true for 4-8 port GE switches used in LANs, Ethernet switches for use in many embedded applications require more specialised functionality, such as VLAN QoS and SERDES PHYs, significantly reducing the number of accommodating vendors, overall shipping volume and, therefore, cost economies.

Additionally, RapidIO technology in general assumes a maximum backplane or board-level channel of 100cm using copper traces on FR4-like materials. Ethernet PHYs for twisted-copper pair must support channels 100 times longer than the RapidIO technology's and, since Ethernet cabling assumes bundling with many other similar pairs, it must tolerate significantly more crosstalk. Together, these result in significantly higher PHY complexity than is actually required for backplane applications. As a result, some switches provide a non-standard SERDES PHY for these applications.

Computing true PHY cost must also be done carefully. For example, one commercially available RapidIO endpoint supporting messaging was only 25% larger than a straightforward Gigabit Ethernet controller without full TOE

capabilities. Likewise, a four-lane RapidIO SERDES is about 50% larger than a single XAUI lane. This suggests that the silicon complexity required for endpoints is comparable.

For its part, RapidIO technology offers 2.5 times more effective bandwidth per link than GE. Yet, the cost per port of a 16-port RapidIO switch is competitive with or better than similar GE switches. For applications requiring more than 1Gbps, the only alternative for Ethernet is 10Gbps. Today, RapidIO technology offers higher effective bandwidth for payloads less than 1024 bytes at lower cost, even without taking into consideration the cost imposed on Ethernet endpoints to process protocol stacks at 10Gbps. Such processing also has a significant impact on power consumption as a GHz-class processor to terminate each GE link increases power by watts.

Virtually any application layer service can be supported by either Ethernet or RapidIO. The difference between the two technologies, however, lies in their individual inefficiencies and the level of hardware processing supported. Ethernet has a long history in the LAN which, because of backwards compatibility, header and protocol inefficiencies, software dependence, complex and proprietary offloading mechanisms, and lack of implementation standards, makes it less than ideal choice for backplane applications.

The RapidIO standard was specifically designed to provide optimised performance for embedded applications. By eliminating unnecessary overhead, implementing protocol processing in hardware, providing a stable specification resulting in standardised implementations, focusing on maximising effective throughput, supporting sufficient quality of service mechanisms and providing cost-effective silicon supported by an extensive ecosystem, RapidIO technology provides the interconnect technology that will carry embedded developers well into the future. ■
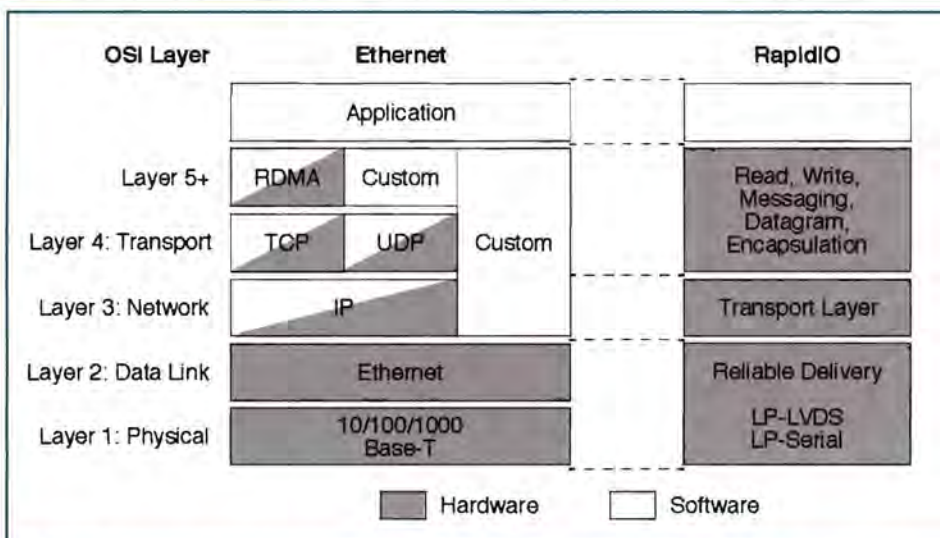

Figure 2: While Ethernet is a fairly flexible standard supporting many optional layer 3+ protocols, lack of a single, uniform implementation results in a diversity of actual implementations and, as a consequence, increased Ethernet stack complexity. Unique partitioning of processing between hardware and software ends up tying developers to vendor-specific implementations

To download a complete white paper on this topic visit www.rapidio.org/outperforms_whitepaper/

CRYSTAL CLEAR

Liquid crystal display (LCD) units used for the purposes of moving picture reproduction are some of the more popular flat panel displays. Like all flat panel displays, LCDs employ a matrix structure in which the active element, in this case a liquid crystal, forming the pixel cell is located at the intersection of two electrode buses.

So, what is a liquid crystal? A liquid crystal is neither crystal nor liquid. It exhibits liquid-like as well as crystal-like properties. This feature is a result of the liquid crystals' comparatively elongated molecules and their structure. Though a liquid crystal is a natural material, the liquid crystal which is used for LC displays is a multi-component mixture that is artificially created by blending of biphenyl, cyclohexane, ester and the like.

MATRIX LCDS

There are two matrix LCD technologies: passive matrix LCD (PMLCD) and active matrix LCD (AMLCD). In the passive-matrix LCD, pixels are addressed directly with no switching devices involved in the process. The effective voltage applied to the liquid crystals must average the signal voltage pulses over several frame times, which results in a slow response time greater than 150ms and a reduction of the maximum contrast ratio. The addressing of a PMLCD also produces a kind of crosstalk resulting in blurred images because non-selected pixels are driven through a secondary signal-voltage path. This place a limit to the number of pixels that may be used in a display and with it a limit on the maximum resolution.

In the AMLCDs, on the other hand, a switching device is used to apply the voltage across the liquid crystal (see Figure 1) and, so, a better response time becomes possible. In contrast to passive-matrix LCDs,

AMLCDs have no inherent limitation in the number of pixels, and they present fewer cross-talk problems.


Figure 1: *Active matrix LCD*

There are several kinds of AMLCD depending on the type of switching device used. Most use transistors made of deposited thin films, which are accordingly called thin-film transistors (TFTs). The most common TFT semiconductor material is made of amorphous silicon (a-Si). Amorphous silicon TFTs are amenable to large-area fabrication using glass substrates in a low-temperature (300°C to 400°C).

An alternative TFT technology, polycrystalline silicon, normally known as polysilicon or p-Si is costly to produce and especially difficult to fabricate when manufacturing large-area displays. Nearly all TFT LCDs are made from amorphous silicon because of the technology's economy and maturity, but the electron mobility of a p-Si TFT is 100 times better than that of an a-Si TFT. This makes the p-Si TFT a good candidate for a TFT array containing integrated drivers, which is likely to be an attractive choice for small, high definition displays such as view finders and projection displays.

TFT CELL DRIVE

In the TFT LCD, switching transistors are provided for each pixel cell as shown in **Figure 2**. One side of each liquid crystal cell is connected to its own individual TFT, while the other side is connected to a common electrode which is made of transparent ITO (Indium tin oxide) material. This is necessary to ensure high aperture ratio. Aperture ratio is the ratio of the transparent area to the opaque area of the panel. A cross-section of a TFT is shown in **Figure 3**.

Unlike the CRTs in which the phosphor persistence allows for continued luminance of the picture even after the electron beam has moved to scan other lines, in flat display applications no such persistence exists and refreshing of pixels to produce natural moving pictures becomes difficult as the number of pixels increases. Hence the need for a pixel cell 'memory'. A charge on a storage capacitor C_S is used for this purpose as illustrated in Figure 2. Each cell consists of three sub-pixels (RGB), normally referred to as cells. Each cell contains a liquid crystal driven by a TFT acting as a switch. The liquid crystal is placed within two electrodes. One electrode is connected to the TFT's source electrode and the other goes to a common electrode.

The TFT-LCD panel is scanned line by line. Each line is selected in turn by a V_{SEL} pulse to the line (or gate) electrode bus. Once a line is selected, the pixel cells along that line can be addressed and their luminance levels set by a voltage applied via a source driver to their corresponding data (also known as source or column) electrode. The source driver supplies the desired voltage level known as the greyscale voltage representing the pixel value, i.e. the luminance of the pixel cell. The storage capacitor C_S is charged and this charge maintains the
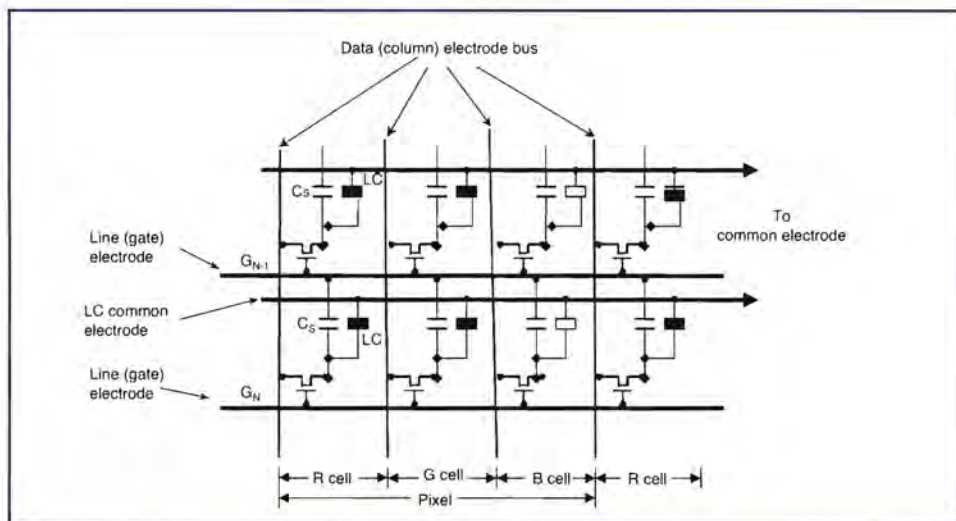
Figure 2: Equivalent circuit for TFT LC

Figure 3: TFT cross-section

luminance level of the pixel cell while the other lines are being scanned. When all the lines have been scanned and all the pixel cells addressed, the process is repeated for the next frame and the picture is refreshed.

Figure 4 shows the operation of a TFT liquid crystal cell where G_N is the currently selected gate line and G_{N-1} is the immediately preceding gate line. The TFT gate is connected to the line (or gate)

electrode bus, also known as the gate bus and the drain is connected to the data (or column) bus, also known as the source bus. Storage capacitor C_S is connected between the current gate line G_N and the immediately preceding gate line (G_{N-1}). For this reason, C_S is known as C_S-on-gate. It forms the drain load for the TFT.

The TFT turns fully on when its gate voltage is 20V and turns off when its gate goes to at least -5V. To select the pixel cell, a 20V pulse, V_{SEL}, is applied to the gate. At the same time, data in the form of an analogue positive voltage V_{DAT} is applied to the drain. For peak white, V_{DAT} is 0V while for pitch black V_{DAT} is a maximum of about 8V. With the TFT on, the source and drain are shorted and V_{DAT} is applied across the liquid crystal. The storage capacitor, C_S-on-gate charges up and this charge is sustained even when the TFT is turned off. This is then repeated for the next line and so on.

The main function of C_S is to maintain the voltage across the liquid crystal until the next line select voltage is applied when the picture is refreshed. A large C_S can improve the voltage holding ratio of the pixel cell and improve the contrast and flicker. However, a large C_S results in higher TFT load and lower aperture ratio. In determining the value of the storage capacitance, account must be taken of the stray capacitance between the TFT's gate and source, G_{GS} which is effectively in parallel with C_S.
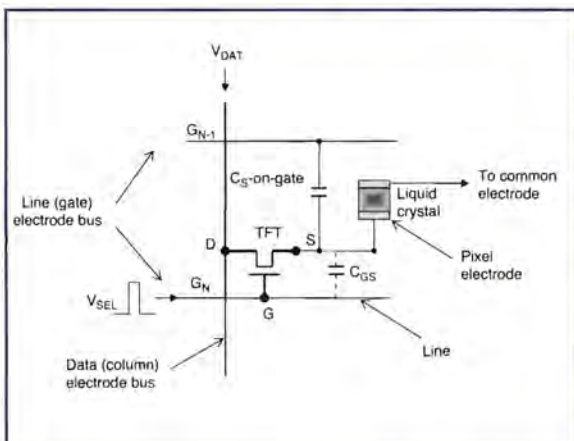

Figure 5: Time rising, Time falling' response of a liquid crystal

RESPONSE TIME

Response time is one of the few areas remaining where the performance of a traditional CRT still holds an advantage over LCDs. CRTs have nearly instantaneous pixel response times, but LCDs tend to be much slower. The result is that the user might see smearing, motion blur or other visual artefacts when there is movement on the screen.

A pixel's response time is the time it takes a pixel to change state. If it is a rise-and-fall response, then it is a measure of the time it takes a pixel to change state from black-to-white-to-black as illustrated in **Figure 5.** More specifically, it represents the pixel ability to change from 10% 'on' to 90% 'on' and then back from 10% 'off' to 90% 'off' again.

Originally, this was the standard way of reporting response times of LCD TVs and computer monitors, and was normally listed as a TrTf (Time rising, Time falling) measurement. Some manufactures started using a gray-to-gray (GtG) measurement for LCD response times which is different from TrTf. There are as yet no standards and manufacturers can state any figure that suits them.

One factor that affects the response time is the viscosity of the liquid crystal material means it takes a finite time to re-orientate its molecules in response to a changed electric field. A second factor is the capacitance of the liquid crystal material is affected by the molecule re-alignment which changes the TFT load and with it the brightness to which the cell ultimately settles.

A good response times starts at around 25ms with some LCD TV manufacturers claiming a response time as fast as 16ms or less. Short response times are required for

Figure 4: TFT cell equivalent circuit

fast moving images such as games. New techniques have been developed to reduce the response time. Such techniques include the use of lower viscosity liquid crystal material. Reducing the cell gap thickness is another technique which results in less of the liquid crystal material to re-orientate, giving a response time as little as 8ms. Thinner cells make production more difficult with lower yields and hence more expensive.

Another technique is to apply a drive signal for a brief duration in order to give the pixel cells a 'jump start' and then reducing it to the required level as illustrated in **Figure 6**. This technique is known as Amplified Impulse provides grey-to-grey transition to be completed up to five times faster than a typical LCD.
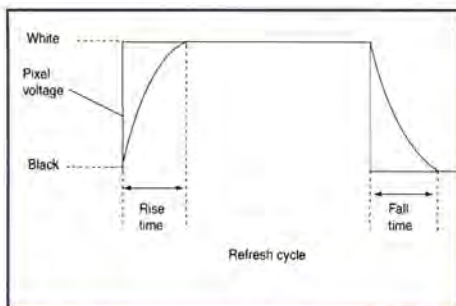

Figure 6: Amplified Impulse LC cell drive technique

For TV images, two techniques have been developed, both attempt to hide the cell transition time. 'Backlight strobing' involves flickering the backlight off momentarily and the 'black frame insertion' introduces a black frame during the LC transition. Backlight strobing also helps improve motion blur caused by the 'sample-and-hold' effect in which an image, when held on the screen for the duration of a frame-time, blurs the retina as the eye tracks the motion from one frame to the next.

By comparison, when an electron beam sweeps the surface of a cathode ray tube, it lights any given part of the screen only for a miniscule fraction of the frame time. It's a bit like comparing film or video footage shot with low and high shutter-speeds. This type of motion blur has come about as manufacturers moved from the traditional resistor type digital-to-analogue converter to

the much more compact sample-and-hold type. Motion blur originating from sample-and-hold in the display can become less of an issue as the frame rate is increased.

POLARITY INVERSION

In liquid crystal cells, it is the magnitude of the applied voltage which determines the amount of light transmission. Such voltage may be DC or AC. Applying a voltage of the same (DC) polarity to a liquid crystal cell would cause electroplating of one electrode resulting in what is known as 'DC stress' causing deterioration in image quality.

To prevent polarisation (and rapid permanent damage) of the liquid crystal material, the polarity of the cell voltage is reversed, a process known as polarity inversion. Polarity inversion may be implemented in three different ways: frame inversion, line (or horizontal) inversion and dot inversion (**Figure 7**). It will be noticed that line inversion incorporates frame inversion as well, since a positive line in one frame becomes negative in the following frame and vice versa.

Unfortunately, it is very difficult to get exactly the same voltage on the cell in both polarities, so the pixel-cell brightness will tend to flicker. This flicker is most noticeable with frame inversion in which the polarity of the whole screen is inverted once every frame resulting in a 25Hz and 30Hz flicker for PAL and NTSC respectively.

Flicker may be reduced by having the polarity of adjacent lines using line inversion thus cancelling out the flicker. Better results may be obtained with dot inversion. In this way the flicker can be made imperceptible for most "natural" images.
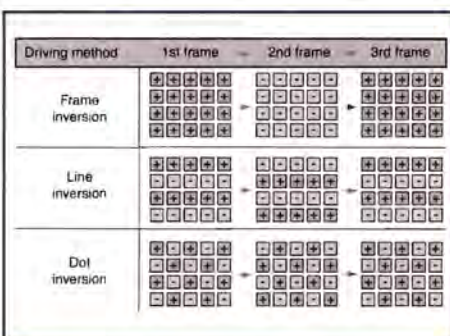

Figure 7: Polarity inversion

Polarity inversion is carried out by inverting both the pixel cell electrode V_P and voltage at the common electrode, V_{COM} frame by frame, line by line or dot by dot. Referring to **Figure 8**, the voltage across the liquid crystal cell V_{LC} is the difference between V_P and V_{COM}: $V_{LC} = V_P - V_{COM}$

When both V_P and V_{COM} are inverted:
$$V'_{LC} = -V_P - (-V_{COM})$$
$$V'_{LC} = -V_P + V_{COM}$$
$$V'_{LC} = -(V_P - V_{COM})$$

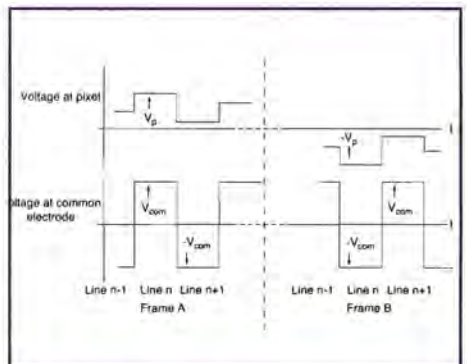**Figure 8:** Voltage across the liquid crystal, $V_{LC} = V_P - V_{COM}$

Figure 9: Line inversion pixel voltages sequence for plain white display

The line inversion sequence for three lines of a plain white screen is shown in **Figure 9**. For line n, the pixel voltage V_P for all the pixels on the line is high to remove the 900 twist of the liquid crystal and remains constant throughout 'line n', since all pixels are at the same luminance level. They are then inverted for the following 'line n+1' and so on. The pixel voltage V_{COM} is also constant over one whole line and inverted for the following line.

For a five-step greyscale display, the pixel values change along the line as shown in **Figure 10**, starting with V_{white} for peak white and $-V_{black}$ for black at the end of the line. This is then inverted and repeated for the next line and so on. V_{COM} on the other hand is constant over one line and inverted over the next. The LC voltage V_{LC}, being the difference between V_P and V_{COM} is that shown in Figure 10 inverted every line. For dot inversion, the pixel voltage and the common electrode voltage will invert on successive dots.

THE BACKLIGHT ASSEMBLY

There are two types of backlight formats: the guided type for screen of 20 inches or less and the direct type for larger screen sizes. Both use cold cathode fluorescent tubes (CCFT) for their low energy consumption and low cost. The guided type is slim and compact but suffers from complicated structure and low light efficiency. On the other hand, the direct type is thick in structure but simpler in structure with high efficiency.

The assembly of the guided type consists of one CCFT tube on either side of the screen with a light guide and a reflector behind the light together with one or more microprisms and one or more diffusers in front of it.

The direct light assembly has more than two lamps as shown in **Figure 11**. The same layers are used as for the guided light type performing the same functions. Direct light diffusers are used to diffuse the light and to avoid seeing the backlight. A transparent ITO (Indium Tin Oxide) sheet connected to ground is used to filter out the noise produced by the lamps.

The CCF tube requires a sine wave with amplitude of few thousand volts and a frequency in the region of 50-0kHz. This is provided by a DC-AC converter, which is in essence an oscillator. While maximum brightness may be obtained by turning the tube fully on, in most applications, there is a need to reduce the lamp's brightness.

There are two basic methods for dimming the CCF tube. The first simply reduces the tube current either directly or indirectly by reducing the voltage applied to it. The second method maintains a constant current but turns the lamp on and off to control its brightness (**Figure 12**). If the inverter is turned on for longer periods than it is off, a brighter light is produced and vice versa.

This technique employs a pulse-width modulated (PWM) waveform to turn the inverter oscillator on and off. A typical tube drive and control signals are shown in **Figure 13**. The frequency of the PWM
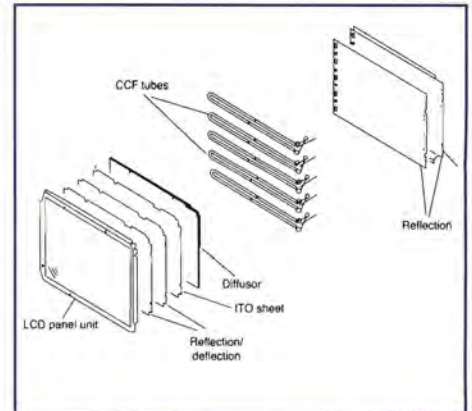
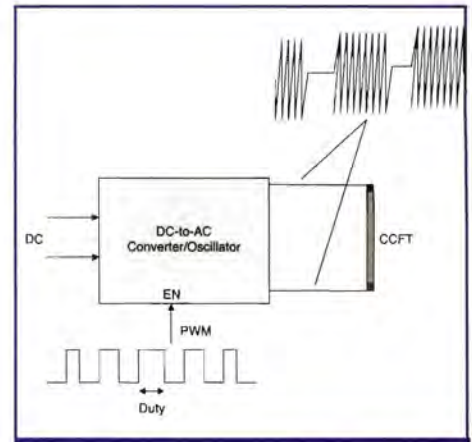

Figure 11: Backlight assembly (direct light type)

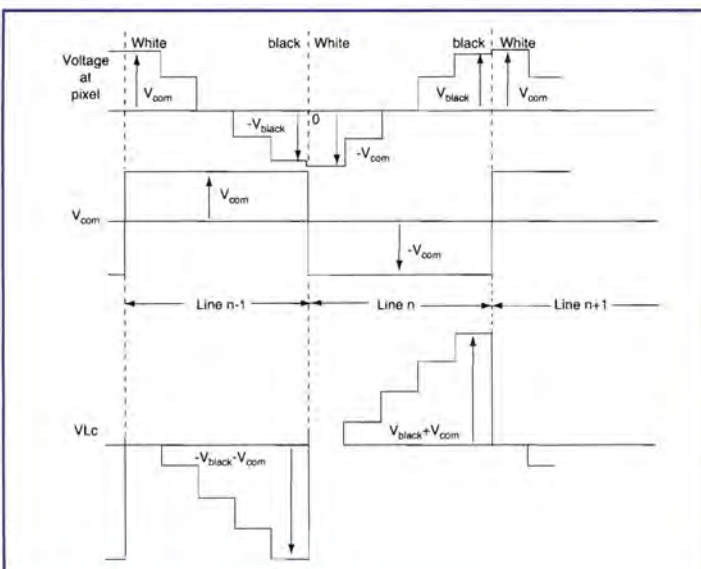*Figure 12:* CCF tube brightness control

*Figure 10:* Line inversion waveform for a five-step greyscale display

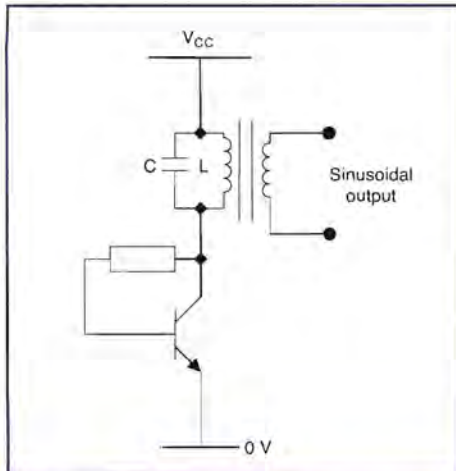*Figure 13:* CCF tube drive and control signals

Figure 14: Tune-collector oscillator as a DC-AC inverter

waveform has to be chosen carefully to avoid interaction with the frame rate. A typical value is 270 and 330Hz and, in the case of PC monitors, it is varied with the frame rate itself to avoid interference with the graphics.

THE DC-AC INVERTER

Essentially, the DC-AC inverter is a tuned collector oscillator (Figure 14). When power is switched on, the transistor conducts feeding energy into inductor L. When the inductor saturates, current ceases and the back e.m.f. forces the current to reverse. Energy in L is now transferred to C. When C is fully charged, charging current ceases causing an opposite back e.m.f. across the inductor

and the capacitor discharges into L with its energy transferring back to L until the inductor saturates and so on. The output across the secondary of the transformer is a pure sine wave.

The basic elements of a DC-AC inverter are illustrated in **Figure 15**. Capacitor C_P is the primary tuned capacitor which resonates with the inductance of the primary winding of transformer T_1. Capacitor C_S is connected in series with the tube to ensure constant current operation. At the frequency of operation, the impedance the tube assembly, together with the ballast capacitor, is a very high making the inverter act as a constant current source. The output is a sine wave with a slight distortion caused by the reactance of the tube.

Because of the very high impedance, measuring devices such as a DVM or a CRO would load the output so much as to render the readings almost meaningless. A current probe should be used to observe the shape of the waveform on an oscilloscope.

Figure 16 shows practical backlight inverter driving two CCF tubes together

with the control chip as used by a Panasonic's 15-inch LCD receiver. The control chip provides the PWM signal to drive two separate inverters, one for each CCF tube. The control chip itself is R10 controlled by a signal from the R10 microprocessor.

For inverter 1, a centre-tapped step-up transformer is used to feed $700V_{rms}$ to drive CCFT1. The tuned circuit is formed by C_1 and the primary of the transformer and
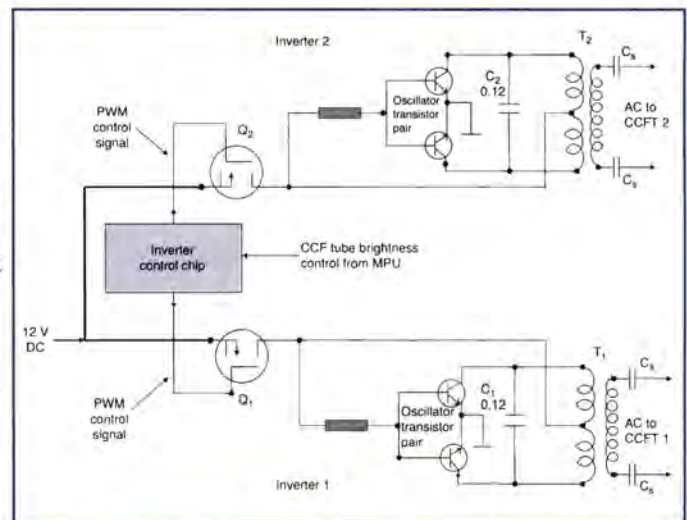

Figure 16: The essential elements of a practical DC-AC inverter

a pair of transistors is used to oscillate back-to-back. The DC power to the oscillator is obtained from switching transistor Q_1. While Q_1 is on, oscillation takes place and the tube lights up. However, when Q_1 is off, the oscillator turns off and with it the tube itself. Switching transistor Q_1 is controlled by a PWM signal from the control chip. The width of the pulse controls the ON/OFF ratio of Q_1 and with it the brightness of the tube.

For inverter 2, Q_2 is the switching transistor and C_2 is the tuning capacitor. The tuning capacitors may be recognised by their non-nominal values, in this case 0.12μF. Capacitors marked C_S are the series capacitors that ensures the lamp presents very high impedance to the inverter. It is normal to include over-voltage/over-current protection, as well as a Panel Enable from the microprocessor controller. ∎


Figure 15: The essential elements of a DC-AC inverter

HANDYSCOPE HS4

Four channels not enough

512 CHANNELS IS THE LIMIT
COMPLETE PACKAGE STARTING AT £ 440

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

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

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

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

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

WWW.TIEPIE.NL

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

HIGH-SPEED SERIAL VIDEO INTERFACE ENABLES RICH GRAPHICS IN MOBILE DEVICES

The emergence of the "personal information device" driven by rapid adoption of mobile phones and their expanding variety of features, capabilities and services is pushing the need for high resolution displays capable of handling increasingly rich video content.

Video messaging, still photos and even desktop-style user interfaces require larger displays in terms of horizontal and vertical pixel densities and the ability to render a higher number of colours per pixel. While yesterday's mobile phones may have touted a 16-bit colour QCIF display, tomorrow's converged devices with not just mobile phone capability but multimedia, computing, navigation and numerous connectivity features such as Bluetooth, Wireless LAN, Near Field Communication or more, may well present all this capability using 18 to 24-bit colour VGA resolutions. Higher resolutions, higher frame refresh rates and overhead in both display and data transmission easily push practical bandwidth requirements into the range above 1Gbps.

Legacy solutions employ either fully parallel solutions or asynchronous serial buses such as SPI to transfer video data. When moving from 16-bit to 24-bit per pixel of video, simply extending the width of the parallel bus is hugely impractical if not only for space reasons, especially if the video signal needs to be routed through a hinge, swivel, pivot or slider type of mechanical display connection.

Serialising the interface to the display seems a natural solution, by solving at the same time the space constraint problem, by reducing the amount of wires, and the EMI problem, by using low-voltage swing differential signalling. However, additional constraints exist:

power consumption is paramount in battery-powered handheld devices, as well as cost. At first glance, placing an additional two components to realise a serial link in an existing design may appear to go counter to the need to minimise power, real estate, component and placement cost. Future solutions, especially for mainstream and high-volume applications, will undoubtedly favour fully integrated serial solutions that will minimise the penalty in power, footprint size and cost.

New consortiums like MIPI (Mobile Industry Processor Interface), have been working for the past few years to standardise such serial links for the cellular handset and similar markets. Standardisation brings the benefit of through industry collaboration, being able to cover a wide range of current and future technical requirements, interoperability between devices and ubiquity of product solutions.

HIGH-SPEED SERIAL INTERFACES

High-speed serial interfaces replace parallel topologies in a wide array of mobile applications today. Many of today's common interconnect standards such as USB and PCI Express are based on serial transmission to achieve speed, physical compactness and link robustness, as do a vast array of implementations less visible to the consumer, such as notebook computer display interconnect, high-speed backplane interconnects and emerging memory bus architectures.

Though different in scope and optimised for best performance in specific environments, high-speed serial interconnects all make use of a few essential elements. Perhaps foremost, several important benefits are all at once

WHERE ARE THE SERIAL INTERFACE SOLUTIONS IN MOBILE DEVICES HEADING TO NEXT, ANALYSES **NIC ROOZEBOOM** FROM NXP SEMICONDUCTORS

achieved by using differential signalling, which provides a substantial reduction in noise emission and allows the signal swing to be substantially reduced, in turn reducing the amount of required signal power.

The ratio at which data is serialised is

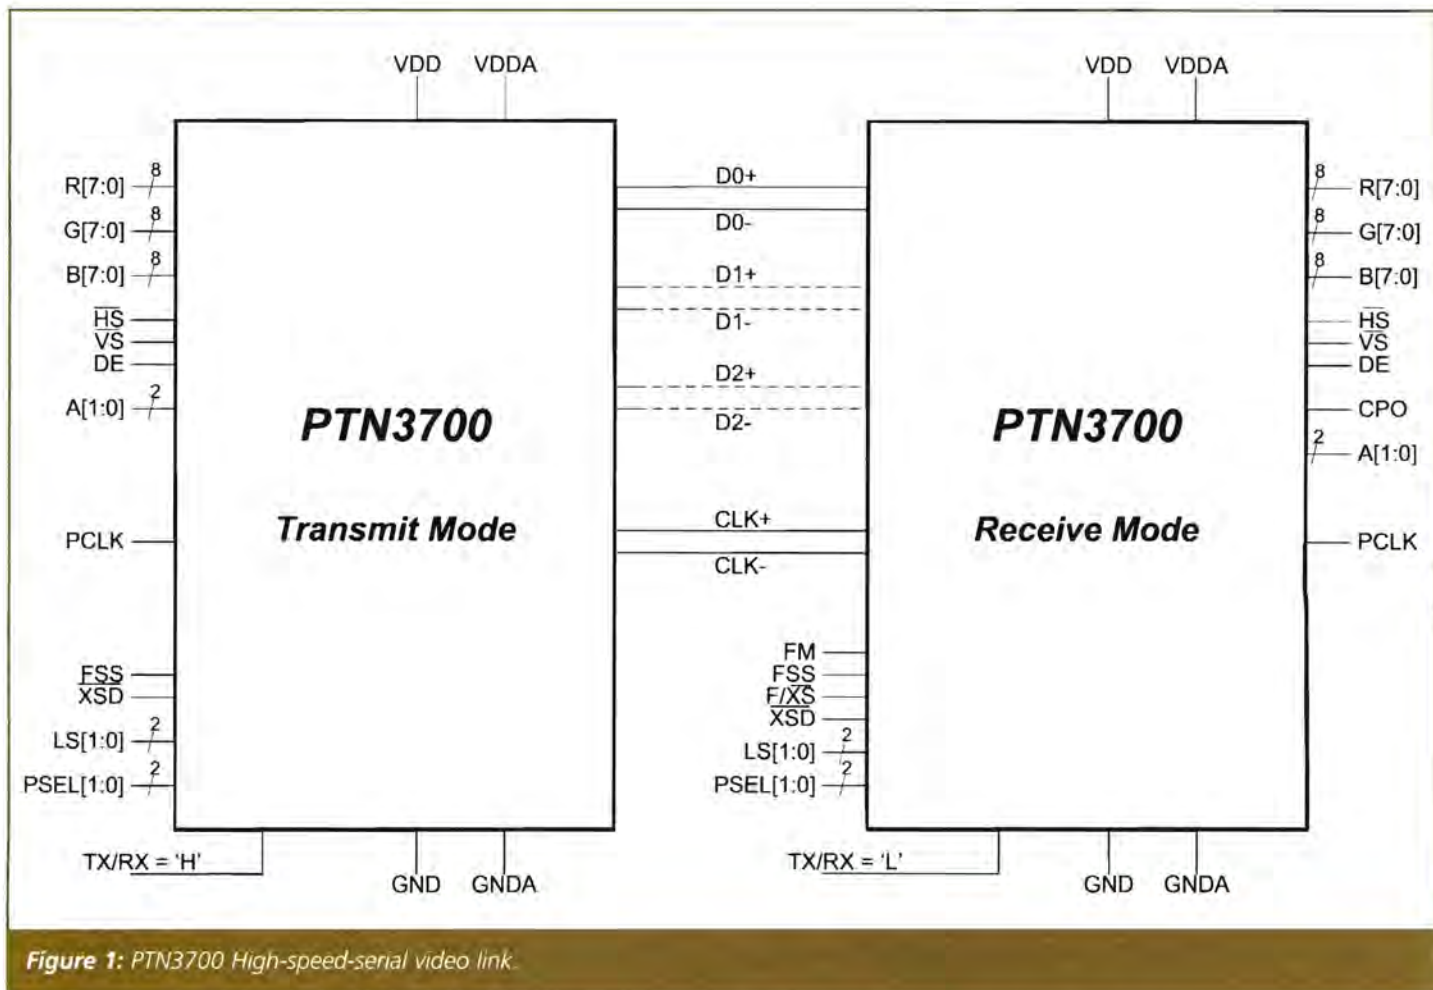*MIPI is playing a significant role in mobile phones*

Figure 1: PTN3700 High-speed-serial video link.

chosen such that per parallel word transmitted, all data bits plus any overhead can be transferred within the parallel clock period. For example, to serially transmit one 24-bit video pixel along with its synchronisation bits without any other overhead, one would need the outgoing serial bit rate to be at least 27 times the incoming pixel clock rate.

The frequency at which video data will be transmitted is bound to two basic quantities: the physical implementation of the video display grid and the display refresh rate. Since the display and display driver need additional time between lines and between the end and start of a frame. Ultimately, the pixel rate can be calculated by multiplying the display refresh rate by the number of horizontal and vertical pixels including overhead.

Should display sizes cause the serial bit rate to exceed what is desirable from an IC implementation or application standpoint, then it is possible simply to

distribute the payload over multiple serial lanes, reducing the absolute signalling rate per lane by that same factor. This makes it possible to scale from lower end display sizes such as QVGA all the way up to high-end display sizes such as XGA, simply by utilising the necessary number of lanes as needed.

Depending on the specific end application for the video serial interface, additional overhead may or may not be needed, at the expense of complexity and efficiency. For traversing relatively short distances, the simplest solution is source-synchronous transmission, where the clock reference for the serial data is transmitted as a separate signal along with the data. When longer distances have to be reached, the difficulty of controlling skew, jitter and other timing issues will increase to the point where it is necessary to use line coding, a process in which the clock reference is embedded into the data stream. This in turn necessitates clock recovery from the data

stream at the receiver end, also increasing complexity and inefficiency, depending on the level of sophistication of the line coding scheme.

Further complexity may be needed when it is necessary to encrypt data for intellectual property protection reasons. But for the purposes of this article, links usually remain short and internal to the mobile device. Moreover, any overhead added to the original data effectively increases the amount of power required to transport each bit, a consideration of paramount importance in battery-powered handheld devices – second to, most likely, space constraints. For these reasons, the arguments are strong to opt for simple source-synchronous transmission as suitable and appropriate for the intended application space.

LANDSCAPE OF SERIAL INTERFACE ICS

The MIPI Alliance has recognised the need to unify today's and tomorrow's

**Figure 2:** *Serial transmission pixel clock referenced*

Figure 3: *Serial transmission pixel clock referenced*

versatile portfolios of MIPI-based solutions.

However, it can be expected that due to the drive for technical innovation and differentiation, vendors will seek solutions that are available now that solve the serial display interface solution in an efficient and cost-effective way. In this environment, "pre-MIPI" solutions from many different IC vendors are seen to be used that, although not providing all desired benefits such as full integration, interchangability, interoperability and multiple sourcing, they do allow setmakers to advance their feature roadmaps faster and enable new functionality.

Solutions are being designed in a stepping stone manner, allowing customers to use currently available products to migrate their product architectures to future MIPI-based solutions. Since redesigns of core chipsets to implement a new interface are costly and take time, add-on solutions mitigate both cost and risk by extending the useful economic life of the existing chipset and allow quicker market introduction of advanced features and increased performance.

In non-mainstream applications that incorporate high-resolution displays, it may not be economically attractive to redesign a processor or system-on-chip to reap the full benefits of a serial display interface. There may also be examples where a core IC is used across a variety of products within a platform, only a subset of which include a fully featured colour display, and for that feature alone it may require a serial interface. In such scenarios, stand-alone MIPI interfaces that can be added on to existing parallel architectures will be able to find useful applications.

FUTURE TRENDS

A serial display solution accomplishes a few important basic goals: serialisation of the interface combined with differential signalling, to achieve low wire count, low EMI and robust data transfer at high bandwidth. Over the next few years, solutions will become available on the market from multiple vendors that are

technical requirements for mobile phones into comprehensive interface standards, so that vendors may benefit from optimal solutions that operate well together, are widely available and can be integrated into core chipsets and peripherals. As of today, fully ratified standards exist for the Display Serial Interface, the Display

Command Set and Camera Serial Interface, for example. For high-volume, cost-sensitive and performance-intensive applications such as mobile phones, vendors of IC solutions (baseband processors, application/video coprocessors, display drivers, interface products) will emerge with increasingly

designed around the MIPI standards for camera and display serial interfaces. This will create an environment where mobile handset makers can design complete solutions based on a modular approach using interoperable components and subsystems.

To enable early introduction of MIPI-enabled designs, bridge solutions may be considered during the time it takes for fully integrated core ICs, display modules and camera modules to become available. Bridge ICs can be useful for a variety of reasons depending on the specific design and product introduction environment. A few typical motivations are discussed here.

First of all, use of new MIPI bridges allow accelerated realisation of MIPI solutions in platforms based on legacy parallel interfaces. In this way, designers can achieve most if not all the technical benefits of high-speed serial solution: wire count reduction, low EMI, low-power transmission enabling high-resolution video to traverse a mechanically challenging interconnect, for example.

Furthermore, attaching a bridge to a design with a legacy interface effectively enables that platform with a standard interface, which in turn can be mated with modules and subsystems built to that standard interface, realising a wider and modular selection of components across multiple vendors and implementations.

As such, standalone bridges may be a valuable part of a flexible ecosystem of standard-based solutions, complementing and enabling platforms revolving around core processors and subsystems that may or may not feature integrated high-speed serial interfaces. This enhances product definition and market introduction effectiveness by adding expandability, flexibility and re-usability of modular solutions.

When implementing high-speed interfaces, especially with signalling rates of over 1Gbps, the design of the basic IP for the physical layer interface as well as its implementation are not quite so trivial. Often in cases where risky analogue parts are being initially implemented in an

Figure 4: MIPI DSI bridge

expensive ASIC, a backup solution is desirable. A separate bridge can, in such cases, mitigate the risk of high-speed mixed-signal IP integration, while providing insurance for timely market introduction of new essential functionality.

Conversely, add-on interfaces provide benefit by prolonging economic lifetime of core chips and thereby preserving significant investment in hardware and software infrastructure. Adoption rates of new interfaces often differ between core ASICs on one hand and peripherals such as displays and cameras on the other. A standalone bridge resolves the discrepancy by being able to re-use an existing host processor with parallel interface for a longer period.

THE BENEFITS

The migration to higher resolution displays and camera sensors in mobile devices is advancing rapidly with the consumer demand for more content-rich media while on the go. This trend necessitates the need for increasingly

higher data bandwidth between host processor and display or camera, even though there's a call for further miniaturisation and lower power consumption. Some of these technical challenges can be surmounted by the use of efficient high-speed serial interfaces.

While integration of serial interfaces such as MIPI DSI or CSI onto host processors and peripheral devices is ideal for power and space considerations, early adoption is greatly helped by standalone devices able to bridge legacy devices to the new standard. Smartly designed standalone bridges, due to their small footprint, low power and easy applicability allow designers to realise most of the technical advantages of high-speed serial transmission, while offering benefits such as extending the useful economic life of existing chipsets and peripherals, allowing market introduction before integrated solutions are available, and providing greater flexibility to add high-speed serial capability across a wide range of platforms, modules and peripherals. ■

> USING SYSTEM MONITOR CAPABILITIES FOUND IN 65NM FPGAS CAN GREATLY INCREASE ENVIRONMENTAL COVERAGE OF FPGA DESIGNS AND OFFER ADVANCED DIAGNOSTICS DURING DEVELOPMENT. ANTHONY COLLINS, STAFF PRODUCT MARKETING ENGINEER AT XILINX GOES INTO THE DETAIL

ENHANCING SYSTEM MANAGEMENT AND DIAGNOSTICS

Consumer expectations run high in the telecommunications industry. When people pick up the telephone, they expect to hear a dial tone and – barring a natural disaster or catastrophic system failure – they always do. This expected level of service is what the industry refers to as high availability. As broadband providers start competing with national and regional telecommunications companies for voice and video services by offering triple play services, consumers expect the same high availability for their digital voice and Internet services as they've come to expect from the local phone company.

High availability results from building redundancy into the system hardware. However, to effectively manage the redundancy, the system must be able to monitor its own operating conditions and, in the event of a failure, switch to backup hardware before customers notice any downtime. Close monitoring of the physical environment allows pre-emptive action to be taken in the event of a failing component. This involves monitoring the physical environment inside the chassis, using various sensors to record such variables as temperature, supply voltages, humidity and cooling performance.

FPGAs are an important building block in a system's high availability infrastructure. Therefore, it requires better monitoring capability of a FPGAs own on-chip environment and its immediate system surroundings.

SYSTEM MONITORING

System monitoring capabilities found in today's 65nm high-performance FPGAs allow designers to easily access information regarding the FPGA on-chip (die) temperature and power supply conditions. This innovative capability also provides access to external sensor information via external analogue input channels. Up to 17 external sensors can be monitored. Access to this information involves little or no design effort, depending on the required functionality. Common functionality like alarms, automatic channel sequencer and data filtering are available within the system monitor blocks, enabling easy development of a solution.

Figure 1 shows the System Monitor

Figure 1: Virtex-5 system monitor

Figure 2: *Access system monitor measurements via the JTAG TAP*

capability in Xilinx Virtex-5 FPGAs. The capability is built around a 10-bit 200kSPS (kilo-sample-per-second) analogue-to-digital converter (ADC). The analogue input range of the ADC is 0V to 1V. At a resolution of 10-bits, the ADC can accurately resolve an input voltage to an accuracy of approximately 1mV.

As seen in Figure 1, both the on-chip sensors and external analogue input channels are connected to the ADC input using analogue multiplexers, which means output voltages from various sensors must be sequentially converted to a digital word by the ADC. These measurement results are written to status registers where they are easily read using the FPGA fabric or externally using the FPGA and PC board JTAG infrastructure.

The system-monitoring control registers can be written/read using the same interfaces. The control registers are used to configure the system-monitoring operation (e.g. select sensor channels for measurement, program alarm limits, set up sensor filtering, etc) The system monitor capability is fully functional shortly after power-up and does not require the FPGA to be configured for correct operation. By default only, the

on-chip sensors are monitored after power-up. However, external analogue inputs can also be enabled. Measurement information can only be accessed via the JTAG Test Access Port (TAP) prior to configuration.

USER ALARMS

One of the useful built-in features of system monitoring capabilities is the ability to generate alarm signals for the on-chip sensors. This allows designers to specify the threshold limits for these alarm signals. A system monitoring capability can be configured to autonomously monitor the sensors and only alert the system when an alarm condition is detected.

It also contains a factory-set alarm condition, called over temperature (OT). If a designer enables the system monitor capability, he or she can request a full chip power-down if a die temperature greater than 125°C is detected. Chip power-up is initiated once the die has cooled to a level specified by the designer. The system monitor continues to operate and monitor the on-chip sensors during chip power down. The OT functionality is disabled by default and must be explicitly enabled by the designer.

CHECKING THE CHECKER

Using the system monitor capability to provide accurate and reliable environmental information requires reliability checks on the measurement data and system monitor operation. The system monitor function offers a number of features to help confirm reliable operation. Built-in auto calibration of the ADC and sensors corrects any drift in the analogue measurement system due to the operating environment. Self-check features also allow the system host to monitor the operation of the system monitor function.

LEVERAGING SYSTEM MONITOR JTAG ACCESS

The system monitor also offers the ability to access the full functionality of the block using the JTAG TAP. By enabling analogue testing and access to analogue information, greater value and efficiencies can be achieved using the system's existing JTAG infrastructure. This access is available before configuration of the FPGA for use as part of a PC board test-scheme, either in production or during normal operation, to facilitate a debug effort.

To support off-chip measurements, e.g. supply voltages and currents on the PC

board, special JTAG commands can be used to enable external analogue inputs prior to FPGA configuration. Even after FPGA configuration, the system monitor function within the FPGA does not require an explicit instantiation in a design, which allows full access to its features for debug work via the JTAG TAP, even at a late stage in the design process. **Figure 2** illustrates a typical diagnostic application where the physical operating environment of the FPGA is monitored during normal operation.

The example in Figure 2 shows the system monitor capability is used to look at IR drop in the power distribution system (PDS) during a period of heavy current demand starting at time t_0. The temperature of the FPGA is also monitored during this high-activity period. Potential issues with the power supply or PC board design can be quickly identified during development. The JTAG access also provides an easy way to

confirm that adequate cooling is in place for a particular design.

A real-time debug and verification system for FPGAs can offer an easy way to access the system monitor capability within the FPGA; however, access can also be incorparted into other JTAG test and programming environments.

SYSTEM INTEGRATION

In addition to convenient access via the JTAG TAP, full access to the system-monitoring function and status registers is also provided via the FPGA fabric. These registers can be configured and read at any time from the fabric. Dual access to the system-monitoring function registered by the JTAG TAP controller and fabric interface is permitted and an arbitration scheme is available to manage possible contention. The contents of these registers can also be defined when the system monitor function is instantiated in a design and initialised

during FPGA configuration. Thus, the system monitor capability can be configured to start up in a user-defined mode of operation post-configuration.

The fabric interface is known as the dynamic reconfiguration port (DRP). The DRP is a BRAM-like parallel, 16-bit synchronous data port. For more advanced applications where greater control over the system monitor capability is required, the DRP allows the function to be easily mapped into the peripheral address space of a hard or soft uP.

Figure 3 illustrates a typical system management application where the soft core processor is running a protocol like intelligent platform management interface (IPMI) and communicating with the system host over management channels like Ethernet or even a simple UART/modem. The system monitor feature also provides an important uP peripheral in the form of a general purpose ADC.

This is the first time analogue peripherals, like those commonly found in microcontrollers, have been integrated into an FPGA. Full control over the ADC operation is supported. The ADC offers a number of sampling modes and can support unipolar, bipolar and full differential analogue input schemes.

IMPORTANT ASPECTS FOR DESIGN

Improving reliability of systems using environmental monitoring is an important aspect of the design concept. By system monitor capability delivers a simplified solution for common on-chip and external environmental monitoring needs. Minimal development and design effort is required to access the functionality. By interfacing the system-monitoring capability to the JTAG TAP controller, the JTAG functionality has been extended into new application areas and has enabled new test capabilities. ■

Figure 3: System monitor (ADC) as a uP peripheral

DOGAN IBRAHIM, LECTURER AT THE NEAR EAST UNIVERSITY IN CYPRUS, DESCRIBES THE DESIGN OF A LOW-COST HANDHELD ECG UNIT BASED ON A POPULAR MICROCONTROLLER AND A GRAPHICS LCD

MICROCONTROLLER BASED HANDHELD ECG UNIT
WITH GRAPHICS LCD

An electrocardiograph (ECG) is a measurement of the electrical activities of the heart muscle as obtained from the skin. The ECG has been used extensively in medicine since its invention in the early 1900s and is currently used to diagnose various heart related diseases, such as tachycardia (fast beating of the heart), backycardia (slow beating of the heart), coronary artery blockage and various other irregular heartbeat patterns. ECG is also used in sports medicine and training for tracking the heartbeat patterns to assist the patient in attaining a desired optimum heart rate.

Figure 1 shows the typical ECG waveform of a normal person. Various parts of the waveform are named with letters such as P, Q, R, S and T. The beginning of the cycle (P-wave) corresponds to the contraction of the atria; the time delay here corresponds to the filling of the ventricles. The ventricles then contract generating part of the waveform known as the 'QRS complex'. Then the depolarisation occurs where the ventricular muscle returns to rest, producing the T-wave.

Since there is no absolute zero reference voltage in the body, it is the difference in potential between two points that is measured. In order to detect the strongest potential difference, the optimum electrode positions are: one on the right shoulder, one on the left shoulder and one on the left leg. This is also known as the Einthoven's Triangle and arises from the work of Willem Einthoven who discovered the principles of the ECG and was the inventor of the first electrocardiograph.

Figure 1: A typical ECG waveform

THE ELECTRONICS

The block diagram of the ECG unit designed is shown in **Figure 2** The unit consists of four functionally separate sections: electrodes, data acquisition unit (DAU), data processing unit (DPU) and display unit (DU). The electrodes are used to sense the electrical potential at pre-determined points of the body. The DAU removes any electrical

noise present in the signal and amplifies the signal several thousand times. The signal is then fed to the DPU which converts the signal into digital form, scales it and sends to the DU where it is displayed on the graphics LCD.

The amplitude of a typical ECG signal is around 1-5mV with a bandwidth of 0.05-100Hz and combined with electrical noise. This signal must be amplified to such a level that it can be plotted or analysed as required and, at the same time, the noise must be removed from the signal by using appropriate filters.

The main sources of error in the measurement of the ECG signal are motion artefacts, 50Hz powerline interference and noise from surrounding electrical equipment (e.g. nearby motors).

The requirements of a typical ECG system can be summarised as follows:

Figure 2: Block diagram of the ECG unit

- Capability to sense signals as low as a few millivolts in the presence of large noise;
- Very high input impedance, > 5MΩ;
- Very low input leakage current < 1μA;
- Noise rejection at 50Hz;
- Noise rejection above 100Hz;
- High common mode rejection ratio;

The ECG signal is so small that it is impossible to separate this signal from noise, unless an instrumentation amplifier is used at the input of the system. Instrumentation amplifiers are high gain amplifiers with very high Common Mode Rejection Ratios (CMRR). An instrumentation amplifier basically consists of a difference amplifier designed to reject common mode signals. Thus, using an instrumentation amplifier we can reject the signals which are common to both inputs of the amplifier, such as the noise signals. CMRR is one of the commonly used parameters which define the efficiency of an instrumentation amplifier. It is defined as:

$$CMRR = \frac{A_d}{A_m}$$

where is the differential gain and is the common-mode gain. Usually, the CMRR is expressed in dB where:

$$CMRR(dB) = 20\,Log\,\frac{A_d}{A_m}$$

The higher the CMRR, the better the performance of the instrumentation amplifier. For example, the AD624 instrumentation amplifier has a minimum CMRR of 140dB, which corresponds to a ratio of 10,000,000. So, if the differential gain is set to be 1000, then from Equation 1 follows:

$$A_m = \frac{A_d}{CMRR} = \frac{1000}{1,000,000} = 0.0001$$

or the common-mode gain is 0.0001. This implies that the amplifier actually does not amplify common-mode signals, but attenuates them by a factor of 1:0001, i.e. stops common-mode signals from reaching

the output, thus eliminating noise at the inputs of the amplifier.

ELECTRODES

The electrodes are an important part of any ECG system as they sense the electrical activity of the heart and pass the received signals to the amplifiers for amplification and processing.

High quality ECG measurement requires use of good electrodes, proper placement of electrodes on the body, good amplifier design and good laboratory and clinical practices. Most commonly used electrode conducting material is made of Ag, Ag-Cl or SiCl. The choice of the electrode material helps reduce the junction potential and, thus, provides a smooth signal with no baseline drift. Ag-Cl type electrodes are commonly used in electrode designs.

Currently, disposable ECG electrodes are in common use which are normally used once and then thrown away. As shown in **Figure 3**, these are small rectangular shaped flexible paper-like materials with one side having conductive adhesive polymer mounted on a backing. A small aluminium foil is left at one end of the electrode for making electrical connections and gel or any kind of adhesive are not required. Usually, plastic clips (see **Figure 4**) with one face containing a conductor (e.g. metal) are used to make external connection to the electrodes. The other end of the clip is connected to a wire lead to send the signal to the ECG unit.

Figure 3:
Disposable
ECG electrode

THE ECG CIRCUIT

The ECG circuit is shown in **Figure 5**. The input part of the circuit is similar to what is proposed by S. Carlson in "Home Is Where ECG Is" in the Scientific American

Figure 4:
Electrode
clips

journal of June 20, 2000. At the input of the circuit an AD624 type instrumentation amplifier is used and the electrodes are connected to the amplifier using current limiting resistors. A set of diodes are used at the inputs to protect the circuit from any over-voltage. The ECG voltage is only a few millivolts and, thus, the diodes are normally non-conducting. Should this voltage become 600mV or higher, the diodes will conduct to protect the system inputs.

The gain of the instrumentation amplifier is set to 1000 by the following connections at the input terminals:
- Connect pins 3,11 and 13
- Connect pins 12 and 16

The amplifier is operated from two 9V batteries, connected to provide ±9V. A passive second order low-pass filter is used at the output of the instrumentation amplifier, designed using two resistors and two capacitors. It was found during the trials that a considerable amount of 50Hz powerline noise was present in the signal. A second order notch filter with a centre frequency of 50Hz and a bandwidth of 6Hz was used to remove this noise.

The filter was designed using the UAF42 type filter chip (OPA2111 data sheet). The required filter components and theoretical filter frequency response were obtained using the filter design software supplied free by Burr-Brown (now TI). **Figure 6** shows the 50Hz notch filter whose components were chosen to be exact to 1% tolerance, which are commercially available. The theoretical frequency and phase responses of the filter are shown in **Figure 7**.

The ECG waveform is bi-directional where the S part of the waveform is negative. But the A/D converter on the PIC microcontrollers is uni-directional, where only the positive analogue voltages in the range of 0 to +5V can be converted to digital.

Figure 5: ECG circuit diagram

In order to convert the ECG signal to uni-directional positive signal, a level shifting circuit was designed using one half of the OPA2111 operational amplifier configured as a difference amplifier, and the negative input is supplied with a fixed voltage using a potentiometer. Thus, by varying the potentiometer we can shift the base of the waveform to make it uni-directional.

Although the gain of the instrumentation amplifier was set to 1000, depending on the subject and the contact potential, this gain may not be enough to bring the ECG voltage to between 0 and +5V. Thus, the other half of the OPA2111 operational amplifier was used to provide further gain, using a potentiometer where the gain can be changed by a factor of about 7. Two diodes are connected to the output of the DAU to protect the inputs of the microcontroller from over-voltage or negative voltage.

The Data Processing Unit is designed using a PIC16F877 microcontroller. The ECG signal is fed to the AN0 analogue channel of the microcontroller. PORTS C, D and E of the microcontroller are used to drive a CRN-12864C type graphics LCD, using a built-in KS0108B controller. The graphics LCD operates from a single +5V supply and consists of 128 pixels horizontal and 64 pixels vertical. The graphics LCD is designed such that the (0, 0) co-ordinates are the top-left point of the display.

THE SOFTWARE

The software used in this project was developed using the PROTON+ compiler,

Figure 6: 50Hz notch filter

developed by the Crownhill Associates Ltd in the UK. The compiler supports most of the PIC range of microcontrollers and takes full advantage of their various features, such as an A/D converter in the 16F87x series, the data memory EEPROM, text and graphics-based LCDs, and so on.

The software developed for the ECG is quite small. The basic operation of the software is described by the following PDL:

```
BEGIN
                Declare variables
                Configure I/O ports
                Configure A/D channel
        DO FOREVER
                Start A/D converter
                Get 128 A/D samples
with 10ms delay
                Call subroutine DISP
        ENDO
END

DISP:
BEGIN
        Wait 1 second
        Clear the screen
        Scale the samples
        Display on graphics LCD
        Return from subroutine
END
```

The program listing is given in **Figure 8**. The following variables are used:

RAMPTR:	a pointer used to index the A/D samples received.
SAMPLES:	an array which stores the A/D samples. This is a word array, i.e. each sample is 2 bytes long (actually 10 bits wide) and up to 128 samples can be stored in this array
x	x co-ordinate of the graphics screen
y:	y co-ordinate of the graphics screen
yy:	temporary word variable
XMAX:	maximum value of variable x
YMAX:	maximum value of variable y

The input-output ports are then configured, where PORT A is configured as the input, and PORT C, PORT D and PORT E are configured as outputs. The A/D converter channel AN0 is configured for 10-bit operation and the pointer RAMPTR is initialised at 0.

The horizontal axis (y axis) of the LCD display is used as the time axis and the vertical axis (x axis) is used to represent the voltage (the amplitude of the ECG waveform).

The main program loop starts with label Loop. The ECG signal is read and converted into digital using the function ADIn and the signal is stored in array SAMPLES, indexed by variable RAMPTR. After receiving a sample the pointer is incremented to point to the next location. This process is repeated until all 128 samples are received and a 10ms delay is used between each sample. Thus, each horizontal pixel corresponds to a sample and the total width of the horizontal axis is 10ms x 128 = 1280ms or 1.28 seconds. In this period, it is expected to see two complete ECG waveforms (each waveform is about 600ms long).

After receiving 128 samples, the program calls to subroutine DISP to plot the ECG waveform. The LCD screen is refreshed at a rate of one second so that it can be seen before cleared. Inside this subroutine, RAMPTR is cleared to 0 so that it points to the beginning of the array. Also, XMAX is defined as 1024 and YMAX is defined as 128. This is because the vertical axis (x axis) is for voltage and the voltage can have 1024 quantisation levels in a 10-bit A/D converter. Also, the horizontal axis consists of 128 pixels where each pixel corresponds to 10ms. Thus, the maximum value along the horizontal axis is 128.

A FOR loop is then used and the ECG signal is displayed using the PLOT instruction. Here, the vertical axis is organised such that the 0 point is at pixel location 56 so that any negative going pulses could also be shown on the display, and the x values are configured to be positive vertically. **Figure 9** shows the x, y

co-ordinates of the LCD screen. This organisation of the x, y co-ordinates was done using the following assignments:

$$x = 56 - 56 * SAMPLES [y] / XMAX$$

Thus, when the ECG sample value is 0, i.e. SAMPLES [ ] = 0, then x = 56 and this is plotted starting from x co-ordinate 56, which is as shown in **Figure 9**.

Similarly, when the ECG sample value is maximum, i.e. SAMPLES [ ] = XMAX, then x = 0 and this is plotted starting from x co-ordinate 0, which is at the top left corner of the screen as shown in **Figure 9**, i.e. the ECG signal is plotted between the x co-ordinates x = 0 and x = 56, where x = 56 corresponds to the new x = 0 point, and x = 0 corresponds to the new x = 56 co-ordinate. Similarly, the horizontal axis (y axis) is scaled by dividing the SAMPLES with YMAX.

The ECG waveform is plotted by using the PLOT instruction where the x co-ordinate is specified first, followed by the y co-ordinate.

Figure 9: x, y co-ordinates of the LCD screen

TEST RESULTS

Testing the ECG unit was not difficult. The electrodes were connected to the subject and it was necessary to adjust the dc shifting of the waveform and the amplification factor by the two potentiometers.

Figures 10a and 10b show the signals obtained at various points in the unit. These waveforms were obtained using the PC-based digital oscilloscope Cleverscope.

Figure 11 shows a typical ECG waveform displayed on the graphics LCD display. Notice that the vertical resolution of the display is not high enough to display the waveform in detail. It is the author's opinion

```
'************************************************************
'* Name    : ECG.BAS                                        *
'* Author  : Dogan Ibrahim                                  *
'* Date    : 31/08/2006                                     *
'* Version : 1.0                                            *
'* Notes   : This is the ECG processing and display program *
'************************************************************
Include "PROTON_G4.INT"
'
' A/D Variables
'
Dim RAMPTR As Word        ' RAM pointer
Dim SAMPLES[128] As Word  ' A/D converter samples
'
' LCD variables
'
Dim x As Float
Dim y As Float
Dim yy As Word
Dim XMAX As Float
Dim YMAX As Float
'
' General variables
'
Dim I As Byte
'
' PORT Directions
'
  TRISA = %11111111        ' All inputs
  TRISC = %00000000        ' All outputs
  TRISD = %00000000        ' All outputs
  TRISE = %00000000        ' All outputs
'
' Set-up A/D converter
'
  ADIN_RES = 10            ' A/D converter 10-bits
  ADIN_TAD = FRC           ' Use internal RC oscillator
  ADIN_STIME = 50          ' 50 us sample time
  ADCON1 = %10000010       ' Set PORT A analog right justified
  DelayMS 1000             ' Wait 1 second for LCD to initialize
  Cls                      ' Clear LCD screen
  RAMPTR = 0               ' Initialize RAM pointer

'=============== START OF MAIN PROGRAM LOOP ===============
Loop:
  SAMPLES[RAMPTR] = ADIn 0   ' Read Channel AN0 (0 to 1023)
  RAMPTR = RAMPTR + 1        ' Increment RAM pointer
  If RAMPTR = 128 Then GoSub DISP  ' Display on LCD
  DelayMS 10                 ' 10 ms delay between samples
  GoTo Loop                  ' Repeat

DISP:
  DelayMS 1000               ' Wait 1 sec to freeze screen
  Cls                        ' Clear LCD screen
  RAMPTR = 0                 ' Re-initialize RAM pointer
  XMAX = 1024                ' Initialize max x value
  YMAX = 128                 ' Initialize the max y value
'
' Plot the samples. There are 128x64 pixels. First scale the screen
'
  For yy = 0 To 127
    x=56*SAMPLES[yy]/XMAX     ' Scale the x values
    x=56-x                    ' Offset within the LCD screen
    y=yy*128/YMAX             ' Scale the y values
    Plot x,y                  ' Plot the points
  Next
  Return                      ' Return from subroutine
  End
```

Figure 8: ECG program listing

Figure 10a: Output of the instrumentation amplifier. The 50Hz noise is clearly visible

• The ECG waveform can be analysed by the microcontroller to detect possible heart problems and the results of this analysis can be displayed at the bottom of the graphics display;

• A serial RS232 type interface can be connected to the unit. It should then be possible to connect the unit to a PC and collect the waveform data. Programs can be developed on the PC to analyse and display this waveform in real time and in large format. Also, the waveform can be analysed to detect possible heart problems and then results displayed on the screen. The waveforms and the results can also be printed from the PC for offline analysis or for storage as hard copy. ∎

securely on the subject for long times;

• Electrodes used seem to be very sensitive to contact potential and this can cause the base level of the signal to drift;

• The LCD resolution was not enough to display a detailed waveform. It is the author's opinion that the minimum resolution should be 128 x 128 pixels;

  The dc level and the amplitude of the displayed signal can be changed using two potentiometers and this

Figure 10b: Output of the ECG unit

that the resolution of the display should be at least 128 x 128 pixels.

RESULTS
The design of a handheld, battery operated, low-cost ECG unit with graphics LCD output has been described. Test results showed that the ECG unit gives good output. Some of the problems encountered during the testing were as follows:

• Although the electrodes used are adhesive it was difficult to keep them

has been useful for getting clear and detailed waveforms. The 50Hz powerline noise was well removed from the waveform using the notch filter.

RECOMMENDATIONS
Although the design of the ECG unit was successful, it can be improved further by the modifications and additions as listed below:

• A buzzer could be provided in the design, which will give sound every-time the QRS-complex is detected;

Figure 11: Waveforms displayed on the graphics display

GOOD NEWS! POLYMER OLED TECHNOLOGY IS ABOUT TO COME TO A STICKY END

By Chris Williams, UKDL

You might think that this headline spells doom and gloom, but far from it. I am pleased to report that cross-disciplinary work in the UK is resulting in polymer OLED (Organic Light Emitting Diode) technology being applied with fantastic success in the medical field to cure skin cancers.

The behind-the-scenes facts are staggering. Some 10% of the UK population, 40% of Americans and a mind-blowing 75% of Australians will suffer from at least one event of skin cancer during their lifetime. Whether these events are due to global warming, changes in lifestyle or other reasons, the facts are incontrovertible – skin cancer is a massive menace and, scarily, the incidence levels are going upwards.

Skin cancers can be removed by operation, or in many cases, they can be treated by photodynamic therapy. Here, a cancerous area of skin is coated with a pharmaceutical cream that contains amino-laevulic acid. This selectively reacts with the cancerous tissue of the tumour and metabolises to a porphyrin, a process that can take three hours or more to complete. The tumour is then exposed to light and the porphyrin undergoes a photochemical reaction that kills the tumour.

Conventional photodynamic therapy uses laser light or high intensity lamps. The patient is required to sit still for a period of several hours, and the intensity of the light

TWO TREATMENTS WITH THE CREAM AND THE OLED BANDAGE HAVE BEEN SHOWN TO TREAT SKIN TUMOURS AS SUCCESSFULLY AS CONVENTIONAL PHOTODYNAMIC THERAPY

can cause some individuals intense pain. So, whilst the process of photodynamic therapy is non-invasive and can be preferable to surgical operation, the process is not necessarily without risk and discomfort. It is also limited to delivery at hospitals or specialist clinics, and the patients are effectively immobilised during the length of the treatment.

Enter Professor Ifor Samuel of St Andrews University in Scotland. A co-inventor of Dendrimer OLED technology, Samuel has a wide range of experience in working with various solution-processable polymer OLED materials in collaboration with Cambridge-based firm CDT.

After chance discussions with Professor James Ferguson, Head of Photobiology at Ninewells Hospital, Dundee, they determined to develop a more user-friendly solution that combined the necessary generation of light to activate the photochemical reaction, but using lower intensity light, packaged in a way that would allow patients to be more mobile.

The result of this early development work is the OLED bandage for curing skin cancer.

In its earliest form that has been used for initial trials to prove its effectiveness, the system comprises a polymer OLED display constructed on a small glass substrate that

Polymer OLEDs are manufactured in solution (see bottle) but can be deposited onto various shapes (see tubes and plates). These samples are photoluminescing as a result of being illuminated by UV light

The polymer OLED device is printed onto a glass substrate which is then incorporated into a "bandage" structure for application to the patient's skin

The polymer OLED bandage in use with its portable power supply

is assembled to a 'bandage' that sticks the assembly to the patients skin, covering the tumour. This bandage also has simple metal foil fins to help dissipate the heat generated as part of the process.

A portable, disposable battery unit powers the bandage assembly. (Regular readers of this column will be aware of the fact that with present-day lighting technologies exhibiting less than 100% efficiency in converting electrical energy to light energy, the balance will be mostly expressed as heat.)

This simple product built from the latest state of the art displays technology operates at much lower luminance levels than the conventional photodynamic therapy units, and patients can be fully ambulatory whilst the treatment is proceeding. So, walk in to the clinic or outpatients department, have the cream applied to the tumour, wait three hours for it to be fully absorbed and for the

metabolising action to have completed, stick on the OLED bandage, turn it on, walk out, and three hours later remove and dispose of the bandage and the battery pack. Two treatments with the cream and the OLED bandage, one month apart, have been shown to treat skin tumours as successfully as conventional photodynamic therapy.

This invention is potentially so radical that the two professors have formed a new company called Lumicure to further develop and manufacture the OLED bandage for mass usage. The device is still not available off the shelf; there must now be full multi-centre clinical trials and the bandage system itself must be developed to be more suited to high volume mass production. One of the key targets will be to develop this technology so that it can be produced on plastic substrates, which will conform more easily to the various curves and shapes of the

human body, and also be more suited for simple disposal, but the expectation is that the first products could become available towards the end of 2009.

What an interesting development – polymer OLEDs will not only be in the displays on a mobile phone, laptop and, eventually, TV screens, but they will enable new markets for Printed Lighting with manufacturers looking to create printed lighting tiles and even electronic wallpaper and, now, polymer OLEDs will come to the rescue in many cases of skin cancer too.

How incredible that a technology can go from being a laboratory curiosity in the late1980s to deliver these wide-ranging societal benefits in such a short time.

Chris Williams is Network Director at the UK Display & Lighting Knowledge Transfer Network (UKDL KTN)

CLEAR AND CONSISTENT

I read John Mayo's review of Ian Sinclair's and John Dunton's "Practical Electronics Handbook" in the August issue of Electronics World magazine with a mixture of amusement and incredulity. Mr. Mayo appears to have been more thorough than Messrs. Sinclair and Dunton. I agree that today, descriptions of USB would be very useful to maintain relevance in computing, as the old printer ports and RS232 are well known and probably going to be phased out soon.

Unfortunately, the use of USB is associated with software drivers and I would go further than Mr. Mayo and say that a book which did describe USB should also have a usable description of suitable software which can be really used by readers, so that a new generation of PC-operated home equipment can be developed.

Regarding the comment about transformers, it is pretty dire to find such an inaccurate picture being painted.

Presumably Mr. Sinclair knows that mains transformers do not have air gaps, in the main; one or two exceptions can be found, but these are not common. If the authors were not aware of this, it is more serious than Mr. Mayo suggests. One would have thought that such a book ought to make it quite clear where air gaps are and are not needed!

My final thoughts on this review are that actually, electronicists are multi-disciplinary, and as electronics spans physics and chemistry, there is no reason why details which apply should not be presented in whatever form is the most appropriate. This is not to disagree with Mr. Mayo, but to agree that if some particular nomenclature is used, it should be at least explained. For example, semiconductor engineering might be thought of as a specialised branch of solid-state physics, but many would still view this as "electronics"; and similarly, battery design could be viewed as a branch of chemistry. I dare say that physicists would claim

semiconductor engineering as "one of theirs" and chemists, battery technology as "one of theirs" (perhaps with a little more justification). But I agree again with Mr. Mayo that terminology should be made clear and books should be consistent.

Well done, Mr. Mayo! I won't be buying this book. But I might write one instead.

John Ellis
UK

MYK AND FRIDGES

I love Myk Dormer's column, particularly his offering on home automation ["Home Automation, a Luddite View", Electronics World, May 2007]. They keep offering an 'Internet fridge' on a local game show [in Australia], and you can see the contestants thinking "what the hell would I want that for?", so they haven't even managed to 'give' one away yet.

Roly Roper
Australia

VOLTAGE-MODE UNIVERSAL BIQUAD FILTER

With the increasing emphasis on the voltage-mode universal biquad filters with single input and multiple outputs or multiple inputs and single output, there is still a need to develop new biquad filters offering new advantages.

In 2004, Horng et al proposed a voltage-mode multifunction biquad with a single input and three outputs which can realise low-pass, band-pass and high-pass filter transfer functions by using two DDCCs, two grounded capacitors and two grounded resistors. On the other hand, in 2005, Erkan et al proposed a voltage-mode universal biquad in *Electronics World* magazine on pp. 48-49, June 2005 issue, with three inputs and a single output by using only one balanced dual input-dual output current conveyor (BDI-DOCC), two floating capacitors and two floating resistors. However, these two configurations can not combine the universal biquad filter with a single input and three outputs or with three inputs and a single output.

Now, a new universal voltage-mode biquad with three inputs and three outputs by using only one BDI-DOCC, two capacitors and three resistors is proposed in **Figure 1**. The proposed circuit can be a universal voltage-mode filter with two inputs and three outputs and, also, it can simultaneously realise voltage-mode notch, high-pass and band-pass filter signals from the three output terminals without any component choice conditions.

On the other hand, it also can act as a universal voltage-mode filter with three inputs and a single output and it can realise five generic voltage-mode filter signals from the same configuration without any component choice conditions. Therefore, the new circuit is obviously more versatile than those with a single input and three outputs or with three inputs and a single output.

By using standard notation, the input-output relationship of a BDI-DOCC is characterised by $V_{X+} = V_{Y1}$, $V_{X-} = V_{Y2}$, $I_{Z+} = I_{X+}$ and $I_{Z-} = -I_{X-}$. Circuit analysis for the new

universal filter shown in Figure 1 yields the following three voltage-mode transfer functions, provided only two voltage input terminals V_1 and V_3 are given by the input voltage signal V_i and V_2 is equal to zero:

$$\frac{V_{o1}}{V_i} = \frac{s^2 C_1 C_2 + G_1 G_2}{s^2 C_1 C_2 + s C_1 G_2 + G_1 G_2}$$

$$\frac{V_{o2}}{V_i} = \frac{s^2 C_1 C_2}{s^2 C_1 C_2 + s C_1 G_2 + G_1 G_2}$$

$$\frac{V_{o3}}{V_i} = \frac{-s C_1 G_2 \dfrac{G_1}{G_3}}{s^2 C_1 C_2 + s C_1 G_2 + G_1 G_2}$$

Thus, the proposed circuit realises a non-inverting notch signal at V_{o1}, a non-inverting high-pass signal at V_{o2} and an inverting band-pass signal at V_{o3}. Note that there aren't any component choice conditions for realising the above three voltage-mode second-order filter transfer functions. Moreover, the use of only one BDI-DOCC and five passive components is simpler configuration than that of Horng et al.

On the other hand, if we give $V_2 = V_i$ in addition to the above two input voltage signals, $V_1 = V_3 = V_i$, the voltage output signal V_{o1} of the proposed circuit shown in Figure 1 is obtained and shown by the following equation:

$$\frac{V_{o1}}{V_i} = \frac{s^2 C_1 C_2 V_1 + s C_1 G_2 V_2 + G_1 G_2 V_3}{s^2 C_1 C_2 + s C_1 G_2 + G_1 G_2}$$

From **Equation 4**, the five standard biquad filter functions could be easily obtained as the same by Erkan et al.

To validate the theoretical prediction of Figure 1, we used H-Spice with TSMC025 process to do the simulation. The supply voltages were $V_{DD} = -V_{SS} = 1.25V$, the biasing voltages were $V_{b1} = -V_{b2} = 0V$. The biasing currents were $I_{B1} = I_{B2} = 50$.

Figure 2 was designed for $f_0 = 1MH_z$ and $Q = 5$ by choosing $R_1 = R_3 = 2$, $R_2 = 10$,

with $C_1 = 15.9pF$ and $C_2 = 79.5pF$. Figure 2 was shown the simulated and theoretical response as low-pass, band-pass, high-pass, notch and all-pass. As it can be seen, there is a close agreement between theory and simulation.

Hua-Pin Chen And Ming-Tzau Lin
Taiwan

Figure 1: *The proposed voltage-mode universal biquad filter*

Figure 2: *Amplitude and frequency responses of Figure 1*

AN EASY WAY TO REDUCE PARASITIC OSCILLATIONS IN FLYBACK CONVERTERS

Figure 1: *Drain voltage composition of the flyback converter*

Flyback operation is well known. The flyback operation is a two-step process. During the ON-time of the switch, energy is taken from the input and stored in the primary winding of the flyback transformer (actually, two coupled inductors). At the secondary side, the catch diode is reverse-biased, thus the load is being supplied by the energy stored in the output bulk capacitor. When the switch turns off, the primary circuit is open and the energy stored in the primary is transferred to the secondary by magnetic coupling. The catch diode is forward-biased, and the stored energy is delivered to the output capacitor and the load.

The main source of EMI from PSUs of this type is the high frequency parasitic resonance (the leakage inductance resonates with drain capacitance – see **Figure 1**).

In experiment results it has been clearly shown that using a slow diode in the circuit of the auxiliary (bias) winding reduces the high frequency parasitic resonance. The power dissipation in the clamp also decreases. However, in addition, a serial connection to V9 a parallel chain "resistor – fast diode" (R5-V10) is also effected. In this case, the chain reduces the high frequency leakage inductance resonance of the drain capacitance through the slow diode in the first time period and then properly works through the fast diode in the following general flyback mode.

A similar effect is achieved by connecting a serial RC circuit (R11, C27) in parallel to the bias winding. Here, the clamp losses and resonance reduce, too. Combined action of these both chains gives fantastic results. This combination works more effectively (see **Figures 2** and **3**).

This design idea worked around the clock more than a year, without a fail. All produced flyback converters are like twins. I checked the operation of these ideas in an earlier designed AC/DC adapter by using STMicroelectronics's VIPer100 IC.

The EMI from the clamper was also reduced and the operation of this highly sensitive receiver is a lot better now. I have increased the sensitivity and yet decreased the unwanted effects.

I have mathematical models of these complex processes, which give good results for a choice of radio-elements for the flayback converters.

THE WORKINGS

See the circuit representation in **Figure 5**.

In an observed time span, the output winding of the transformer is short-circuited on its inner resistor R2 (there is a charge of in the output bulk capacitor C_{out}). Thus the primary inductance is equal to the leakage inductance of the transformer.

The parasitic oscillation is a product of the parallel oscillatory circuit, and is formed by the leakage inductance of transformer and total capacitance – a sum of drain capacitance of the power MOSFET (in our case it is that of the IC) C_{oss} and inner transformer capacitor C_T of the winding I.

The resonance frequency of this oscillatory circuit with damping (ω_b) is:

Figure 2: *Proposed circuit design*

Figure 2: *Circuit representation*

$$\omega_D = \omega_o\sqrt{(1 - r1^2(C_\Sigma / Ls))}$$

ω_o is the resonance frequency of this oscillatory unloaded circuit:

$$\omega_o = 1/(\sqrt{(C_\Sigma / Ls)})$$

where Ls is the leakage inductance of transformer and C_Σ is the cumulative capacitance:

$$C_\Sigma = (C_{oss} + C_T)$$

C_{oss} is the output (drain) capacitance of the MOSFET and C_T is the inner capacitance of the primary winding (winding I). We can decrease the capacitance C_T by using a proper transformer design. That way we can calculate the estimate resonance frequency of the parasitic oscillation.

The intensity of this parasitic oscillation depends on the Q-factor. The Q-factor is a function of characteristic impedance (ρ) of the oscillatory circuit. Unloaded Q is:

$$Q = \rho/r1$$

where $\rho = \sqrt{(Ls/C_\Sigma)}$

A circuit damping in this case is

$$d_o = 1/Q = r1\sqrt{(C_\Sigma /Ls)}.$$

In our case we have an additional auxiliary winding III (it is the bias winding), which is loaded onto some resistance. This resistance is the product of all parts, which are connected to the bias winding at the time of the parasitic oscillation. We must use a value of capacitance C corresponding to the switching frequency of the converter and the estimate resonance frequency of the parasitic

oscillation. This says that the capacitive reactance of C must be the lowest of estimations of the resonance frequency of the parasitic oscillation but higher than the switching frequency of the converter. In this case we assume that the sum resistance of load in auxiliary winding (winding III) is:

$$R_\Sigma = R_{VD}//(R + 1/j\omega_oC)$$

where R_{VD} is the equivalent resistance of the chain V9, V10 and R5 in the starting period of time of the flyback mode (see **Figure 3**).

R is an additional resistor for suppression of the parasitic oscillation (R27 in Figure 3) and C is the capacitor in series to R, $1/j\omega_oC$ it is its reactance on estimate resonance frequency of the parasitic oscillation ω_o.

Equivalent resistance of the bias winding R_Σ is transformed to the primary winding. Value of this new shunting resistance is dependant on the quotient turns of the windings:

$$R_{sh} = R_\Sigma *(n1/n3)$$

Here, n1 it is number of turns of the primary winding and n3 it is number of turns of the bias winding. In that way we have the oscillatory circuit with losses. Its Q-factor (Q_{eq}) will decrease and the circuit damping (d) will increase according to value of this shunting resistance (R_{sh}). In this case the circuit damping is:

$$d = r1\sqrt{(C_\Sigma /Ls)} + (1/Rsh)\sqrt{(Ls/C_\Sigma)} = d_o + (1/Rsh)$$

It is very general at this stage as the values of C_Σ and R_{sh} are nonlinear. They are functions vs time (C_{oss} and R_{VD}) and voltage (C_{oss}). But it is a suitable means to have some result.

Another thing, however, is that the circuit damping will not increase up to infinity in case $R_{sh} = 0$, because here it will work as an active resistance of the bias winding (resistor R3 in **Figure 5**) and the process will take another course. But, in this instance I don't think that when R_{sh} is close to 0Ω is a real situation. We need to remember about the rate of rise of the MOSFET's drain voltage, which also influences the parasitic oscillation.

This method gives not only reduced parasitic oscillation but also a decreased EMI in the flyback converter. We have lower power dissipation in the clamp circuit, because this design idea's starting point is the second step of the flyback conversion. We have decreased stray coupling between the MOSFET and the PWM controller, especially if we are using an integrated microcircuit for our design. I had this bonus when I used TOP250Y in the high power flyback converter. This method will not increase a MOSFET drain voltage, but it is decreased.

Examples
Graphic examples are shown in Figure 3. They are photos of the circuit design of Figure 2.
Vladimir Rentyuk
Ukraine

Parasitic oscillation when using the usual circuit with a fast diode on the bias winding side

Parasitic oscillation when using a slow diode circuit together with the additional RC-chain on the bias winding side

Figure 3: *Graphic examples of my idea*

This series of Tips 'n Tricks addresses the challenges with a collection of power supply building blocks, digital level translation blocks and even analogue translation blocks. Throughout the series, multiple options are presented for each of the transitions, spanning the range from all-in-one interface devices to low-cost discrete solutions. In short, all the blocks a designer is likely to need for handling the 3.3V supply voltage, whether the driving force is complexity, cost or size.

NOTE: The tips 'n' tricks presented here assume a 3.3V supply. However, the techniques work equally well for other supply voltages with the appropriate modifications.

TIP 1: 5V → 3.3V DIRECT CONNECT

5V outputs have a typical VOH of 4.7 volts and a VOL of 0.4 volts, whilst a 3.3V LVCMOS input will have a typical VIH of 0.7 x VDD and a VIL of 0.2 x VDD.

When the 5V output is driving low, there are no problems because the 0.4 volt output is less than in the input threshold of 0.8 volts. When the 5V output is high, the VOH of 4.7 volts is greater than 2.1 volt VIH, therefore, we can directly connect the two pins with no conflicts if the 3.3V CMOS input is 5 volt tolerant.

If the 3.3V CMOS input is not 5 volt tolerant, then there will be an issue because the maximum volt specification of the input will be exceeded. See other tips for possible solutions.

Figure 1: 5V tolerant input

TIP 2: 5V → 3.3V WITH DIODE CLAMP

Many manufacturers protect their I/O pins from exceeding the maximum allowable voltage specification by using clamping diodes. These clamping diodes keep the pin from going more than a diode drop below Vss and a diode drop above VDD. To use the clamping diode to protect the input, you still need to look at the current through the clamping diode.

The current through the clamp diodes should be kept small (in the microamp range). If the current through the clamping diodes gets too large, then you risk the part latching up. Since the source resistance of a 5V output is typically around 10 ohms, an additional series resistor is still needed to limit the current through the clamping diode as shown **Figure 2**. The consequence of using the series resistor is it will reduce the speed at which we can switch the input because the RC time constant formed the capacitance of the pin (CL).

If the clamping diodes are not present, a single external diode can be added to the circuit as shown in **Figure 3**.

Figure 2: Clamping diodes on the input

Figure 3: Without clamping diodes

TIP 3: 5V → 3.3V ACTIVE CLAMP

One problem with using a diode clamp is that it injects current onto the 3.3V power supply. In designs with a high current 5V outputs and lightly loaded 3.3V power supply rails, this injected current can float the 3.3V supply voltage above 3.3V. To prevent this problem, a transistor can be substituted which routes the excess output drive current to ground instead of the 3.3V supply. **Figure 4** shows the resulting circuit.

The base-emitter junction of Q1 performs the same function as the diode in a diode clamp circuit. The difference is that only a small percentage of the emitter current flows out of the base of the transistor to the 3.3V rail, the bulk of the current is routed to the collector where it passes harmlessly to ground.

The ratio of base current to collector current is dictated by the current gain of the transistor, typically 10-400, depending upon which transistor is used.

Figure 4: *Transistor clamp*

TIP 4: 5V → 3.3V RESISTOR DIVIDER

A simple resistor divider can be used to reduce the output of a 5V device to levels appropriate for a 3.3V device input. An equivalent circuit of this interface is shown in **Figure 5**.

Typically, the source resistance, RS, is very small (less than 10 ohms) so its effect on R1 will be negligible, provided that R1 is chosen to be much larger than RS. At the receive end, the load resistance, RL, is very large (greater than 500kohms) so its effect on R2 will be negligible provided that R2 is chosen to be much less than RL.

There is a trade-off between power dissipation and transition times. To keep the power requirements of the interface circuit at a minimum, the series resistance of R1 and R2 should be as large as possible. However, the load capacitance, which is the combination of the stray capacitance, CS, and the 3.3V device input capacitance, CL, can adversely affect the rise and fall times of the input signal. Rise and fall times can be unacceptably long if R1 and R2 are too large.

Neglecting the affects of RS and RL, the formula for determining the values for R1 and R2 is given by **Equation 1**.

Figure 5: *Resistive interface equivalent circuit*

- Stray capacitance = 30pF
- Load capacitance = 5pF
- Maximum rise time from 0.3V to 3V ≤ 1µS
- Applied source voltage Vs = 5V

The calculation to determine the maximum resistances is shown in **Equation 3**.

$$\frac{V_S}{R1 + R2} = \frac{V_L}{R2} \quad \text{; General relationship}$$

$$R1 = \frac{(V_S - V_L) \cdot R2}{V_L} \quad \text{; Solving for } R_1$$

$$R1 = 0.515 \cdot R2 \quad \text{; Substituting voltages}$$

Equation 1: *Divider values*

The formula for determining the rise and fall times is given in **Equation 2**. For circuit analysis, the Thevenin equivalent is used to determine the applied voltage, VA, and the series resistance, R. The Thevenin equivalent is defined as the open circuit voltage divided by the short circuit current. The Thevenin equivalent, R, is determined to be 0.66*R1 and the Thevenin equivalent, VA, is determined to be 0.66*VS, according to the limitations imposed by Equation 2.

$$t = -\left[R \cdot C \cdot \ln\left(\frac{V_F - V_A}{V_I - V_A}\right) \right]$$

Equation 2: *Rise/fall time*

Where:
t = Rise or Fall time
R = 0.66*R_1
C = Cs+CL
V_I = Initial voltage on C (VL)
V_F = Final voltage on C (VL)
V_A = Applied voltage (0.66*Vs)

Solve Equation 12-2 for R:

$$R = -\left[\frac{t}{C \cdot \ln\left(\frac{V_F - V_A}{V_I - V_A}\right)} \right]$$

Equation 3: *Example calculation*

Substitute values:

$$R = -\left[\frac{10 \cdot 10^{-7}}{35 \cdot 10^{-12} \cdot \ln\left(\frac{3 - (0.66 \cdot 5)}{0.3 - (0.66 \cdot 5)}\right)} \right]$$

Thevenin equivalent maximum R:

$$R = 12408$$

Solve for maximum R_1 and R_2:

$$R1 = 0.66 \cdot R \qquad R2 = \frac{R1}{0.515}$$

$$R1 = 8190 \qquad R2 = 15902$$

TIP: OPTIMISING RESONANT-RESET FORWARD CONVERTERS WITH SYNCHRONOUS RECTIFICATION

By David Burgoon, Linear Technology

One way to a small footprint, high-efficiency DC/DC converter for moderate power output is via a resonant-reset forward converter with synchronous rectification (often used in isolated industrial and telecom applications). The resonant capacitor facilitates zero-voltage turn-off of the primary switch. High efficiency comes from careful choice of the resonant capacitor value and use of fast synchronous rectifier drive circuitry. The converter shown in **Figure 8**

Figure 6: *Typical drain voltage waveform of primary switch with 0.5nF resonant capacitor at 72V DC input. Note the peak voltage 230V and total reset and return time of 2.0us*

Figure 7: *Drain voltage of primary switch with 2.2nF resonant capacitor at 72V DC input. Note the peak voltage is down to 180V and total reset and return time of 1.6us*

delivers 3.3V at 30A from a 48V source with over 92% efficiency.

The resonant capacitor must be small enough to permit the transformer to reset during the available reset time, which is the time that the transformer voltage is not being controlled by the primary switch or the active forward rectifier. This means that the energy in the primary and leakage inductances must be dumped into the effective capacitance across the transformer primary, predominantly the resonant capacitor. For a typical resonant reset waveform, this requires slightly more than a quarter of the period of the resonant frequency:

$$f_{res} = 1 / [\, 2 * pi * sqrt(L_{pri}C_{res})\,]$$

as the inductor energy is depleted when the capacitor reaches its peak voltage. Allowing approximately half of the period permits the primary switch drain voltage to return to the input voltage, resulting in lower turn-on switching loss.

Another important consideration in selecting the resonant capacitor value is that it determines the peak voltage on the primary and secondary forward switches. Using the capacitor value to set the peak switch voltage is often overlooked as an opportunity to use a lower voltage MOSFET(s) with lower switching and/or conduction losses. Unfortunately, increasing the capacitance beyond the reset requirement carries the penalty of additional reset and capacitor return time requirements, which, if not available, can reduce efficiency or require a reduction in switching frequency.

Use a synchronous rectifier drive circuit with appropriate timing, such as the LTC3900, to reduce the reset and return time considerably. This is a result of turning the forward rectifier off at or near the peak of the leakage inductance resonant ring, thereby starting the primary inductance reset at the peak voltage, rather than zero. **Figure 6** shows a typical resonant reset waveform; note the peak voltage of 230V exceeds the 200V MOSFET's rating, and the total reset and return time is 2.0µs. **Figure 7** shows the waveform of the optimised reset circuit. The peak voltage is now down to 180V and the total reset and return time is down to 1.6µs. Input power measurements at 48V in showed a savings of 2.4W.

Figure 8: *Test circuit*

+Vin

36 - 72 Vdc

-Vin

1.0uF
100V

VIN

T1
23.4x20.1x9.4mm planar

2

7

4
3

5

10

11

Si7450DP

2.2nF
630V

10V

L2
1.0mH

BAS21

1

BAS21

6

68uF
20V

0.03
1W

VX

2xSi7336DP

2xSi7336DP

845

1.00K

562

VX

B0540W

B0540W

VB

+VOUT

L1
0.85uH

+Vout

2x100uF
6.3V

3.3Vdc at 30A

-Vout

-VOUT

15
1/4W

26.1K

D1
10V

0.47uF
25V

-VOUT

VIN

30K

383K

10V

150
1/4W

66.5K

1uF

100pF

Vref

A

MMBT2369

50k

8.66K

CT

Vref

820

120K

B

120pF

VCC

UVLO

GND SS

CT

SPRG

RLEB

FB

DPRG

VREF

CT

220pF

68nF

10K 33K

DRVA

CS

DRVB

SDRB

SDRA

COMP

LTC3723EGN-1

6

10

4

2

3

11

150K

Vref

0.47uF

4.7nF

Fsw=275k

220pF

T2

1

4

8

5

560

560

8

FG

SYNC

GND

CS+

CS-

CG

VCC

TIMER

LTC3900ES8

U1

5

1

2

3

4

8

7

6

1nF

Vref

820

MOC207

6

5

2

10nF

330

4.7nF

VB

VIN

COMP

OPTO
GND

FB

OC

LT4430ES6

UN1

+VOUT

330

27.4K

1K

470pF

6.04K

47nF

VIN

2.2nF
250V

-VOUT

2.2nF,630V, TDK C3216JB2J222K (1206)
1uF,100V TDK C3225X7R2A105M (1210)
100uF, 6.3V TDK C3225X5R0J107M (1210)
D1 MMBZ5240B
L1 Pulse PA1294.910
L2 Coilcraft DO1608C-105
T1 Pulse PA0810.007 (7:6:6:1:1T)
T2 Pulse PA0184 (1:1T)

100W 36 to 72Vin to 3.3V/30A Forward Converter

INTERFACING PIC MICROCONTROLLERS

Embedded Design by Interactive Simulation

MARTIN BATES
NEWNES

This is Martin Bates's second book after "*PIC Microcontrollers, an Introduction to Microelectronics*" which dealt with the 16F84 microcontroller.

The new book introduces some more modern and capable parts, particularly the 16F877, together with a broad coverage of typical peripherals and interfaces which enable the device to be applied in the real world.

As an aid to learning, Bates makes use of the Proteus VSM graphical simulation tool from Labcenter Electronics. There is no CD-ROM included with the book, but simulation work files for all examples are available from the companion website at **www.picmicros.org.uk.**

introduction to the PIC16F877 microcontroller. The architecture is covered in sufficient detail, as are techniques of coding. Good coding practice is encouraged, including clear code layout and commenting, and the use of programming aids such as flowcharts and pseudocode. The focus is on assembly language processing, but the alternative of C language coding is also introduced in context. Microchip's own MPLAB is used for code entry and compilation, and there is a good grounding in important topics for the student, such as number systems and parameter types.

> THE BOOK IS WRITTEN BY AN EDUCATOR, WITH THE CLEAR INTENT OF HELPING COLLEGE AND UNIVERSITY STUDENTS GET TO GRIPS WITH MICROCONTROLLERS IN GENERAL AND THE PIC16 PARTS IN PARTICULAR

A free demo version of the Proteus VSM simulator can be downloaded from the Labcenter website. This enables the reader to run the simulations illustrated in the text. To be able to edit and alter the simulations, or enter your own, it is necessary to register the program. Labcenter offers a special edition to readers that will simulate the 16F877 microcontroller for £49. The next option, at a cost of about £150, enables simulations of circuits, peripherals and PIC code for a small selection of PIC16F series parts. To simulate the full 16F series costs about £300, and to simulate all supported PIC parts, about £1,000.

The book begins with a very practical

On the hardware side, the addition of a variety of interfaces and peripherals leads to some useful real-world simulations, for example a keypad and display configured to create a simple pocket calculator. Interfaces including RS232, I2C and SPI are covered, as are basic analogue, motor control and amplifier circuits that might be used in conjunction with a microcontroller.

The concluding examples are the realisation of two systems that are representative of the kind of complexity that one would expect of real world systems. A collection of sensor circuits are introduced, with associated coding, that together could comprise a useful weather monitoring station. A versatile microprocessor system

board, including many of the elements introduced earlier, such as a keypad and numeric display, external memory, and analogue and digital interfaces, completes the course.

The book is written by an educator, with the clear intent of helping college and university students get to grips with microcontrollers in general and the PIC16 parts in particular. The breadth of coverage is impressive, and the depth is appropriate for this audience. The use of graphical simulation, not just of the microcontroller, but of whole systems, animates the learning process. Accomplished programmers could debate the relative merits of graphical system simulation against a microcontroller-focused simulation such as MPSIM, but for a student it does help to bring the code to life.

This is not a book for the seasoned programmer looking to extend his knowledge, but as a course for the aspiring embedded programmer, this book draws into a single resource most of the topics required to create a successful design. The addition of a graphical simulation tool that allows the student to model an entire breadboard on screen is certainly a valuable learning aid. In the early stages of learning it is very easy to become discouraged when a carefully assembled breadboard stubbornly refuses to co-operate. The simulation is much more quickly put together and gives the possibility of seeing inside the microcontroller.

Alex Barrett

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

Q What is China RoHS and when did it come to force?

A China RoHS, or to give its official name – *Measures for the Administration of the Control of Pollution (caused) by Electronic Information Products (EIPs)* – came into force for products manufactured from 1st March 2007.

The legislation applies to products sold in China but excludes products imported into the country for re-export or manufacturing of products for export. It also excludes components imported to be used for manufacturing and future sale in China.

Q Which products are in scope of China-RoHS?

A The scope of China RoHS is much broader than EU RoHS and includes all "Electronic Information Products" (EIPs). These include radar equipment, IT, telecom, production equipment used for making EIPs, some types of test instruments, medical devices, electronic components such as resistors and ICs, batteries, PCBs, materials and certain household appliances. The Chinese government has published guidance which lists over 1800 different products. China RoHS requires all EIPs sold in China to be labelled.

Q What substance restrictions apply?

A At present there are no substance restrictions. However, there will be restrictions for certain specified products that will be listed in a soon to be published catalogue. It is likely that the six EU RoHS substances – lead, cadmium, mercury, hexavalent chromium, PBB and PBDE (not Deca-BDE) – will be restricted although the legislation allows for other, additional substances to be included at the discretion of State.

The markings that are required on all EIPs sold in China differ depending on whether or not they contain any of these substances above the maximum permitted levels.

Q When will the catalogue be published and what is likely to be included?

A No date has been announced but it is likely that this will be towards the end of 2007 or early 2008. The catalogue will specify the types of product that have substance restrictions. It will define which substances are restricted, a date from which restrictions apply and any exemptions that will be permitted. The catalogue will be reviewed annually in order to incorporate new substances of product types that fall within the scope of China RoHS. It is not yet known which products will be in first iteration of the catalogue.

Q My products are EU RoHS compliant, so will they comply with China RoHS?

A To be compliant with China RoHS all EIPs must be marked. At present no substance restrictions apply but if RoHS substances are present this must be indicated by the markings now. A key difference between EU and China RoHS is that the China RoHS marking requirements do not have exemptions; the substance is either present or not and, so, if a product is EU RoHS compliant by exemption, it will still need to be marked if RoHS substances are present above defined China RoHS maximum concentration values (MCVs).

Q How will China RoHS affect the design of products?

A Initially, it will have no effect as it is a marking requirement only. However, the design of electrical equipment will be affected in the future by China RoHS in several ways. Products that are listed in the catalogue will need to be designed and built without the specified restricted substances. This will mean that lead-free components and solders will be used. This can affect design in several ways but the two main ones are: (i) when lead-free versions of a component are not available, an alternative circuit design may be needed, and (ii) some heat sensitive components cannot be used as they will be damaged by the higher lead-free soldering process and, so, an alternative design will be required.

The China RoHS legislation states that the Chinese government will introduce compulsory design standards for EIPs. No drafts are available yet but this is believed to be a measure that could adopt some of the principals of the EU Eco-design approach. Following good eco-design practices now will pre-empt any future measures that will be introduced. The Chinese government is also expected to introduce compulsory standards to define the recyclability of products. No drafts are available yet but it is expected that toxic and hazardous materials should be avoided and only materials that can be recycled should be used. This is ahead of China publishing its own WEEE legislation.

Please email your questions to:
svetlana.josifovska@stjohnpatrick.com
marking them as RoHS or WEEE.

DaVinci DVEVM Now With GHS IDE

Each Texas Instruments (TI) Digital Video Evaluation Module (DVEVM) is now shipping with Green Hills Software evaluation version of its MULTI integrated development environment (IDE) for the DaVinci technology. Developed in partnership with TI, the Green Hills Software MULTI IDE provides a comprehensive suite of development tools and operating system-aware debugging for TI's TMS320DM644x systems-on-chip (SoC) processors.

MULTI for DaVinci technology integrates an ARM compiler from Green Hills along with TI's Code Composer Studio compiler for the TMS320C64x+ DSP core. This means that a single instance of MULTI can compile and debug for both the DM644x ARM9 core and TI's C64x+ DSP core.

MULTI supports advanced kernel and OS-aware debugging for Monta Vista's Linux as well as for Green Hills Software's INTEGRITY and velOSity real-time operating systems. For TI's C64x+ DSP core, MULTI is integrated with TI's DSP/BIOS real-time kernel, including support for full DSP/BIOS kernel awareness.

TI's DaVinci technology is a complete offering of video-optimised processors, software, development tools and support ecosystem, tailored for a simplified development of efficient digital video and audio. compliant.
www.thedavincieffect.com.
www.ghs.com.

Kiosk Printing Heads in a New Direction Amplifier

The applications for a kiosk which prints cards are many and varied – cashless payment/vending systems being another ideal use for a kiosk placed in an office, gym, leisure centre or staff canteen. A unique, high quality kiosk card printer has been developed by DED specifically for these markets, capable of printing one off, individually-designed cards or several of the same design if required.

Based on the highly successful P330i plastic card printer from Zebra, the D330i has been adapted to hold up to 400 cards (0.76mm), allowing for minimal maintenance and supervision. It features a standard USB interface and Ethernet option and is capable of printing cards of between 0.25mm and 1.00mm thick. Print speed for a full colour, edge-to-edge card is less than 25 seconds and for monochrome card is just four seconds.

The D330i also features an adjustable card low sensor which can be used to monitor the level of remaining cards in the hopper. A range of card encoding options are available including a three-track magnetic stripe encoder, contact smart card station and a contactless (proximity) card encoder.
www.ded.co.uk

Customised Microwave 'Kicker' Amplifier

The new TMD PTT4200 a high-speed pulse 'kicker' amplifier from TMD Technologies is a critical part of an ultra fast feedback system being developed by the Oxford University and STFC team for steering the nano-sized electron and positron beams into collision at the proposed International Linear Collider (ILC).

A key part of the feedback system, the PTT4200 amplifier from TMD drives a stripline kicker to provide the correcting kick to the particle beams in the linear accelerator.

The amplifier features a very high power differential output stage capable of providing +/- 70A peak output currents to the kicker with a delay, including rise-time and kicker inductance, of only 35ns. To do this the amplifier employs multi-parallel output channels, with the input stage isolated to prevent ground loops.

Amplifier power consumption is minimised by de-powering the circuitry outside the pulse period – achieved by using a separate trigger input to enable the amplifier before each pulse. The very compact design greatly simplifies installation in the confined space around the beam tunnel and permits the shortest possible lead distance to the kicker.
www.tmd.co.uk

Programmable Eddy Current Probe Driver

Monitran has launched the MTN/DSPD-40, an innovative and extremely versatile eddy current probe driver.

Unlike traditional eddy current probe drivers, which are typically purely analogue and which have to be manually 'tuned' to the application, the MTN/DSPD-40 is fully programmable. Internal memory contains two calibration tables, used to tune the sensor for use with specific materials, for example, carbon steel, stainless steel and aluminium. Users can also load their own tables via the driver's RS485 serial port, or supply Monitran with samples of the material to be tested and its engineers will create the tables.

The driver has a multimode power supply and multimode analogue output, set by the user. Data collection and parameter selection can be through any standard terminal program and the driver has selectable filtering.

It weighs 160g and is packaged in a DIN rail mounted plastic enclosure (measuring 80 x 60 x 25mm) with flameproof and waterproof versions available. Ideal applications for eddy current probes include industrial, aerospace, automotive and scientific research.
www.monitran.com

Active Probes For Digital Oscilloscope Range

Yokogawa Europe has introduced two new active probes to complement the recently introduced DL9040L and DL9140L digital oscilloscopes.

The two new probes are the PBA1000, with a probe frequency bandwidth of 1GHz, and the PBA1500, with a probe frequency bandwidth of 1.5GHz. The PBA1000 combines with the 500MHz DL9040/L oscilloscope to achieve a 500MHz bandwidth measurement system, while the PBA1500 combines with the 1GHz DL9140/L oscilloscope to achieve a 1GHz bandwidth measurement system.

"We have identified the bandwidth sector from 500MHz sector to 1GHz as a major growth area in the oscilloscope market," said Clive Davis, Product Marketing Manager for oscilloscopes with Yokogawa Europe. "The popularity of our DL9040/L and DL9140/L oscilloscopes in this market sector is proof of this, and this market penetration should be boosted still further by these high-performance and cost-effective probes."
www.yokogawa.com

Motor-Driver IC With Hall-Element Commutation For Vibration Motors

The new A1442 from Allegro MicroSystems Europe is a full-bridge motor-driver integrated circuit, optimised for driving vibration motors in applications such as mobile telephones, pagers, electronic toothbrushes, hand-held video game controllers and low-power fan motors.

The new device is designed to drive low-voltage bipolar brushless DC motors, with commutation of the motor being achieved by the use of a single Hall-element sensor which detects the rotational position of an alternating-pole ring magnet.

A high-density CMOS semiconductor process allows the integration of all the necessary electronics, including the Hall-element sensor, the motor-control circuitry and the full output bridge. Low-voltage design techniques have been used to achieve device functionality down to 1.8V drive voltage.

The A1442 is a fully integrated single-chip solution. A soft-switching algorithm is used to reduce audible switching noise and EMI interference. A micropower 'sleep' mode can be enabled by an external signal to reduce current consumption for battery management in portable electronic devices.
www.allegromicro.com

Beta Layout Ltd

www.pcb-pool.com
The Best Value Online PCB Prototyping
Service Available:
Instant on line quotations & ordering (no pre
registration).
We offer the following:
• No minimum quantity. • No tooling
charges. • No drill Limitations. • 1 to 6 layers
(prototype quantities) • 2 layers (small pro-
duction batches) • Fr4,1.6mm,35 um,HASL
(Pb free/Rohs/Weee) • Soldermask /
Silkscreen (optional) • Leadtimes from 2 -8
working days • Full DRC on all orders (we
manually check every file !!!). • PCB-POOL®

accepts direct outputs from 15 major layout
softwares (see our website for details)
Download our fully functional PCB LAYOUT
software FREE of charge.
Free Phone : 0800 3898560

Sky systems ltd

www.sky-pcb.com
Sky Systems Ltd. was established with the
vision and promise of providing manufactur-
ing high quality, cost effective solution and
one-stop service to meet the most demand-
ing of our customers' requirements.
We offer the followings:
• 1-12 Layers • Fr-4 / Cem-3 • HAL(Lead
Free),Flash Gold, Electroless Gold Plating,
OSP ,Immersion Silver, Immersion Tin • Gold
Finger • Soldermask • Silkscreen • Routing /
Punching / V-Cut • Online quotation • FREE
PCB Prototype with quantity orders • Short

lead time • Fast worldwide delivery •
Flexible quantity • CHINA Factory
For more information or request a quote
today from our web site.

DataQuest Solutions Ltd

www.dqsolutions.co.uk
As specialists in the supply of ultra high-
speed instrumentation cards for the PC
and industrial chassis, we provide an
advanced product range with the following
features:
• PCI, PCI-X, PCI-express & PXI formats.
• Signal capture (A/D) with sampling rates
up to 500M samples/sec.
• Simultaneous sample and hold with up
to 16 bits resolution.
• Waveform generation (D/A) to 125 mil-
lion samples per second.

• Digital I/O and pattern generation for a
wide range of logic levels.
• Storage of data to deep on-board RAM
or PC memory.
• Fully portable systems for many types of
card.

Surplectronics

www.surplectronics.com
Surplectronics is a brand new electronics
components e-Store with a difference. We
specialise in popular electronic compo-
nents from leading manufacturers at rock
bottom prices that you won't find else-
where.
New lines are arriving every day and many
of our products are one-off specials that
may not be repeated.
Our products are all new, unused and are
fully guaranteed. Ordering is simple, with
UK shipment free for orders over £50 and

payment is completed credit card. Goods
are shipped that day for prompt delivery.
Surplectronics - serious about surplus. Tel:
0870 850 5258

DB Technology

www.dbtechnology.co.uk/
Anechoic chamber and open area test site.
• Compliance Tests
• Fixes included. FCC Listed.
• Flexible hourly booking available.
• Rapid, accurate pre-compliance tests.

Telnet Ltd

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

TELNET
1 Stoney Court, Hotchkiss Way, Binley
Industrial Estate, CV3 2RL
Tel: 024 76650702 Fax: 024 76650773

Mobile: 07860 400683
Email: dales@telnet.uk.com

FTT

www.ftt.co.uk/PICProTrng.html
FTT (a Microchip Consultant Partner &
Training Partner) has developed a range of
courses – both distance learning and
instructor led – covering Assembly and C
Programming of PIC16, PIC18, PIC24 and
dsPIC microcontrollers. For each processor
family there are both C and Assembly pro-
gramming courses at: • FOUNDATION
LEVEL & • INTERMEDIATE LEVEL. For infor-
mation about these courses, advanced
courses and workshops such as: •
Advanced C Embedded & Real Time
Programming • TCP/IP & Ethernet • USB •

CAN • DSP • Motor Control • programming
using Real Time Operating Systems such
as uCOSII, CMX & Salvo and • other micro-
controllers, please inquire.
Tel: 020 8669 0769 Email: pictrng@ftt.co.uk

Microchip Technologies

http://www.microchip.com/
Microchip Technology Inc. is a leading
provider of microcontroller and analogue
semiconductors, providing low-risk product
development, lower total system cost and
faster time to market for thousands of
diverse customer applications worldwide.
Microchip designs, manufactures, and mar-
kets a variety of high performance compo-
nents for high volume, cost-effective embed-
ded control solutions, including 8- and 16-bit
PIC® microcontrollers; dsPIC® digital signal

controllers; development kits; serial
EEPROMs, more than 350 mixed-signal ana-
logue and interface products; KEELOQ secure
data transmission products; and the
PowerSmart® family of smart battery man-
agement products. Microchip's product solu-
tions feature compact size, integrated func-
tionality, and ease of development.

ADVERTISING ENQUIRIES CONTACT MATTHEW ON: 020 7933 8980