european business press

www.electronics-eetimes.com

February 2011

EE Times europe

For the creators of technology

Reaping the benefits of energy harvesting 10

Tabula unveils 3D-PLD EDA tool 8

Special Focus: Embedded Computing 33

Authorized distributor for more than 440 of the world's most trusted manufacturers.**

The industry's broadest product selection available for immediate delivery

Find contact and ordering information for your region at
digikey.com/europe

*A shipping charge of €18.00 (£12.00) will be billed on all orders of less than €65.00 (£50.00). All orders are shipped via UPS for delivery within 1-3 days (dependent on final destination). No handling fees. All prices are in euro and British pound sterling. If excessive weight or unique circumstances require deviation from this charge, customers will be contacted prior to shipping order. **Digi-Key is an authorized distributor for all supplier partners. New product added daily.

EE Times europe

PUBLISHER
André Rousselot
+32 27400053
andre.rousselot@eetimes.be

EDITOR-IN-CHIEF
Julien Happich
+33 153907865
julien.happich@eetimes.be

EDITORS
Philip Ling
+44 1622746580
phil.ling@eetimes.be

Christoph Hammerschmidt
+49 89 44 45 02 09
chammerschmidt@gmx.net

CONTRIBUTING EDITORS
Paul Buckley
+44 1962866460
paul@activewords.co.uk

Jean-Pierre Joosting
+44 1202543405
jean-pierre.joosting@eetimes.be

CIRCULATION & FINANCE
Luc Desimpel
luc.desimpel@eetimes.be

ADVERTISING PRODUCTION & REPRINTS
Lydia Gijsegom
lydia.gijsegom@eetimes.be

ART MANAGER
Jean-Paul Speliers

ACCOUNTING
Ricardo Pinto Ferreira

REGIONAL ADVERTISING REPRESENTATIVES
Contact information at:
http://www.electronics-eetimes.com/en/about/sales-contacts.html

EUROPEAN BUSINESS PRESS SA
144 Avenue Eugène Plasky
1030 Brussels - Belgium
Tel: +32 (0)2 740 00 50
Fax: +32 (0)2 740 00 59
www.electronics-eetimes.com
VAT Registration: BE 461.357.437
RPM: Brussels
Company Number: 0461357437

european business press

CONTENTS FEBRUARY 2011

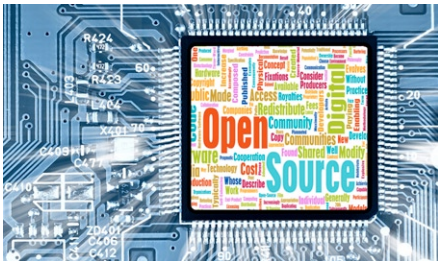

Managing open source licensing for semiconductors. 50

OPINION

NEWS & TECHNOLOGY

DESIGN + PRODUCTS

SPECIAL FOCUS: POWER MANAGEMENT

SPECIAL FOCUS: DATA STORAGE

SPECIAL FOCUS: EMBEDDED COMPUTING

READER OFFER

ELECTRONIC ENGINEERING TIMES EUROPE is published monthly in 2011 by European Business Press SA, 144 Avenue Plasky, 1030 Bruxelles, Belgium Tel: +32-2-740 00 50 Fax: +32-2-740 00 59 email: info@eetimes.be. VAT Registration: BE 461.357.437. RPM: Brussels.
Volume 13, Issue 2 EE Times P 304128
It is free to qualified engineers and managers involved in engineering decisions – see: http://www.electronics-eetimes.com/subscribe

Software is the new hardware

By Philip Ling

While the importance of embedded software is now well understood by most semiconductor manufacturers, and the increasing dominance of the ARM architecture is encouraging more of them to differentiate themselves through the software support they offer, the same evolution is not readily apparent within the distribution world.

Unlike just a few years ago, when Independent Device Manufacturers (IDMs) chose to keep the higher yielding accounts in-house, today the distribution market handles an increasing amount of the IDM's business, and accounts for the majority of semiconductor sales. Despite this, it maintained a hardware-centric business model, until now.

In a bold attempt to align its MCU sales with its overall business share in the distribution market, Silica has launched the 'Core 'n More' programme. The intention is to increase its share of the MCU market from 9% to 15% by 2014, through a homogeneous strategy that will encompass all of its business activity – not just sales – and be rolled out across all European regions.

Silica is very clear about its intention to focus on the ARM architecture to achieve this growth, but it will be in the provision of software expertise that it expects to differentiate itself from other distributors, something that isn't going to directly provide a revenue stream as there is currently no intention to start selling software tools and services directly. Instead the strategy is to partner with other companies with expertise in this area, while also building up its own in-house expertise.

Silica's regional vice president for central Europe, Karlheinz Weigl, explained that despite the MCU market being the biggest in Europe for distributors, Silica hadn't focused on it until now. The launch of the 'Core 'n More' programme is the culmination of a year's planning but isn't expected to provide a return on that investment – in the form of an average of 15% market share across Europe – until 2014. The three-year strategy reflects Silica's position in the MCU market; for whatever reason it isn't selling as many MCUs as its competitors.

Weigl explained that while Silica has eight out of the top ten MCU suppliers, it didn't have a structured approach to selling

Karlheinz Weigl, Silica's Regional Vice President for Central Europe: 'Spreading an homogeneous strategy across all of Europe'

MCUs, unlike the rest of its line card. As MCUs are rarely used in isolation, Weigl is hoping the new strategy will also increase the company's attach rate', or the number of peripheral components sold with every MCU. At the moment the attach rate is around 2.2 to 1, or 2. 2 components sold to support every MCU, by 2014 Weigl would like that ratio to be 5 to 1.

Today the MCU market is highly fragmented across 8-, 16- and 32-bit architectures, but Weigl expects that to change in favour of 32-bit, powered almost exclusively by ARM-based devices. So while Silica currently sells more 8-bit devices than 32-bit, it believes the growth potential is in the high end, which means that all distributors – including Silica – should benefit from greater sales simply by providing 32-bit MCUs. Its strategy to increase its market share will, therefore, actually mean it needs to grow quicker than the rising market.

Part of the strategy to achieve that is to make software the key differentiator, through a European Software Support Team. Silica intends to hire regional software experts who will develop in-house expertise on all the major software development environments, operating systems and tools. FAE training will also be key, but Weigl was very clear about which architecture would get the biggest investment; ARM. And it isn't all ARM-based devices, but the 'sweet spot' of Cortex-M3, -M4, -A8 and -A9. This notably excludes Microchip's PIC32 family, which is based on the MIPS core, and the Cortex-R series of ARM processors, which ARM has recently extended with the introduction of the Cortex-R5 and -R7.

All of ARM's Cortex cores are based on the ARM v7 instruction set architecture, making them binary compatible –something ARM's customers insist upon due to the massive investment they have in their embedded software, particularly in cellular codecs which is one application area where the Cortext-R4 has been popular.

The architectural enhancements found in the new Cortex-R cores have been predominantly led by the needs of software engineers, providing hardware-assist for functions that would otherwise need to be coded and thereby making the software engineers' job 'easier'.

But because the Cortex-R isn't currently in demand within the mainstream industrial market, it doesn't immediately fall within the remit of the Core 'n More programme. Similarly, because MCUs based on the Cortex-Mo are, relatively speaking, simpler to use, they aren't targeted as an area where Silica's new strategy can provide the crucially important differentiation. ■

Philip Ling is Editor for embedded electronics at *EETimes Europe* - reach him at phil.ling@eetimes.be

ST sees automotive and consumer electronics converge

EXECUTIVE INTERVIEW

By Christoph Hammerschmidt

FOR STMICROELECTRONICS, the current automotive boom translated into record results. Traditionally, the largest European semiconductor vendor has a strong commitment to automotive electronics. EE Times Europe talked with Joseph Notaro, ST's Director of Automotive Application, about the company's strategy in this demanding market.

EE Times Europe: What does the automotive market represent for ST?

Joseph Notaro: Historically and since the merger between STS and Thomson in 1987, the automotive sector has always been one of ST's key segments. We have a comprehensive approach of the automotive segment and I suspect that our product portfolio is one of the broadest among our competitors. Traditionally, we have achieved between 13 and 15 percent of our revenues with automotive products.

EE Times Europe: Where do you see ST's strength in the automotive market?

Notaro: There are two main aspects. One is more technical, one is more behavioural. We have a good balance between commitment and flexibility and we have always been strong in making strategic alliances.

EE Times Europe: Can you name some of these alliances?

Notaro: For instance we have alliances with tier ones such as Bosch and Magneti Marelli, but also with other third parties such as Mobileye and Navteq. We've also joined forces with one of our competitors, Freescale, to develop next-generation microcontrollers. Automotive electronics starts to converge with other market segments, such as the multimedia segment and industrial applications. In all of these segments, ST ranks among the top suppliers.

EE Times Europe: Unlike most of your competitors, ST's presence in the automotive industry spreads evenly across various segments such as infotainment, powertrain, body control, safety etc. Where is your focus in terms of application expertise and perhaps in technology?

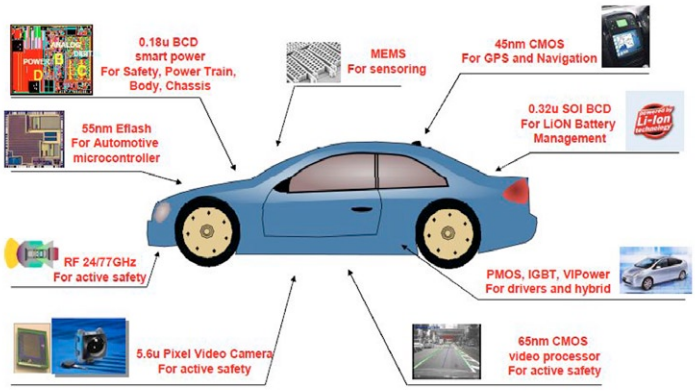

ST's automotive product range covers more or less all in-vehicle segments.

Notaro: This has to do with the way we are structured in terms of R&D. We have common development teams but then different application teams each one focusing on the segments you mentioned. For instance, we have an infotainment team which maintains a good relationship and internal idea exchange with our consumer business and telecom group but, they are focusing more on the high end of in-car multimedia and navigation systems. And we have our traditional automotive electronics division where the focus is more on powertrain, safety, chassis and car body. Thanks to ST's unique structure, we're also becoming a big player in the ADAS area (Advanced Driver Assistance Systems). This includes topics such as blind spot detection or lane departure warning. We make automotive-qualified CMOS cameras, which is quite unique among pure semiconductor companies. We also sell radar chips, 24/26GHz, transmitters and receivers and we develop 77/79 GHz parts for our customers.

EE Times Europe: ST also has a strong position in the market for consumer MEMS. Wouldn't it make sense to leverage this market for automotive applications too?

Notaro: Definitively. The MEMS market in automotive is large, but the growth opportunities have been more outside of our automotive activities. ST has made the strategic decision to focus first on the consumer market. Every tablet computer, every game console today contains an ST MEMS. We have shipped more than one billion MEMS devices to the consumer market. We know that the automotive market takes more time and has different requirements in terms of quality and reliability. We do sell MEMS in smaller quantities to our automotive customers but we remain cautious in this market, because we want to be successful without taking any risk.

EE Times Europe: I learned that the current MEMS supply is a bit difficult. Which dynamics drive this market?

Notaro: We don't see issues in the automotive MEMS market. More generally, the supply situation is difficult for semiconductors. The situation is getting better, but the automotive market has really been booming. There has been a big dip in production but we are catching up quickly with the increased demand. Emerging economies are driving the market at a time when in-car semiconductor content is growing significantly.

Advanced safety, blind spot detection, cameras in vehicles, all these devices have penetration rates much higher than we predicted a few years ago, so do infotainment

systems. Smart lights that can swivel, level, or automatically switch between high beam and low beam are another example of new electronic content in cars. The acceptance and the push from the carmakers have been incredible which means the demand for all this optional equipment is increasing much faster than anyone had predicted.

EE Times Europe: What is ST's strategy with respect to microprocessor architecture? Is there something like a preferred architecture at ST?

Notaro: For sure. If you look at the traditional automotive domain, including powertrain, safety, chassis and car body, the PowerPC architecture is ST's chosen architecture. We have a joint development activity with Freescale on the PowerPC. Our goal, our vision is to make the PowerPC the de-facto standard on the market in these applications. We've developed applications that went to production late 2010. On the infotainment market where the convergence between the consumer and telecommunications segments is remarkable, the ARM architecture is more present. With our Cartesio family, our strategy for these markets is really ARM-based. This ARM-based platform for GPS, connectivity, storage, and audio processing provides everything you would expect to find on a multimedia processor. We use the same platforms for in-vehicle navigation systems and for PNDs. Garmin, for example, is powered by our Cartesio family.

EE Times Europe: Which technology trends do you see in the automotive market for the years ahead? Which domains will grow the fastest?

Notaro: The main drivers for automotive semiconductors are in-vehicle multimedia and new environmental and safety legislation. There is a lot of activity to reduce CO_2 emissions; some reductions are mandated by the EU, some by the North American government. There are also mandates in Japan and in China to reduce CO_2 emissions.

EE Times Europe: And how can semiconductors contribute towards this end?

Notaro: There are two different angles to look at it. There is a lot of hype around electrical vehicles (EVs) and today, this represents a tremendous growth potential for the semiconductor content of vehicles. The number of semiconductor devices necessary to drive an electric motor is much higher than what a traditional combustion engine would require. But for EVs, the biggest challenge is to increase the energy density and reduce the weight of the battery at reasonable costs. At the same time, there is a lot of pressure to increase the efficiency of the existing internal combustion engine. We see a lot of opportunities in what we call vehicle electrification, whereby different types of loads such as LEDs replacing incandescent lamps or start-stop systems maximizing the efficiency of the car, must be driven by different devices. You need more sensors, more actuators and more power devices. Controlling these loads also requires microcontrollers with more computing power, able to run more advanced algorithms dealing with multiple sensor interfaces.

EE Times Europe: Everybody is talking about the connected car. How will the connected car affect the semiconductor content and even the driving experience?

Notaro: The technology is already there. The challenge will become the infrastructure. This is another interesting trend in automotive. The car is not considered anymore an object independent of the rest of the world. With the electric car, the vehicle becomes a part of your everyday cycle. You plug it at home, you drive, you connect to get information, you connect the car to your home electronics to download music from your music server at home etc. This will change our driving habits. Today, traffic updates are too slow and not very reliable. Real-time traffic updates, information about charging stations for electric vehicles, all of this will change the way we drive. The infrastructure will be an enabler for the electric and the connected car.

EE Times Europe: What developments are you expecting in the power electronics segment?

Notaro: The power semiconductor area is probably the biggest growth area for Electric Vehicles (EVs) because of the amount of power to control. The more efficient your power semiconductors are the more efficient your electric car can be. Efficiency translates into range, how far you can go with your EV with a given battery size. Today, the trend is towards higher voltages; we talk about 400V batteries. Dealing with these batteries, you need semiconductors that can handle up to 1000V.

But this may not be sufficient. We are working on new compound materials

Figure 2: Joseph Notaro, ST's Director of Automotive Application sees synergies between all application fields in automotive electronics.

such as silicon carbide and gallium nitride power transistors which really increase the components efficiency, helping maintain the current size of power semiconductors. SiC could work at higher temperatures which means that you could handle higher powers for a given die size. This could lead to big advantages in cooling your electronic module. Doing away with liquid cooling for EVs would translate into big cost and weight savings. I think that within the next five to ten years, GaN and SiC will be the big breakthrough in vehicle power electronics and of course in industrial markets.

EE Times Europe: Will ST bring SiC or GaN semiconductors to market within this time frame?

Notaro: We are already in production today with SiC diodes. Some technology breakthroughs are coming that will enable us to produce true silicon carbide MOSFETS within the time frame I mentioned. It is challenging to go up in voltage range with GaN but we have lots of research in this direction. Being an engineer, I am convinced that there are very few technical hurdles that cannot be overcome. And when we produce a device we always have in mind high volume, low cost, high quality. This is one of our differentiators.

EE Times Europe: This brings me to the next question: What is ST's manufacturing strategy? Do you manufacture all semiconductors in-house?

Notaro: We will be one of only two companies that will have 55nm embedded Flash to be produced in-house, one hundred percent. All our smart power and mixed signal devices are developed and manufactured in-house. This shows our commitment to the automotive sector. ■

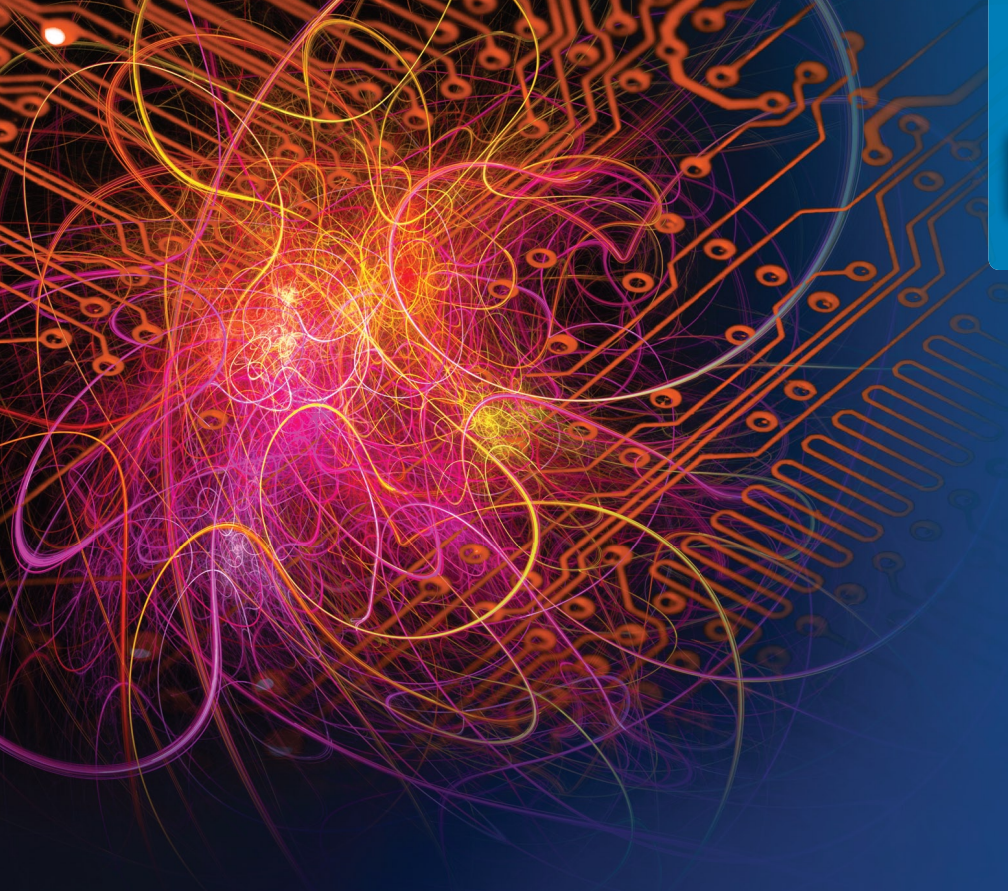

mouser.com
Semiconductors and electronic components for design engineers.

We've got what design engineers need. **The future.**

You'll find the newest products for your newest designs.
Get what's next at Mouser.

embedded world 2011
Exhibition&Conference
... it's a smarter world

mouser.com

MOUSER ELECTRONICS®
a tti company

Nürnberg, Germany, 1 - 3rd March 2011
Hall 9, Stand 168

Tabula takes 3D programmable logic sky-high with cloud-based EDA tool

By Julien Happich

ABOUT A YEAR AFTER announcing its Spacetime's 3D architecture, Tabula is unveiling Stylus, a secure cloud computing environment designed to unleash the power of its ABAX 3PLD devices. Since the public launch of the 3D programmable logic devices (3PLDs), Tabula has been busy qualifying its silicon for production on TSMC's 40nm process and shipping dozens of functional samples to early adopters. What was missing for the mass adoption of the clever time-driven multiplexed reprogrammable logic was a reliable design tool that could take regular RTL code and port it seamlessly to the ABAX devices. That's what Stylus is about, hassle-free with all the compute power of the cloud.

According to Alain Bismuth, Tabula's Vice President of marketing, all the physical devices have exceeded the benchmarks of their initial announcement and 2011 will be the ramp up year for volume production. "Tabula has not only delivered on its promise to deliver more performance at an unmatched price point, but we are now able to bring a free tool environment that will enable any designer to program our chips transparently without having to deal with the specifics of its time-driven architecture" he commented.

In the course of 2010, the company gained several design-wins and some of its chips are already running on base-stations. "We have seen many telecoms and consumer OEMs who were still favouring ASICs or ASSPs versus traditional FPGAs, on a cost and performance basis, turn to the ABAX devices and benefit from the re-programmability at a price-point an order of magnitude lower than current FPGAs offer" Bismuth added. This complements the company's view that with a 4x signal processing throughput and a 2.5x logic density achieved through up to eight folds of sub-clock cycle reconfiguration, its devices are stepping up the design shift from ASICs to FPGAs. This is one notch beyond the trend that traditional FPGAs manufacturers have been riding to gain market share against ASICs.

2010 helped tune the in-house development tools and validate Stylus across many design scenarios. As a dedicated synthesis and place-and-route package supporting 3PLD devices, Stylus offers an intuitive, familiar design environment for ASIC and FPGA designers through a browser-like GUI.

At the heart of Stylus is Spacetime-optimized synthesis coupled with timing-driven 3D place-and-route that automatically and transparently maps standard RTL directly into Spacetime.

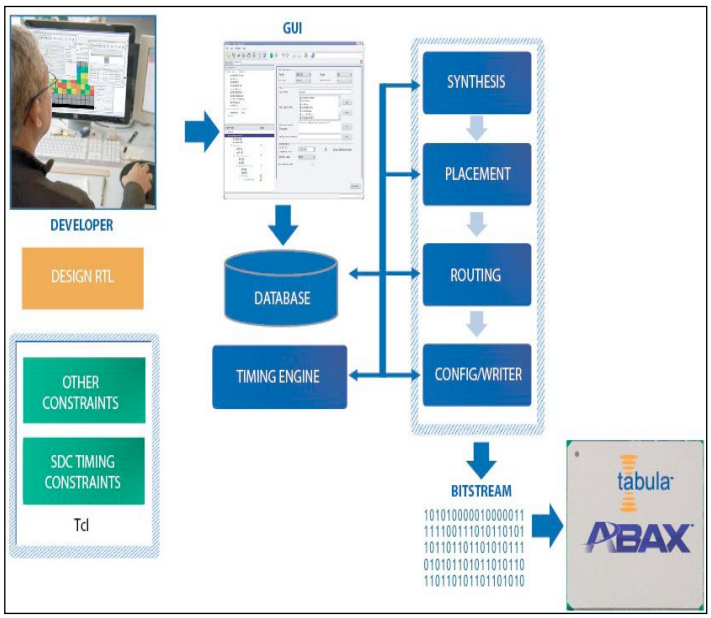

Fig1: The cloud-based Stylus design flow accepts standard VHDL/Verilog/ System Verilog inputs

The tool accepts standard VHDL/Verilog/System Verilog inputs as well as SDC constraints. Several productivity-enhancing features have been implemented, including advanced floor-planning, static timing, and power analysis. Debug logic can be integrated on-chip and direct probing of select nodes give users an inside-out view of the circuit's behaviour.

IP cores from Tabula's leading IP partners can be integrated and supports design libraries for industry-standard logic simulators.

The company opted for a cloud-computing environment to ensure that designers always have access to the latest software without either the overhead associated with IT management or the expense of compute servers. The web access is also meant to enable real-time, on-site-like, technical support without file transfers. To further accelerate device propagation, Tabula also announced a comprehensive development kit featuring the ABAX A1EC04 3PLD, a part providing 390k LUTs, 5.5 MBytes of RAM blocks, and 48 channels of 6.5Gbps SerDes. The platform will enable designers to test a rich variety of parallel and serial interfaces such as DDR3 and PCI Express Gen2, as well as embedded systems based on the popular V1 ColdFire soft CPU. ■

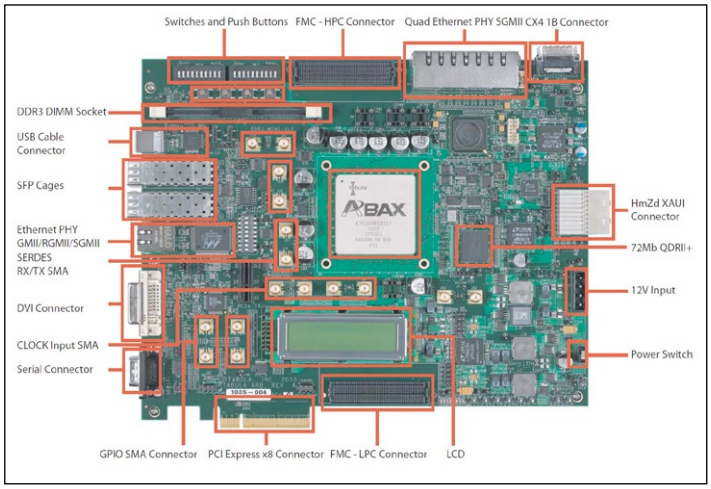

Fig2: A feature-rich development board with the ABAX A1EC04 3PLD.

Join the element14 Treasure Hunt

Visit element14 for your chance to

Win an iPad

element14

GET YOURSELF CONNECTED

www.element14.com

Seeing Green: A look into 2011 at growing technology trends inspiring new product design

TECHNOLOGY REMAINS ever changing. It's shaped by our demands for smaller, faster and smarter products engineered to make our lives easier. Not only that, it needs to be responsive to our planet's rising environmental concerns. Terms like reduced carbon footprint, green energy and energy harvesting are the new mantra being chanted by today's companies and engineers.

To understand how green technology is impacting today's product design, we consulted Mouser Electronics, a leading global semiconductor and electronic components distributor that works closely with design engineers and buyers. Who better to ask than a company known for its rapid introduction of the latest products and most advanced technologies used for prototype development?

Future of SSL looks bright

When talking with Mouser, our conversation turned to the topic of lighting. Solid State Lighting (SSL) took the spotlight and centered on light emitting diodes (LEDs). In recent years, Mouser has seen the SSL market take off. Decreases in price, energy savings as a result of greater efficiency and longer life have led to a number of applications being developed.

Where traditional lighting is used, an SSL solution can be implemented. SSL offers more brightness per watt while producing less heat. The size of SSL devices can be extremely small, for use in more places, as well as surface mounted for greater design flexibility. Finally, the life of an SSL is longer (lasting for more than 35,000 hours) in comparison to the incandescent bulb (averages only 750 hours).

LED lighting is forecasted to account for 16% of the worldwide lighting market by 2013, accounting to over 14.62 billion Euros in revenue. This is based on a 25% year on year growth. Greentech Media Research is even predicts a 30% growth rate in 2011.

What's behind this growth? The largest use of LEDs has been mobile devices. This will continue to grow in 2011 as the mobile industry ramps up its development of smartphones. Automotive manufacturers are switching to LEDs for daytime running lights due to greater visibility and energy efficiency. Another factor is the upcoming sales ban on incandescent lights. Although the deadline varies from country to country, this planned ban will lead to less incandescent lights being sold by 2012 – spearheading alternative solutions to fill the void as the Edison bulb gives way to SSL.

Energy harvesting ripe for growth

In the quest to find alternative energy sources, much of the buzz centers on micro-energy harvesting. Simply put, it's the capturing of tiny amounts of energy from the environment through vibration (piezoelectric and oscillating mass), temperature differentials (TEGs or thermoelectric generators), light (solar panels) and RF (induction through ambient electromechanical waves), along with other sources. The concept of snaring free or waste energy and converting it into reusable power is sparking industry-wide attention. Mouser sees the energy harvesting trend gaining more momentum in the year ahead. The cost of harvesting has now decreased to where it's a practical alternative to traditional power sources; the seeds for growth are firmly planted.

Energy harvesting enables long-term, maintenance-free operation of low-power electronic devices, especially wireless sensors. It holds the promise of replacing batteries with green, virtually infinite power for wireless devices. One company is already in the process of converting waste heat into a sustainable, maintenance-free power supply for wireless sensor devices. A built-in chip thermogenerator takes a few degrees of temperature differential and harvests thermal energy to operate the wireless sensor node. By employing wireless sensors, design engineers can create a hands-off solution that lasts the lifetime of the application. Therefore, users avoid having to maintain hundreds, thousands of batteries.

Ambient RF energy also stands at the forefront as a means of supplying power to wireless devices. Energy from RF transmitters offers the unique benefits of providing predictable and consistent power over distance without being connected to a power source. Not to mention, ambient RF energy is currently available from billions of radio transmitters worldwide, including mobile phones, handheld radios, mobile base stations, plus radio and television broadcast stations. This is leading to the creation of

new, battery-free devices and even allows battery-powered devices to be wirelessly trickle-charged. Devices with this capability have the advantage of mobility while charging.

The numbers are convincing. ABI and iSupply estimate the number of mobile phone subscriptions has surpassed 5 billion globally, and ITU calculates there are over 1 billion subscriptions for mobile broadband. Also, consider the number of WiFi routers and wireless end devices such as laptops. The Darnell Group, a market research firm, estimates 200 million energy harvesters and thin-film batteries will be in use by 2012. Overall, energy-harvesting components are forecasted to exceed 2.93 billion Euros by 2020. The ground is fertile for new product design.

Next-Gen processors turn green

Energy harvesting isn't the only green energy solution taking form. In energy management, the industry has seen a huge movement towards components with lower-power consumption. The trend to reduce power consumption and extend battery life is taking precedence. Manufacturers of microcontrollers (MCUs) are in a literal race when it comes to providing high performance at decreasing power consumption. Simply, MCUs are going green.

Manufacturers are engineering MCUs that draw nanoamperes of current in sleep mode. In fact, the newest 32-bit microcontrollers can consume one quarter of the energy of other 8-, 16- and 32-bit MCUs. Extensive clock gating within the chip design is key in reducing power requirements, but so is smart use of different power modes and well designed systems.

What's advancing the growth for green MCUs in 2011? A report from GBI Research points to the demand for low-power multicore processors in laptops and smartphones. The rapid growth of smartphones usage is driving the MCU market. The number of smartphones in use globally is expected to hit 1.7 billion by 2014, driven by demand in emerging markets of Asia-Pacific, notably China, India, as well as Central and South America with Brazil leading. Another fact, laptops have already overtaken 2009 worldwide desktop sales. New product categories like tablets and netbooks have registered strong initial sales, increasing their stake in the global PC market, furthering the demand for greener MCUs.

Keeping pace with what's next

Mouser is taking the lead in providing new, energy friendly technologies to meet the increased demand for high-performance semiconductors and electronic components with lower energy consumption. The company's Advanced Technology Group is concentrating on energy-harvesting components and Solid State Lighting along with leading-edge advancements in ultra-efficient MCUs for developing innovative product designs and prototypes.

"Along with providing these emerging technologies, Mouser is deeply committed to producing solution-based content on its website as a resource to the design community", Russell Rasor, Mouser VP of Advanced Technology, states. "These are all technologies that require much more than a list of part numbers... designing for 2011 is going to require a lot of thought. Design engineers can review downloadable data sheets, supplier-specific reference designs, application notes, technical design information and engineering tools for new product design".

About Mouser

If technology is made to transcend boundaries, shouldn't the company that delivers it be too? Mouser thinks so and has 14 locations globally with plans to expand further.

More specifically, Mouser has offices strategically located throughout Europe, including Germany, France, U.K. and Italy. This is in direct response to strong European design community. Other factors behind Mouser's exceptional growth in Europe has been the company's unyielding focus on providing best-in-class service, excellent product information, rapid delivery and order accuracy combined with the newest, most cutting-edge products from top suppliers. The result: Engineers can design in the latest technologies into their prototypes – giving an edge in features, product lifecycle and speed to market.

The engine driving much of Mouser's growth is www.mouser.com. It houses in excess of 1.8 million parts online from more than 400 leading suppliers. The site is updated with new products every day and clearly identifies components Not Recommended For New Design (NRND) and parts planned for obsolescence. The site also provides an industry-first interactive catalog, data sheets, supplier- specific reference designs, application notes and many other useful tools. Plus, the website has more than 650 product microsites, covering detailed information and key features on the newest components ready for design. Furthermore, it supports 16 different languages and currencies. All these capabilities have earned Mouser the best distributor website for technical information, search and ease of use.

It's no surprise that Mouser has been receiving a lot of industry recognition in Europe. Recently, Mouser was individually honored by both Bourns and Elektra as the 2010 European Catalogue Distributor of the Year. Additionally, Mouser is continually strengthening its broad linecard, featuring such leading suppliers as TI, Maxim, Panasonic and EPCOS, to name a few. What's more, Mouser and its parent company, TTI, Inc., even caught the keen eye of Warren Buffett, joining his Berkshire Hathaway family of companies in 2007. All of which translates into even greater confidence and reassurance for today's European buyers and design engineers in search of the newest products for their newest designs – harvesting what's next is what Mouser does best. ■

WIRELESS COMMUNICATIONS

ESA on analog self-steered antennas

By Paul Buckley

BULKY PRESENT GENERATION satellite dishes and ground terminals could become relics of the past thanks to research currently being conducted for the European Space Agency (ESA) by Queen's University Belfast's Institute of Electronics, Communications and Information Technology (ECIT) which aims to develop discrete self-aligning flat antennas.

The research could lead to a one-size-fits all solution that can be optimised for a variety of technologies presently used to deliver satellite broadband and television to travellers as well as customers in broadband 'not spots'. ECIT is currently working on an 18 month ESA project with the aim of developing a completely self-contained solid-state self-steering antenna that is much lighter and less power hungry than current alternatives.

The team being led by Professor Vincent Fusco plans to complete work on a 1.6GHz demonstrator - capable of providing transfer rates of 0.5Mbits/s - with a power requirement of just 2 watts. The device will ultimately have the capability to operate at 20-30 GHz in order to provide much greater bandwidth. The design currently being worked on is a 4x5 element planar array measuring 30 cm by 40 cm and just 12 mm deep.

The circuits are entirely analogue and incorporate specially adapted phase locked loop circuits. By contrast, conventional circuits convert incoming signals to digital, process them electronically and then convert them back to analogue. This however limits their frequency, and increases their complexity, cost and power requirements.

Queen's University built the world's first 65 MHz self steered antenna a number of years ago. Since then, it has built a close relationship with ESA to whom it is now the main supplier of quasi-optical filters.

Dr Buchanan the lead engineer on the project who received ESA's Best Young Engineer award for his work in the field said: "The work is especially exciting because it has involved taking a piece of pure university research and bringing it into the real world.

We believe that self-tracking antennas offer the prospect of much simpler and more cost effective alternatives to other current approaches. That, we believe, makes them ideally suited to a variety of end uses".

"For example, satellite broadband aircraft antennas are extremely complex. They need to be linked into the plane's onboard navigation system in order to find the satellite."

"We believe that across these applications the solution we are currently working on could reduce power consumption by a factor of 10, weight by a factor of five and cost by a factor of four". ■

Molybdenite better than graphene as an alternative to silicon

SEMICONDUCTOR PHYSICS

By Julien Happich

SMALLER AND MORE energy-efficient electronic chips could be made using molybdenite. EPFL's Laboratory of Nanoscale Electronics and Structures (LANES) publishes in the journal Nature Nanotechnology a study showing that this material has distinct advantages over traditional silicon or graphene for use in electronics applications.

The research carried out in the LANES has revealed that molybdenite, or MoS_2, is a very effective semiconductor. This mineral, which is abundant in nature, is often used as an element in steel alloys or as an additive in lubricants. But it had not yet been extensively studied for use in electronics.

"It's a two-dimensional material, very thin and easy to use in nanotechnology. It has real potential in the fabrication of very small transistors, light-emitting diodes (LEDs) and solar cells," says EPFL Professor Andras Kis.

He compares its advantages with two other materials: silicon, currently the primary component used in electronic and computer chips, and graphene.

One of molybdenite's advantages is that it is less voluminous than silicon, which is a three-dimensional material. "In a 0.65nm-thick sheet of MoS_2, the electrons can move around as easily as in a 2nm-thick sheet of silicon," explains Kis.

"But it's not currently possible to fabricate a sheet of silicon as thin as a monolayer sheet of MoS_2." Another advantage of molybdenite is that it can be used to make transistors that consume 100,000 times less energy in standby state than traditional silicon transistors.

A semiconductor with a "gap" must be used to turn a transistor on and off, and molybdenite's 1.8 eV gap is ideal for this purpose. This gap offers a great level of control

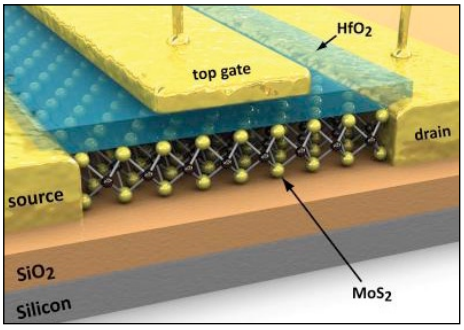

A digital model showing how molybdenite can be integrated into a transistor. Credit: EPFL

over the electrical behaviour of the material, which can be turned on and off easily. The existence of this gap in molybdenite also gives it an advantage over graphene.

Considered today by many scientists as the electronics material of the future, the "semi-metal" graphene doesn't have a gap, and it is very difficult to artificially reproduce one in the material. ■

Our High Performance Solution
Inverterised Air-conditioning reference board

Infineon is renowned for offering best in class discrete devices and ICs – now with the inverterised air conditioning reference board, Infineon can present system expertise in the fast growing inverterised air conditioning market.

Key features and benefits of Infineon's Air-conditioning reference board

- Size and thermally optimised reference platform for inverterised air-conditioner system to drive 2 inverter stages on one PCB
- Innovative cooling for high power SMD IGBTs
- Displaying › 97% PFC efficiency
- Full power electronic SMD assembly for high capacity production

For further information please visit our website:

[www.infineon.com/aircon]

CEA-Leti ramps up 300mm line for 3D-integration

By Julien Happich

LAST JANUARY, CEA-LETI inaugurated a brand new 300mm line dedicated to shared research and development on 3D-integration, with plans to go well beyond the current 2.5D solutions that represent wired die stacking or other similar a posteriori system-in-package assemblies.

By adding this technology to its existing 300mm CMOS R&D line, Leti can offer heterogeneous integration technologies to customers on both 200mm and 300mm wafers. The dedicated line includes 3D-oriented lithography, deep etching, dielectric deposition, metallization, wet etching and packaging tools and will be available for worldwide customers and partners. The line will enable Leti to apply its 3D-integration generic processes on 300 mm wafers, with a 3D toolbox that includes a large portfolio of through-silicon vias (TSVs), and advanced capabilities in alignment, bonding, thinning, and interconnects in specific integration schemes for manufacturing optimized die stacks and building efficient advanced-systems solutions. This will be done in close collaboration with local design and characterization platforms.

EETimes Europe got the chance to visit the new clean room and get a flavour of what's to come in 3D integration, with several key industry partners invited to showcase their research and development strategies during a two-day seminar organized at Minatec in Grenoble.

CEA-Leti's CEO Laurent Malier observed that while process nodes are still going to shrink, more and more companies are considering vertical integration and through-silicon vias (TSVs) to further reduce the footprint of their systems at any given process node. The two main benefits that integrators are looking for, apart from device compactness, are an increased number of I/Os for a higher data bandwidth between the stacked chips and the possibility to achieve the best compromise in die area between different technology nodes.

"The new line has been defined so that it can be used across several generations of process nodes, with a real capacity to produce working prototypes," said Malier. "This will enable us to open up new paths of optimisation, for example getting the best of an established 130nm or 65nm process side by side with the best of a 40nm process without paying the cost of a full IC built at 40nm." This sort of compromise on technology nodes will occur more and more often, as long as stacking older processes becomes a cheaper and more power efficient alternative to the next more expensive node being used in one single layer.

"As a research lab, we don't aim for mass production but having a pilot line could be the next step for small prototype production, to fill the gap between development and mass production," Malier added. The recent lab's partnership with Presto Engineering to develop test and analysis capability for 3D semiconductor devices will be key from a supply chain test point. By stacking dies, the wafer test costs go up because it gets harder to access the chip data for debugging repairing etc... Presto Engineering will be working on improving yield by finding test methodologies as well as parameterizing the stacking processes.

CEA-Leti's CEO Laurent Malier (left) receiving the fist 3D-ready 300mm silicon wafer from Thierry Mourier (right), 300mm TSV line manager.

For the last five years, CEA-Leti has been evaluating building blocks for 3D integration, it has also developed test circuits to validate its 3D processes and the first 3D-ready 300mm silicon wafers are already out. Moving to a 300mm line will help understand the impact of TSVs on advanced CMOS design rules (both mechanical and electrical). The building extension will be finished in May and the line will be fully operational for CEA-Leti's partners before the end of 2011.

The research centre's partners will be able to test and characterise their design ideas with working assemblies. This physical line will go hand-in-hand with CEA-Leti's earlier initiative to develop a 3D EDA toolbox together with R3Logic, and possibly open 3D interfacing standards. What the EDA vendor wants is to encapsulate all the properties of the stacking process in one single 3D technology file to be manipulated across the whole design flow. That way, a tool could look at different libraries and databases and identify what the best process would be for each circuit block and optimize different areas for thermal dissipation and timing constraints within a stacked architecture.

"With the toolbox development and the line, we have set up something unique in Europe, and we'll be able to offer technical support both at design and at manufacturing level," said Dr. Ahmed Jerraya, head of design programs at CEA-Leti.

Further developments on CEA-Leti's roadmap include demonstrators for silicon interposers and, on the 2012 horizon, demonstrators for 3D Network-on-Chip (NoC) routers enabling intelligent TSVs and maybe logic-capable TSVs for reconfigurable 3D routing between stacked CMOS dies.

With shrinking process nodes, connecting the dies' I/Os in flip-chip packages or even to a silicon interposer will require finer and finer solder bumps and copper pillars. Using old-generation lithography equipment could be the next step to pattern the minuscule in-package redistribution layers. This is Replisaurus Technologies' approach; with an electrochemical pattern replication (ECPR) technology the company presented during the seminar. The ECPR technology (see figures 1a to 1d) combines the precision and resolution of advanced lithography with the efficiency of electrochemical deposition into one single electrochemical metal printing step, thus eliminating many steps from the traditional process. A master electrode with an insulator pattern is pre-filled with an anode material (Cu for example). By applying an external potential, the anode mate-

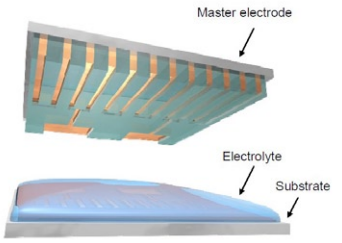

Figure 1a. Master with anode material and substrate with seed layer and electrolyte on.

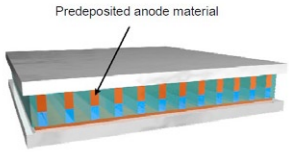

Figure 1b. Master and substrate in contact.

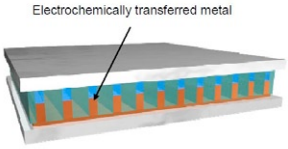

Figure 1c. Electrochemical metal transfer.

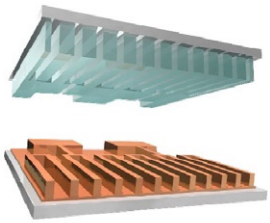

Figure 1d. Separation, followed by seed etch.

rial is transferred through an electrolyte onto the seeded conductive substrate. Separating the master from the substrate leaves an inverse metal replica of the master electrode pattern.

Only the seed conducting layer must be etched to leave the final trace. The company recently signed a multi-year agreement with Leti to optimize the development of reusable master electrodes for applications such as integrated passives, copper pillars and 3D integration. ■

LED surge propels MOCVD market

By Mark LaPedus

LED PRODUCTION

INVESTMENT BANKING firm Barclays Capital has raised its forecast in the metal organic chemical vapor deposition (MOCVD) equipment arena from 900 tools shipped in 2011, up from 800 in its previous forecast. The total MOCVD tool market was close to 800 unit shipments in 2010, up by quadruple figures from 2009, said C.J. Muse, an analyst from Barclays Capital, in a report.

The new forecast follows an ongoing surge for LED demand. MOCVD is one of the critical tools used in LED production. Aixtron, Nippon Sanso and Veeco are battling each other for share in the MOCVD arena. The MOCVD market is not cooling down, as some had thought. "Although sentiment on the LED equipment vendors has been fairly negative over the last few months, driven by concerns about China subsidies coming to a halt, our work suggests that this is not the case," he said.

Back in October, "we suggested that China MOCVD demand was likely to grow substantially from (about) 270 tools in 2010 to (about) 450-500 in 2011. This drove our overall MOCVD market estimate of (about) 800 tools in 2011," he said.

"However, having compiled a bottoms-up analysis of the MOCVD market, our estimates now point to China 2011 demand of (about) 575 tools. And assuming MOCVD demand in the rest of the world declines by (about) 39 percent year-over-year in 2011 to (about) 325 tools, we now estimate that the total MOCVD market is likely to reach (about) 900 tools," he said. ■

LOW-TEMPERATURE TRANSISTORS

Organic field-effect transistors with bilayer gate insulator show high stability

By Julien HAPPICH

RESEARCHERS FROM the Georgia Institute of Technology published a new method of combining top-gate organic field-effect transistors with a bilayer gate insulator, allowing the transistor to perform with incredible stability while exhibiting good current performance. In addition, the transistor can be mass produced in a regular atmosphere and can be created using lower temperatures, making it compatible with the plastic devices it will power.

The research team used an existing semiconductor and changed the gate dielectric because transistor performance depends not only on the semiconductor itself, but also on the interface between the semiconductor and the gate dielectric.

"Rather than using a single dielectric material, as many have done in the past, we developed a bilayer gate dielectric," said Bernard Kippelen, director of the Center for Organic Photonics and Electronics and professor in Georgia Tech's School of Electrical and Computer Engineering.

The bilayer dielectric is made of a fluorinated polymer known as CYTOP and a high-k metal-oxide layer created by atomic layer deposition. Used alone, each substance has its benefits and its drawbacks.

CYTOP is known to form few defects at the interface of the organic semiconductor, but it also has a very low dielectric constant, which requires an increase in drive voltage. The high-k metal-oxide uses low voltage, but doesn't have good stability because of a high number of defects on the interface.

So, Kippelen and his team wondered what would happen if they combined the two substances in a bilayer. Would the drawbacks cancel each other out? "When we started to do the test experiments, the results were stunning. We were expecting good stability, but not to the point of having no degradation in mobility for more than a year," said Kippelen.

The team performed a battery of tests to see just how stable the bilayer was. They cycled the transistors 20,000 times. There was no degradation.

They tested it under a continuous bias stress where they ran the highest possible current through it. There was no degradation. They even stuck it in a plasma chamber for five minutes. There was still no degradation. The transistor conducts current and runs at a voltage comparable to amorphous silicon, the current industry standard used on glass substrates, but can be manufactured at temperatures below 150°C, in line with the capabilities of plastic substrates. ■

Apple eyes mobile payments over NFC

By Rick Merritt

MOBILE NFC

APPLE INC is in discussions with retailers and contract manufacturers about supporting mobile payments on future iOS devices such as the iPhone5, according to one analyst. The company is one of several expected to deliver secure transaction services on mobile devices in 2011.

"We don't know Apple's plans, but engineers we have met are scaling up to support near-field communications, and have Apple's blessing but not a purchase order yet," said Richard Doherty, principal of consulting form Envisioneering (Seaford, N.Y.). "Apple also has been in dialog with its existing and new retailers" about a new mobile payment service, he added.

Doherty said Apple could decide to ship iOS products such as iPhones or iPads as early as this year using NFC to complete secure retail transactions using an extension of its iTunes service.

"Internally at Envisioneering, we are calling it iCash," said Doherty.

Apple could significantly lower the costs credit card companies charge retailers to verify and complete transaction, a major source of irritation for retailers. Such a service "could be a game changer," he said.

"Tens of billions of dollars that flow through Apple in the next several years" if it successfully launches such a service, he added. Whether Apple does go ahead with a 2011 launch is unclear. "Anyone who says they have knowledge of Apple plans is a liar because Tim Cook and Steve Jobs won't decide until a few days or weeks before a product roll out," said Doherty.

Apple has a unique opportunity to popularize mobile payments as it has popularized mobile Web access and touch-screen technology. That's because it can control of some of the key elements such as handsets, the operating system and an online payment service needed to deliver mobile payments.

Word emerged last year of Apple patents that reference implementations of NFC. In addition, Apple reportedly hired an NFC specialist last year.

But Apple is far from alone. At the Web 2.0 conference last year,Google's Eric Schmidt showed NFC on the company's next-generation handset, suggesting mobile payments as one of its uses, and the chief executive of Research in Motion said RIM will support NFC.

Others are lining up partnerships and laying the technical plumbing needed to pilot secure mobile payment systems starting in 2011. A story in the current issue of *EE Times Confidential* described the history behind the emerging services and the elements players will need to ride this new wave of the mobile market. ■

INTELLECTUAL PROPERTY

ARM plays catch-up in graphics IP

By Peter Clarke

YOU DON'T NEED to tell IP pioneer ARM Holdings that success in the core licensing business requires best-in-class engineering, patience and the prescience to aim your technology at a sweet spot five years down the road. The result of building up a solid roster of licensees with favorable royalty terms ought to be a high- margin, low-cost business with plenty of momentum.

In the graphics IP space, that has been the payoff for multimedia IP powerhouse Imagination Technologies Group, which has seen chip volumes for its PowerVR cores take off in mobile devices and auto applications. And it is the challenge now facing ARM and its Mali core.

As Moore's Law drives chip transistor counts higher, discrete graphics processors are losing sockets to embedded implementations and heterogeneous, compute/graphics processing arrays. Meanwhile, graphics applications are mushrooming in consumer, computing, industrial and automotive electronics.

It's therefore no surprise that processor companies have targeted graphics. AMD acquired ATI Technologies in 2006 for $5.4 billion; Qualcomm then bought ATI's mobile graphics division from AMD for $65 million. And ARM (Cambridge, U.K.) created a media processing unit in June 2006 by purchasing Falanx Microsystems. Despite all the jockeying, graphics IP still looks like a one-horse race. Estimates vary, but Imagination's graphics core shipments are thought to approach 250 million chips a year. ■

Printable solar cell technology manufactured from cheap, non-toxic materials

SOLAR ENERGY

By Julien Happich

OXFORD PHOTOVOLTAICS, a company recently spun out from the University of Oxford by Isis Innovation Ltd., has developed a new solar cell technology that is manufactured from cheap, abundant, non-toxic and non-corrosive materials and can be scaled to any volume.

Harnessing the sun's energy, the solar cells are printed onto glass or other surfaces, are available in a range of colours and could be ideal for new buildings where solar cells are incorporated into glazing panels and walls.

Isis Innovation is Oxford's technology transfer company, responsible for creating new technology companies based upon Oxford research. By combining earlier research on artificial photosynthetic electrochemical solar cells and semiconducting plastics Oxford PV can now create manufacturable solid-state dye sensitized solar cells. The device is a form of thin film solar technology, a relatively new development in solar energy generation.

Leading thin film technologies are currently hampered by the scarcity of minerals used. Other dye-sensitized solar cells are being held back by the volatile nature of liquid electrolytes. Oxford PV's technology replaces the liquid electrolyte with a solid organic semiconductor, enabling entire solar modules to be screen printed onto glass or other surfaces.

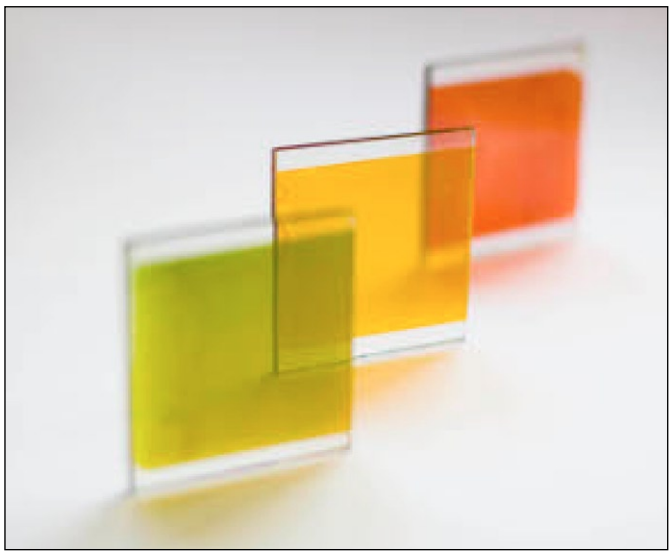

Green is the most efficient "semi-transparent" colour for producing electricity, although red and purple also work well. The materials used are plentiful, environmentally benign and very low cost.

Oxford PV predicts that manufacturing costs of its product will be around 50% less than the current lowest-cost thin film technology and expects its new mechanism will eventually match the unsubsidised cost of electricity generated from fossil fuels. The technology could revolutionise the incorporation of photovoltaic materials into windows and walls and other parts of buildings.

CEO Kevin Arthur said: 'This technology is a breakthrough in this area. We're working closely with major companies in the sector to demonstrate that we can achieve their expectations on economic and product lifetime criteria.'

The technology was developed by Dr Henry Snaith, of Oxford University's Department of Physics, who said, 'One of the great advantages is that we can process it over large areas very easily. You don't have to worry about extensive sealing and encapsulation, which is an issue for the electrolyte dye cell.' ■

Circuits and algorithms to keep your silicon cool

SPECIAL FOCUS: POWER MANAGEMENT

By Christoph Hammerschmidt

ENERGY EFFICIENCY in electronics is the paramount topic of the "Cool Silicon" research initiative that currently occupies more than 60 research institutes, universities and commercial companies in the "Silicon Saxony" region.

The recent presentation of intermediate results showed that there are some promising approaches to reduce power consumption in a broad range of applications, from base stations for mobile networks to microprocessors. One of Cool Silicon's focuses is mobile networks. In existing handset implementations as well as in the base station infrastructure, the researchers found many starting points to improve energy usage. "Energy saving circuit technologies and design methodologies have a huge economic potential," explained Frank Ellinger who leads the chair for circuit design and network theory at the Dresden Technical University. "This is because today's systems are incredibly dumb in terms of efficient use of energy."

An example that supports this thesis is the power amplifier design used in today's handsets and base stations. Power amplifiers account for more than 30 percent of the total energy consumption in mobile networks. Focusing on this alone with smart energy management systems would strongly reduce the operator's electricity bills. The same applies to the handsets. According to Ellinger, such systems transmit with full power even if the base station is close-by, although the connection would not suffer any degradation if the RF power level were reduced. "Under normal circumstances, full transmission power is required in less than one percent of the connections", said Ellinger.

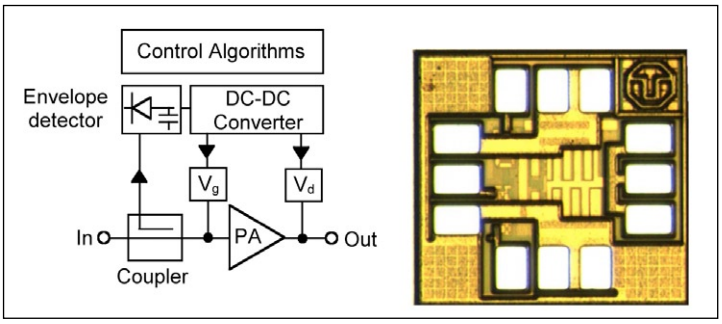

Figure 1: Controlling output power by signal level - Block diagram and silicon implementation of the power-saving PA power regulator for mobile handsets.

His researchers found a smarter way to handle the transmission power. Their system measures the signal strength and then adjusts the supply voltage and current to the PA accordingly. "The challenge is that this control process needs to be done at very high speed, within nanoseconds", Ellinger explained.

Within the scope of the research initiative, Ellinger's team has designed an integrated circuit that does this job. In its current version, the very small and cheap implementation is optimized for mobile handsets. Trials in UMTS networks showed that handsets equipped with the power control technology consumed up to 50 percent less power. Commercialization is in sight. Besides research institutions, chip maker NXP is involved in the project.

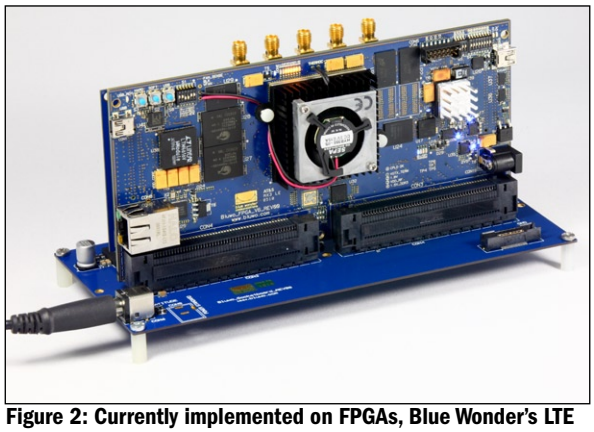

Figure 2: Currently implemented on FPGAs, Blue Wonder's LTE modem promises to be very energy efficient.

Another point of attack for efficiency improvements is the signal processing circuitry inside base stations. Typically the processing is achieved by multi-core processors where voltage and clock frequency can be adjusted to the workload. Power consumption depends not only on supply voltage and current, but is also proportional to the clock frequency. In today's microprocessors, voltage as well as clock frequency are adjusted automatically within certain limits. However, these two parameters are the same across all chip regions and functional units of the chip.

The researchers found that this scheme is much too coarse to be effective. Their multiprocessor SoC design allows fine-grained supply voltages and clock frequencies for the various cores and functional units according to their actual needs. "Today, such an optimization is rather unusual, even in standard high-power microprocessors", said Cool Silicon cluster coordinator Thomas Mikolajik. The design developed by the Dresden Technical University provides a 45/65nm CMOS MPSoC optimized for either UMTS (3G) or LTE (4G) mobile networks. The difference between 3G and 4G in terms of power consumption is that data rate and overall processing complexity for 4G are significantly higher, explained Gerhard Fettweis, professor at the Technical University's Vodafone chair. This poses much higher requirements on computing power and on-chip networking bandwidth which in turn have a direct influence on power requirements. The fine-grained power management system promises a massive reduction in power consumption, believes Fettweis.

In the realm of mobile communications, Blue Wonder Communications is also looking into energy efficiency. The company, which has become a subsidiary of chipmaker Intel since former parent company Infineon sold its mobile communications business unit to the microprocessor vendor from Santa Clara, has developed an LTE

Will your terminal pass PCI-PTS 3.0?

Yes, with the help of Maxim's PCI-PTS reference design

Minimize your BOM cost

- Meshless design with no cover required for CPU and external memories
- End-to-end solution featuring Maxim's high-performance analog and mixed-signal products

Achieve PCI-PTS 3.0 certification faster

- Precertified, secure hardware platform is easy to adapt
- Maxim's secure Linux® OS, plus precertified EMV™ and cryptographic libraries

Benefit from the industry's most secure 32-bit microcontrollers

- Patented on-the-fly encryption of external memories
- Integrated secure NV SRAM with instant zeroization technology
- Intrusion detection and environmental tamper sensors

Visit us at CARTES 2011 in Hong Kong!
Go to www.maxim-ic.com/CARTES2011 for more information.

Linux is a registered trademark of Linus Torvalds.
EMV is a trademark owned by EMVCo LLC.

MAXIM DIRECT™
www.maxim-ic.com/shop

AVNET Memec
www.avnet-memec.eu

SILICA™ An Avnet Company
www.silica.com

MAXIM INNOVATION DELIVERED®

For free samples or technical support, visit our website or call +3531 2235500.

modem featuring advanced power management. The modem is currently implemented as an FPGA-based demonstrator called BlueGate whose hardware is only activated during the active data transfers. Typically, these transfers take only a few milliseconds in uplink as well as in downlink. In addition, sophisticated algorithms have reduced the data transfer time, leading to even shorter periods during which the modem is active and draws power. On a side note, these algorithms support the current LTE standard 3GPP release 8 as well as both LTE technologies globally in use LTE-FDD and TD-LTE. According to Blue Wonder, this is a unique feature in today's LTE landscape. ■

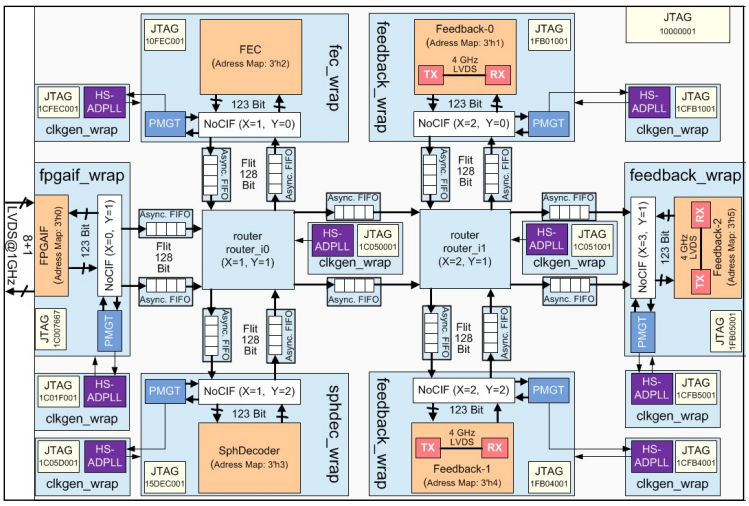

Figure 3: In complex circuits such as this multicore processor for LTE base stations, energy consumption needs to be applied in a much more differentiated way than in today's mainstream microprocessors.

An RTL to GDSII approach for low power design

By Aveek Sarkar

THE STATISTICS on infrastructure and computing needs for the Internet are sobering. The volume of data that is uploaded onto YouTube every minute of the day exceeds 35 hours of video. Facebook, if it were a nation, would be the third most populous after China and India.

Amazon.com withstands onslaughts from hordes of Black Friday online shoppers by having a massive compute infrastructure that sells practically everything. All these services are enabled through the creation and maintenance of large data centers that cater to the increasing amount of content and traffic on the Internet that is driven by the proliferation of broadband and mobile handset connectivity. If you consider the Power Usage Effectiveness (PUE) metric used to judge the efficiency of servers and cooling systems, the current ratio prevalent in the industry is 1.9, this means that almost half the power in a data center is used to cool servers with only the remaining half available for the computations themselves.

As the compute power of handheld devices approaches that of traditional computers through the use of GHz+ processor cores and increased levels of functionality, the power signature of such devices must be carefully controlled to not only make the products commercially viable, but to make them competitive in the market. If the power density of a handheld device exceeds a narrow band, it will be too hot and must be re-designed because a human hand will not be able to hold the device for any period of time.

So, in order to be competitive, IC designs that power data centers and mobile handheld devices not only need to control and reduce their overall power consumption, but also must excel in other metrics such as performance per watt or performance per Gbps. As consumers become more aware of the power their electronic systems consume, whether based on environmental concerns, cost sensitivity or application needs, system and component designers must re-think their design goals and methodologies. Depending on the application (server, networking, 3D graphics, handheld device, etc.), the specific area of power consumption to be addressed for a particular design can be different. For example, in a processor or SoC for a handheld device, very low levels of operational and standby power targets are critical design criteria.

Several techniques and methodologies have emerged to target low power design needs, however they do not contribute to reducing the operational or standby power numbers. They either have a marginal impact, or come too late in the design cycle. This approach towards power is usually ad hoc and is superseded by timing, area and routing considerations. To meet the challenges outlined earlier, a new design methodology that considers power as a design target should be adopted and followed through the entire design chain; starting from the micro-architecture definition and continuing through the RTL design period all the way to physical implementation and sign-off. By employing this process early, you can identify areas of opportunity for power reduction and benefit from the freedom to implement circuit changes that will help in power reduction.

Additionally, power consumed by the chip (at block or full-chip level), should be tracked through the entire design process to ensure convergence on the power goals set for the design. However, to be successful this methodology must go hand-in-hand with other design targets including performance, area, noise and timing. So the design changes proposed by the low power methodology should be simulated and reviewed within the context of these other important design requirements. A comprehensive holistic approach to low power design that analyzes power throughout the entire process; identifying and implementing changes in the circuit that realize power reduction with consideration for other design goals, is a "design for power" methodology.

Using a design for power methodology helps achieve the following:

(a) It ensures that power is considered

Aveek Sarkar is vice president of product engineering and support at Apache Design Solutions Inc - www.apache-da.com he can be reached at aveek@apache-da.com

from the beginning of the design, starting with architectural considerations.

(b) Identifies any circuit changes necessary to achieve or exceed power reduction goals.

(c) Isolates "power bugs" that can cause excessive current draw, but would otherwise pass all other checks (i.e., formal verification, etc.).

(d) Verifies the suggested reduction techniques do not adversely affect other design targets such as area, timing and noise.

(e) Tracks power throughout the design process; through "power regressions" that help immediately flag any design changes that may be functionally correct, but that can cause power increases.

Once the power target goals are achieved in the RTL stage, the logic and circuit changes and "power bugs" are identified and addressed, then extensive verification must be done. This will ensure that the low power techniques introduced perform and deliver the expected savings, but also that they do not adversely impact the design.

For example, the amount of predicted power savings that can be achieved through clock gating may not be realized due to the introduction of clock buffering (during the synthesis phase), that is needed to meet timing and clock skew constraints. The margin available to prevent accidental device turn-on continuously shrinks as designers use much lower supply voltages to reduce the power consumption and heat dissipation. Increased levels of noise in the power and ground networks from circuit changes such as clock and power gating that are used to achieve low power targets, impacts an already reduced margin. Referring back to the example of using clock gating to reduce operation power, increased levels of clock buffering can cause high dynamic voltage drop in the design if the placement of these clock buffers is not done optimally.

Power gating is often used for standby power reduction. However, if the implementation of power gating is not properly done, it can affect the performance and functionality of the design, and waste more power. For example, if the power gate is sized improperly, it can introduce excessive drop through the power gate (when sized smaller than needed), or it can cause excessive leakage in the standby mode (if sized larger than needed). Additionally, if the turn-on sequence of the power gates is not well controlled, it can cause a high "rush-current" in the device that generates noise through coupling in other parts of the chip.

Accurate estimation and prediction of voltage drop in the chip is very important for low power designs. So if the system in which the chip operates is not considered in the simulation through incorporating the package and PCB parasitics, the on-die voltage drop analysis is incomplete.

The design, optimization and verification of a low power chip must be done in conjunction with the design and optimization of the package that the chip will reside in, and the PCB that the packaged chip will go on. This can be achieved by providing accurate models of the package and board to the chip team, and by providing an accurate model of the chip to the package and board team. These models must capture all relevant electrical parameters (e.g. all switching current and parasitic information), and must be silicon validated for accuracy.

So, as you look to reduce operational and/or standby power, re-think your design methodology first. Establish power as a design target, starting from the micro-architecture and RTL design process. By leveraging an analysis driven optimization approach, you can explore different power saving modes and not be restricted to only one methodology.

Using RTL based power estimation numbers that are available on-die, you can initiate power grid and package design planning and prototyping.

Once RTL optimizations are done and a synthesized netlist is available, layout based power integrity analyses must be performed at the full-chip level, along with the package and PCB models, to quantify the success of the RTL stage optimizations and ensure that voltage drop in the chip is contained. In parallel, the package/PCB must be optimized considering the impact of the die to ensure the power, thermal and signal integrity of the system.

In conclusion, the successful design and delivery of a low power chip requires a comprehensive design for power methodology that impacts not only the design of the IC but also of the entire system to meet the needs and demands of a power-conscious society. ■

Power factor corrected LED driver IC

delivers flicker-free dimming for industrial lighting

iWatt has expanded its family of digital LED driver ICs with its first power-factor-corrected (PFC), primary-side-regulated, leading- and trailing-edge-dimmable, AC/DC Digital PWM Controller targeting 120 V/230 VAC offline commercial and industrial LED lighting applications. The iW3614 (3 to 15W) is designed for LED drivers used in space-constrained incandescent replacement lamps, offers hot-plug support and flicker-free compatibility with existing wall dimmers, and incorporates features to assure power efficiency, durability, and reduced size and component costs. The device features a two-stage PFC scheme, on a single chip. A front-stage boost converter enables a power factor over 0.90 (typically 0.94), while the following PWM driver stage minimizes ripple current at the output. iWatt's digital control, primary side control, quasi-resonant switching and high switching frequency combine to reduce thermal loss, component count, driver size and (BOM).

iWatt

www.iwatt.com

150W AC-DC power supply

in a compact 101.6x50.8x32mm form factor

XP Power is launching the ECP150 series of efficient and compact open frame 150W AC-DC power supplies. Measuring 101.6x50.8x32mm, the units deliver a power density of 14.88W per cubic inch. The series comply with UL/EN/IEC60950-1 and UL/EN/IEC60601-1 safety standards for medical and IT equipment. The units have a no load input power of less than 0.5W and a typical efficiency of 91%, helping the end equipment comply with internationally recognized energy efficiency initiatives. Operating from a universal input of 90-264 VAC, there are five single output models available within the range, covering the nominal outputs of +12, +15, +24, +28 or +48 VDC. A 12V / 0.5A output is also provided to power an external fan if required. Up to 100W output is achieved by convection cooling alone, and with only 15 CFM of forced cooling the full 150W output is available. The ECP150 operates between -20 and +70°C.

XP Power

www.xppower.com

Telecoms power management IC

for Icera's E400/E450 data card platform

Maxim Integrated Products has introduced power-management ICs (PMICs) for LTE/WCDMA/GSM/GPRS/EDGE data cards based on Icera's E400/E450 platform. Delivered in a 3.75x3.20mm, 42-bump WLP, the MAX8982A/MAX8982X devices integrate four high-efficiency step-down converters and nine low-dropout linear regulators (LDOs) to power all RF and baseband circuitry.

The integration completely eliminates external components, including the external converter normally required for GSM power amplifier (PA) modules. They thus reduce solution size by 60% and BOM cost by over $1.50, making them suitable for USB dongles (MAX8982A) and PC cards (MAX8982X).

To increase design flexibility and reduce system power consumption, the MAX8982A/MAX8982X allow system designers to program individual output voltages and enable/disable all regulators using the I²C interface or the dedicated control input (PWR_REQ). These PMICs also feature three current regulators with eight dimming current options (up to 24 mA) and an embedded flash timer. All LDOs include programmable voltage options and offer greater than 60 dB PSRR, less than 45 microvolts of output noise, and minimal cross-coupling noise.

Maxim Integrated Products

www.maxim-ic.com

Sub-microamp quiescent current LDOs

extend battery life in wireless devices

Analog Devices' ADP16x series of low-dropout regulators address the size, power dissipation and power-supply-rejection requirements of wireless mobile modules. The ADP160, ADP161, ADP162 and ADP163 extend the battery life of portable devices by drawing only 560 nA (typical) at no load and 42 µA of quiescent current (typical) at full load.

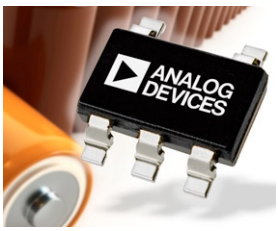

The LDOs provide 70 dB PSRR (power-supply-rejection ratio) at 60 Hz, which is more than twice that of competing LDOs at full load current. The new ADP16x series LDOs are designed for applications such as power meter reader/data terminals, portable industrial, medical measurement devices and remotely located equipment operating from batteries or solar power. all four devices provide up to 150mA of output current and operate from 2.2 to 5.5V supplies. The ADP160 and ADP162 fixed-output voltage and ADP161 and ADP163 adjustable-output voltage regulators are available in TSOT (thin–small-outline transistor) packages, while the ADP160 and ADP162 are additionally available in WLCSP (wafer-level, chip-scale package) packages. The ADP160 and ADP161 also include an output switch to discharge the output load capacitor to zero volts when the device is turned OFF to ensure microcontrollers are in a known state for restart.

Analog Devices

www.analog.com

Step-down DC/DC converter

1.5MHz operation for 93% efficiency

ROHM Semiconductor has introduced the BP5275 series of step-down DC/DC converter modules that integrate all required external components, including input/output capacitors, into a compact, high heat dissipation package. This makes them suitable for use as general-purpose power supplies in a variety of electronic devices. Currently, multiple LDOs, switching regulators, and numerous other electrical components are essential in order to provide stable electrical power to internal circuits. However, the relatively large amount of heat generated by each component requires separate heat sinks or additional substrates to facilitate heat dissipation. Multiple tedious circuit design processes, including selection of external components based on phase compensation, FET voltage, and heat dissipation characteristics, are also necessary, increasing development time and costs.

The BP5275 series was developed using an in-house 1.5MHz frequency switching regulator IC and synchronous rectification system. It has an efficiency of 93% for 6V to 5V conversion). As a result, mounting area is reduced to 1/6th of the conventional size. The package enables direct heat dissipation from the devices to an aluminum heat sink which can be mounted to increase output current capability up to 800mA. In addition, the 3-terminal, pin-compatible configuration enables major increases in power supply efficiency without requiring complete re-design.

Rohm Semiconductor

www.rohm.com/eu

Bus converters offer 98% efficiency

deliver 9.6 or 12V from a nominal 48V input

Vicor Corporation has introduced the IBC048 series of VI BRICK Intermediate Bus Converters, which offer double the power density and half the conversion loss of competitive devices. The devices are available as drop-in replacements for industry standard eighth-brick and quarter-brick 5:1 and 4:1 intermediate bus converters, initial product offerings include 300 W and 500 W models, providing 9.6 or 12 V outputs from a nominal 48 V input with a voltage range extending from 38 to 55 V. The new VI BRICK bus converters use the same ZCS/ZVS engine that has powered V.I Chips into high performance applications with tier-one customers. IBC048 modules cut transient response time from 200 to 20 microseconds, eliminating the need for bulk capacitors. Superior density and conversion efficiency enable competitive advantages for OEMs seeking to maximize performance across a range of applications, including computing, data storage, networking and Power-over-Ethernet.

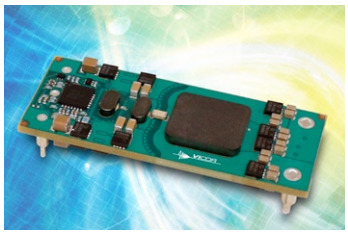

Vicor Corporation

www.vicorpower.com

THE VERY LATEST TECHNOLOGY...
WE'VE GOT IT FIRST!

- Over 100 new technologies added online every day
- The broadest product range in European distribution, over 520,000 products from 3,500 suppliers
- element14 - The first online technology portal and e-community for Design Engineers
- Now receive all of your quotes online with eQuotes
- i-Buy – FREE online eprocurement solution

www.farnell.com

element14

www.element14.com

Farnell

Design with the best

A Premier Farnell Company

Synchronous step-down DC/DC MCU

digital power system management via I²C/PMBus

Linear Technology launched a dual output high efficiency synchronous step-down DC/DC controller with I2C-based PMBus interface for digital power system management.

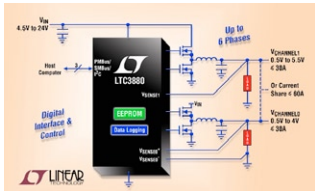

The LTC3880/-1 allows for digital programming and read back for real-time control and monitoring of critical point-of-load converter functions. Programmable control parameters include output voltage, margining and current limits, input and output supervisory limits, power-up sequencing and tracking, switching frequency and identification and traceability data. On-chip precision data converters and EEPROM allow for the capture and nonvolatile storage of regulator configuration settings and telemetry variables, including input and output voltages and currents, duty cycle, temperature and fault logging. The LTC3880/-1 can regulate two independent outputs or be configured for a two phase single output. Up to 6 phases can be interleaved and paralleled for accurate sharing among multiple ICs, minimizing input and output filtering requirements for high current and/or multiple output applications. An integrated amplifier provides true differential remote output voltage sensing, enabling high accuracy regulation, independent of board IR voltage drops.

Linear Technology Corporation

www.linear.com

24V power modules

adjustable switching frequency for flexible design

National Semiconductor's new Simple Switcher power modules can drive high output voltage applications in a variety of markets including industrial, communications infrastructure and military.

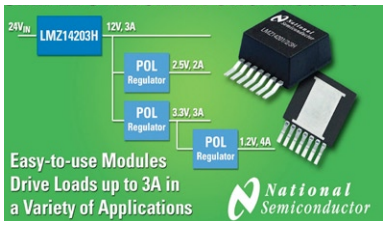

The LMZ14201H, LMZ14202H and LMZ14203H devices provide the efficiency of a synchronous switching regulator with the simplicity of a linear regulator, eliminating the external inductor and complex layout placement challenges typical of switching regulator designs. The devices accept an input voltage rail between 6 and 42V and deliver an adjustable and highly accurate output voltage up to 24V. A single resistor adjusts the switching frequency to enable greater flexibility of design. Pin-to-pin compatible with other family members, the modules integrate a shielded inductor and feature efficiency up to 97%. The integrated inductor alleviates EMI concerns as the modules comply with both CISPR 22 Class B radiated and conducted emissions standards.

The chips drive loads from 1A to 3A and are well-suited to generate intermediate bus voltages in systems using a distributed point of load (POL) architecture.

National Semiconductor

www.national.com

Isolation amplifiers for current sensing

find applications in motor drives and inverters

The ACPL-790B, ACPL-790A and ACPL-7900 optical isolation amplifiers from Avago Technologies are well suited for current and voltage sensing in AC and brushless DC motor drives, industrial inverters, servo motor drives, wind power generation, solar panel power systems and general analog isolation. As current flows through the external resistor in a motor drive implementation, the resulting analog voltage drop is sensed by an ACPL-790B/790A/7900 isolation amplifier, and it allows a proportional output voltage to safely be created on the other side of the optical isolation barrier.

The ACPL-790B precise isolation amplifier provides up to 0.5% high gain accuracy, and offer 200kHz bandwidth and 1.6 ms fast response time to enable capture of transient signals in short circuit and overload conditions. The devices operate from a single 5-V supply that is compatible with 3.3-V outputs. This performance is delivered in a compact DIP-8 package that is suitable for automated assembly. High gain tolerance options of ±0.5% (ACPL-790B), ±1% (ACPL-790A), and ±3% (ACPL-7900) are available. The isolation amplifiers are compliant to UL 1577 5000 Vrms/1-minute rating, IEC/EN/DIN EN 60747-5-5 and CSA industrial safety standards, while 15 kV/ms common-mode transient immunity ensures less torque ripple in motor control applications.

Avago Technologies

www.avagotech.com

Low profile power conversion IC

can supply up to 65W with no heat sink

Power Integrations has announced its TOPSwitch-JX power conversion IC in an innovative eSOP low-profile power package. The surface-mount package is suited for applications without heatsinks up to 65W in compact, open-frame installations such as slim LCD TV auxiliary power, set-top boxes, PC standby, and DVD players.

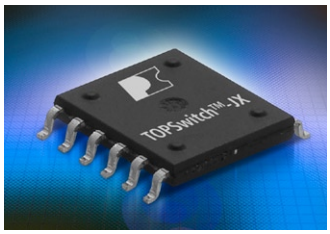

In notebook, netbook, and printer adapters, the thermally efficient eSOP package can supply up to 40W using only the PCB for thermal management. The eSOP package features an exposed die-attach pad that may be reflow soldered to the PCB during assembly. This permits the copper ground plane and thermal mass of the board to act as a heat sink. The 12-pin package has a 119mm² footprint, yet maintains safety creepage and clearance distances required by international standards. TOPSwitch-JX ICs incorporate both a PSU flyback controller and a 725 V MOSFET power in a single package, and feature a novel multi-mode control algorithm that maximizes power efficiency across the entire load range. High efficiency at full power minimizes power wasted during normal operation and reduces the complexity and expense of thermal management on the system.

Power Integrations

www.powerint.com

Integrate Touch Sensing Quickly and Easily

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

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

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

Capacitive Touch Keys and Sliders

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

Touch Screen Controllers

- Fully processed touch coordinates
- **Projected Capacitive technology**
  - Multi-touch enabling gestures
  - Low cost MCU implementation
  - Wide operating voltage: 1.8-5.5V
  - Low operating current 1.5 mA at 5V typical
- **Analog Resistive technology**
  - Lowest system cost, easy integration
  - Universal 4, 5 & 8-wire solution with on-chip calibration
  - I²C™, SPI, UART or USB interfaces
  - Low power "touch to wake-up" feature

GET STARTED IN 3 EASY STEPS

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

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

Projected Capacitive Development Kit - DM160211

Analog Resistive Touch Screen Development Kit - DV102011

Intelligent Electronics start with Microchip

microchip DIRECT
www.microchipdirect.com

www.microchip.com/mtouch

MICROCHIP

System design using NAND flash memory

By Jim Cooke

NAND FLASH error correcting code (ECC) has been on the rise since NAND was first introduced. Although it's not a new issue, the ECC required to support newer multilevel (MLC) and three-bit-per-cell technologies is becoming increasingly difficult for system designers to keep up with. ECC has historically been used to improve the overall data reliability of NAND subsystems. However, as NAND cells shrink, fewer electrons are stored per floating gate. To compensate for the increasing bit error rates of these smaller geometry cells, ECC requirements have to dramatically increase to maintain the desired system reliability.

As system requirements for ECC increase, the number of gates necessary to implement the logic also increases, as does the system complexity. For example, 24 bits of ECC requires about 200,000 gates, while 40 bits of ECC requires about 300,000 gates. It is estimated that in the future, advanced algorithms will approach close to 1 million gates - see figure 1.

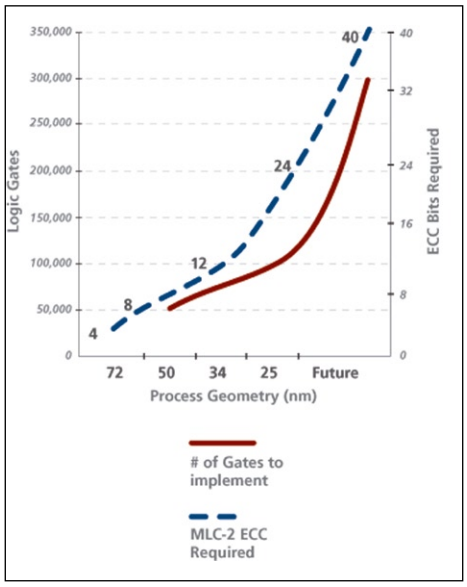

Figure 1: ECC increases as process geometries shrink.

Jim Cooke is senior technical marketing manager at Micron Technology – www.micron.com – he can be reached at jcooke@micron.com

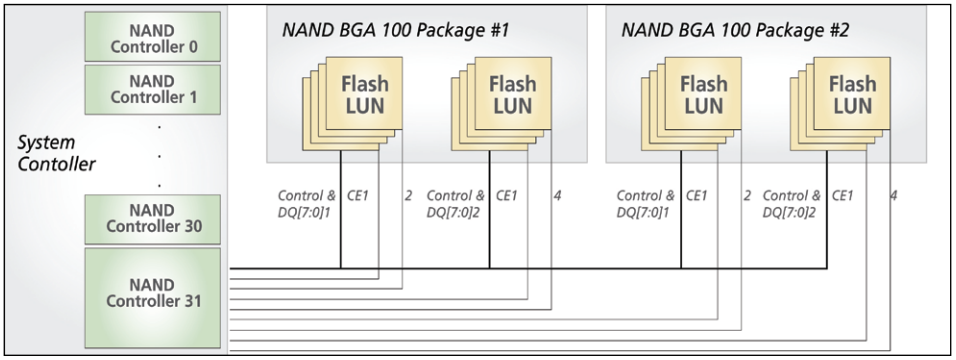

Figure 2: Typical SSD implementation using two 8-die, 100-ball BGA packages.

Many high-performance flash systems require multiple channels of NAND to reach the desired performance. In these systems, each channel typically has its own ECC logic. For example, a 10-channel SSD may have 10 channels of ECC logic implemented. If 60 bits of ECC were required for each of the 10 channels, the result would be 3 million gates just for the ECC logic.

NAND interface choices

The NAND interface has traditionally been an asynchronous interface. Although interface speeds have improved up to 50 MHz in recent years, not much else has changed on this interface. Several years back, Micron and several other forward-thinking companies joined together to form a NAND flash organization that was focused on simplifying the myriad of timing and command specifications offered by the industry. The Open NAND Flash Interface (ONFI) developed the first version of their specification, ONFI 1.0. While there are many advantages to the original ONFI 1.0 specification, one of the biggest is the ability for the host to electronically detect the type of flash device that is connected, as well as other important parameters, like timing modes, page size, block size, ECC requirements. This feature has been carried forward to all of the ONFI specifications and remains an important aspect of all ONFI standards.

Another significant accomplishment of the ONFI organization was the development of the synchronous NAND interface, also known as ONFI 2. ONFI 2.2 currently supports up to 200 mega transfers per second (200 MT/s) using a DDR, source-synchronous interface. After powering up, it can be used in asynchronous mode. However, for higher performance, the host can interrogate the flash device to see if it is able to support the higher-speed synchronous interface before changing to it.

Direct NAND solutions

Implementations that connect NAND directly to the host processor or SSD controller are responsible for managing the NAND. Hardware manages the ECC, while software typically performs all block management and wear-leveling operations. At first this may seem like a disadvantage. However, with today's typical embedded processors running at speeds of hundreds of megahertz and often over one gigahertz, these high-performance processors can accomplish block management much faster and can take advantage of deterministic, multithreading techniques to improve performance. In addition, with the host managing the Flash device directly, the host software can make real-time decisions that can help eliminate exposure to unexpected power failures.

As shown in figure 2, the ONFI 2.2 specification (200 MT/s) was designed to accommodate up to 16 standard NAND loads. A typical implementation of this would be using two 8-die NAND packages. The standard 8-die, 100-ball BGA package includes two separate NAND buses (DQ[7:0]1 and DQ[7:0]2), with each bus having four NAND flash wired together. Each of the four die stacks are

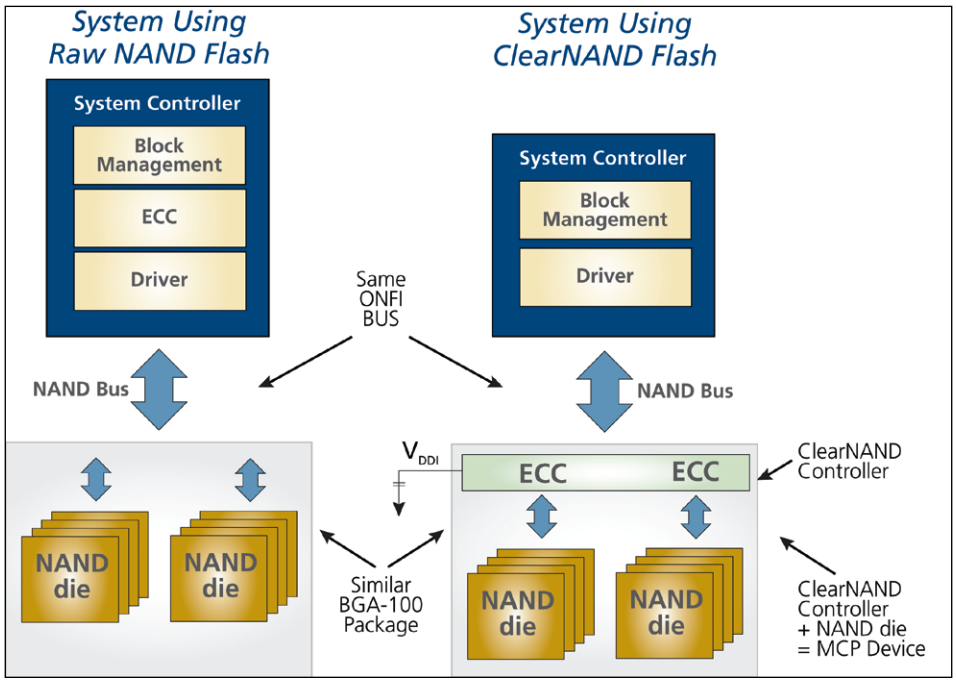

Figure 3: Standard NAND vs. ClearNAND.

controlled with two chip enables. A typical design would wire the two data or DQ buses together, forming a single 8-bit data bus for each package. A maximum configuration would consist of two 100-ball BGA packages, each containing eight die. Each of these standard 100-ball BGA packages requires four chip enables (CE#) to select a specific NAND die. Thus, the host or SSD controller needs to supply eight chip enables to support this configuration.

ClearNAND solutions

Figure 3 shows two system implementations: a traditional system where the processor or SSD controller is interfacing directly with NAND and a system using ClearNAND flash. Both implementations use the same ONFI hardware interface and similar 100-ball BGA ballouts. The ClearNAND example includes a thin controller package with the NAND die in an MCP. The ClearNAND controller implements the ECC required by the NAND inside the MCP package. Using the same ONFI asynchronous or synchronous interface allows designers to migrate easily from standard NAND to ClearNAND Flash.

Micron offers two versions of ClearNAND flash: standard and enhanced. Standard ClearNAND flash, suggested mainly for consumer devices, implements the required ECC and provides a traditional asynchronous ONFI bus for easy migration.

Enhanced ClearNAND flash manages ECC, in addition to offering several performance-enabling features that are of most value to enterprise applications. It also supports both the asynchronous and synchronous versions of the ONFI 2.2 interface and is available in densities up to 64GB.

By abstracting the ECC, both versions of ClearNAND flash will be able to handle the additional ECC that future versions of NAND will require. This will eliminate the need for designers to continually redesign their circuitry to keep up with manufacturers' latest NAND ECC requirements.

Enhanced ClearNAND flash

Figure 4 shows the Enhanced ClearNAND architecture. It supports a single ONFI 2.2 interface and up to 200 MT/s command, address and data bus. The VDDI decoupling capacitor is common in e·MMC products and other devices that include a controller. It is required to decouple the internal voltage regulator. For backward compatibility with standard NAND devices, the VDDI connection is located on an unused pin. The ClearNAND controller supports two internal flash buses, one for even logical unit number or LUNs and one for odd LUNs. These two independent flash buses can operate at up to 200 MT/s. In addition, each bus has its own ECC engine and can manage simultaneous READ or WRITE operations on the two buses. It is envisioned that future versions of the controller will support the ONFI 3 specification, which is targeting up to 400 MT/s. Many of these basic features are covered in the following discussion of the four advanced functions offered by enhanced ClearNAND: volume addressing, electronic data mirroring, interrupt function and internal copyback.

Volume addressing

Volume addressing allows a single chip select or chip enable (CE#) to address up to 16 ClearNAND volumes. Each ClearNAND controller can support up to eight die packaged in an MCP. The ClearNAND controller provides a buffer for the host or SSD controller access. The enhanced ClearNAND design, shown in figure 5, offers an eightfold improvement in density while maintaining or improving signal integrity and reducing the active number of chip enables that are required. This is because a single ClearNAND controller represents only one load to the SSD controller, but supports up to eight NAND die in the MCP package.

There are two aspects to the volume addressing concept. The first is establishing the volume address for each of the ClearNAND

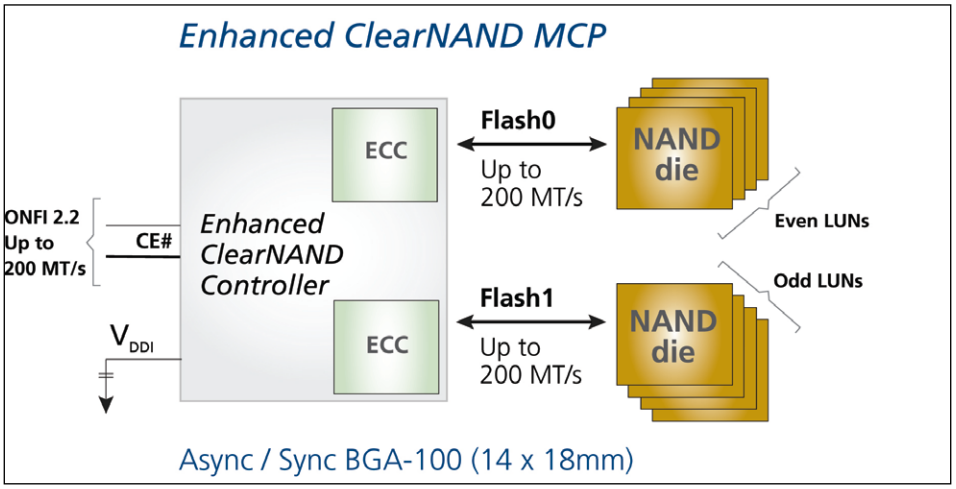

Figure 4: Enhanced ClearNAND architecture.

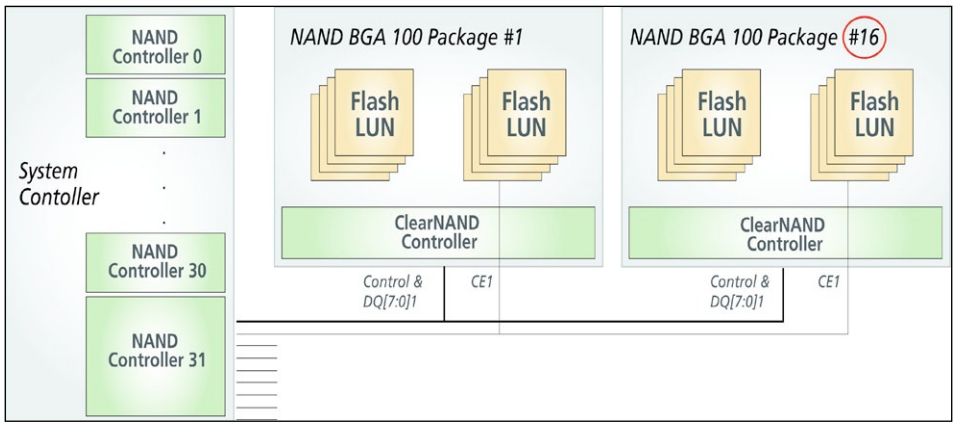

Figure 5: Typical SSD implementation using up to eight dies in a 100-ball BGA ClearNAND packages.

packages. The volume address is appointed only once at initialization and is maintained until the power is cycled. The second aspect is the volume select command itself. This is a new command that is followed by a single byte (actually only 4 bits) volume address. Once the intended volume is selected, it remains selected until another volume is selected. The chip enable pin savings can be significant. For example, a 32-channel SSD requires eight chip enables to control two 8-die standard NAND packages. The 32-channel example would require a total of 256 chip enable pins whereas the enhanced ClearNAND volume addressing feature can address the same amount of NAND using only 32 chip enable pins. What's more, these same 32 chip enable pins can be used to address eight times the density.

Electronic data mirroring

Enhanced ClearNAND supports electronic data mirroring, which allows the data bus signal order to be electronically remapped to one of two configurations. This is particularly useful for high-density designs with ClearNAND devices mounted on both the top and bottom of the PCB. The ClearNAND package is able to electronically detect if it is mounted on the top or bottom of the PCB. This is accomplished using a specific initialization or reset sequence. For example, it is common practice to issue a reset or FFh command to the flash device on power-up. To accomplish the electronic DQ mirroring, the host must follow this FFh command with the traditional READ status (70h) command. The top die detects this command sequence as FFh-70h; the bottom die recognizes this same sequence as FFh-0Eh and can establish that it is the bottom package and reorder its data bus to align directly under the top die. This not only improves PCB routing but also improves signal integrity.

Ready/Busy# redefined to interrupt

Enhanced ClearNAND flash redefines the existing ready/busy# pin to be an interrupt pin. As shown in figure 6, the interrupt# signal, which is still open-drain, provides a real-time interrupt when the ClearNAND volume or die becomes ready. Designers can use this interrupt signal to provide real-time status to the host or SSD controller. In larger configurations supporting multiple ClearNAND packages on a single bus, the interrupt# signals can be wired together. When the host or SSD controller detects an interrupt, it can simply interrogate each of the ClearNAND packages or volumes to learn which volume has posted the new status. This interrupt function can save signals on the host or SSD controller while improving the ability for the SSD controller to respond to status updates.

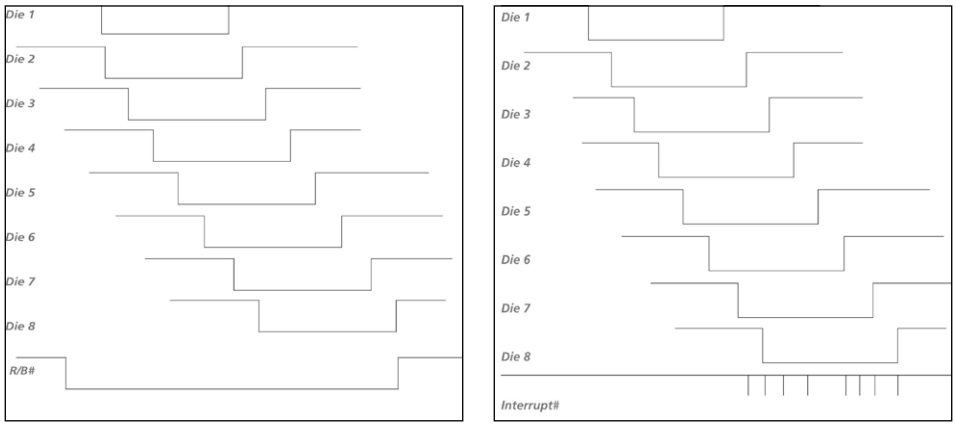

Figure 6: Enhanced ClearNAND flash redefines the existing ready/busy# pin (6a) to be an interrupt pin (6b).

Internal copyback

The last but perhaps most noteworthy feature of enhanced ClearNAND flash is the internal Copyback function, also known as internal data move. This function can provide a significant advantage in SSD systems when it comes to wear leveling or garbage collection operations; that is, the process of collecting fragments of data scattered throughout the various pages and blocks of the NAND and coalescing them in a more streamlined block or sequence of blocks. It is similar to the old hard disk defragmentation utility. Referring back to figure 2, when using standard NAND, moving data fragments from one block to another typically requires the following sequence of operations:

The SSD controller issues a READ command and source address to access the source page of data. The SSD controller inputs data from the NAND device while calculating and making any necessary ECC corrections. Any updating of data or metadata is usually accomplished after this step. The SSD controller calculates and appends the new ECC information before issuing a new PROGRAM command, destination address, and data sequence, which will store the data in the new NAND block.

In this sequential operation, the bus is busy while moving the source and destination data from and to the flash device. The timing of these operations can be significant. An ONFI 2.2 synchronous bus operating at 200 MT/s would require about 41²s to move the data, assuming an 8K page. Since the data has to be moved from and to the flash device, we double this time to 82²s, which doesn't include the ECC overhead. While this sequence is being carried out, the ONFI Flash bus is busy and cannot be used for other operations.

Enhanced ClearNAND flash is different in that it supports internal ECC. Using this built-in ECC allows the Copyback operation to be performed internal to the enhanced ClearNAND package, assuming the source

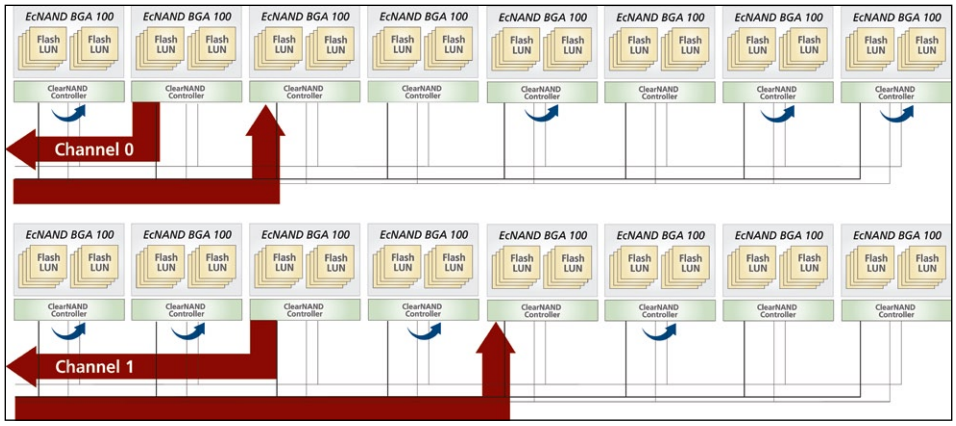

Figure 7: Copyback using enhanced ClearNAND flash

and destination of the data are within the ClearNAND package. The SSD controller is still responsible for issuing the commands and addresses as well as any modified data or metadata. The data movement is handled by the ClearNAND controller and does not tie up the external ONFI bus. If the SSD controller is able to keep its wear leveling and garbage collection operations within a single ClearNAND package, it can have significant performance advantages.

Figure 7 shows an example using enhanced ClearNAND on two ONFI channels labeled channel 0 and channel 1. On both SSD channels, we can see that four of the internal data move operations are occurring simultaneously without the external ONFI bus being used for the data movement. This frees the SSD controller and ONFI bus to move data from one ClearNAND package to another, if necessary. Depending on your architecture, some percentage of these operations may need to go between ClearNAND packages or even between ONFI buses. Taking advantage of the internal data move operation can provide a significant performance improvement for garbage collection and wear-leveling operations.

Conclusion

Micron's Enhanced ClearNAND Flash provides additional performance and features while eliminating the impact of NAND's ever-increasing ECC requirements. Because enhanced ClearNAND supports a ballout similar to the standard 100-ball BGA NAND devices, it's possible to design your product to support both. An example would be to include enough ECC in your host SSD controller to support SLC NAND flash directly and select enhanced ClearNAND flash for your multilevel cell needs, where ECC can present more of a challenge.

The volume addressing feature of enhanced ClearNAND enables higher densities to be addressed using fewer chips, saving potentially hundreds of pins in SSD implementation. Electronic data mirroring simplifies PCB design and routing while improving the signal integrity of the ONFI bus. The intelligent interrupt function provides for real-time status updates to the SSD controller and minimizes the polling for firmware. The dual internal NAND flash buses provide improved Copyback operations, which in turn improve performance. ■

USB hub and card reader applications in the automobile

By Henry Muyshondt

USB AND FLASH MEMORY cards have become ubiquitous in the consumer and industrial worlds. USB implements a high speed serial bus that runs at up to 480 Mbps. Many operating systems provide native support for this technology with many hundreds of millions of devices shipped to date.

But USB is not only used to transfer data between devices, it also provides a means to charge portable devices. As consumers expand their digital lifestyle to have their content always available, more and more devices take advantage of the economies of scale afforded by the explosion of interconnections that ensue. Car makers are embracing this trend as their vehicles integrate into the digital world.

Vehicles are also becoming storehouses of content and information. They can include large amounts of storage capacity for entertainment content and navigation information. One of the most popular memory formats today is Secure Digital (SD). The SD interface is also used in embedded applications to attach devices like WiFi (or WLAN), Bluetooth transceivers, and GPS receivers with an SDIO interface. SD memory can be used to replace rotating media like hard disks, CD and DVD media. A state of the art 32-GB card holds the equivalent of close to 7 DVDs! Car makers can use memory cards both as a connection to consumers and as a mechanism to upgrade different systems within the vehicle, be they navigation systems or any other devices that require software.

The USB and Flash media interfaces are thus very useful in automotive applications. For this purpose, SMSC recently introduced the TrueAuto Quality USB82640 and USB82660 USB hub and card reader combination products, specifically designed for the stringent requirements of the automotive environment.

Henry Muyshondt is director of business development at SMSC - www.smsc.com

Automotive quality requirements

Before getting into the specific functions of the interfaces, let's first consider automotive quality requirements. Devices intended for the automotive market have to be designed, validated, characterized, qualified, fabricated and supported specifically for use in automotive applications. Cars have very long lifecycles and any failure in the field is very costly in terms of repair time and customer satisfaction.

TrueAuto is SMSC's proven automotive quality process. When integrated circuits initially designed for consumer applications, are used in automotive applications, they have to be qualified according to the Automotive Electronics Council's qualification requirements (AEC-Q100). This standard, however, only covers minimum common requirements for the qualification of an automotive IC. Many car companies and tier one automotive suppliers require extensive additional qualification tests, as AEC-Q100 alone does not lead to the ultra-low defect rates that they require. In addition, AEC-Q100 primarily focuses on the qualification phase of the product cycle of an IC. Other phases such as the design and production of the IC, customer support and the handling and investigation of returns are not covered in detail. In order to reach the automotive goal of near zero defect rates, all phases of the IC product cycle need to be addressed thoroughly. TrueAuto robustness begins with SMSC's design for reliability techniques within the silicon IC itself. Automotive-grade excellence and testability are designed into the IC. The IC is fully characterized over many operating parameters to prove the quality of the design under various conditions. What's more, product qualification is focused on the most demanding customer expectations and meets or exceeds automotive reliability standards and customer specific requirements.

Memory for storage

Portable memory cards are used by passengers to transfer information created on computers, portable media players or cameras into the car. Car makers also incorporate gigabytes of microcode into some of today's most sophisticated vehicles. Storage is also required for navigation systems' map data. Maps for a large country, like the United States, can fit in under 2GB of memory. An SD card of this size can be purchased at retail for less than $5.00, making it very cost effective compared to the typical DVD player used in many automotive navigation systems. In addition, reliability is increased as there are no moving parts associated with it. The high speed data transfer enabled by an SD interface can simplify software updates for other components in the car, like a head unit. Yet these in-box use cases must provide data access with true automotive-grade reliability, whether the memory devices connect to internal peripherals or provide external consumer access.

Combining a USB hub with a reader

SMSC's USB82640 chip provides two downstream USB 2.0 ports plus a single Secure Digital (SD)/MultimediaCard (MMC) or Sony MemoryStick memory card interface. The US882660 provides a second SD card interface and is capable of operating up to an ambient temperature of 105°C. The second memory card interface enables navigation data and entertainment content to be accessed simultaneously when stored in popular memory card formats. The SDIO

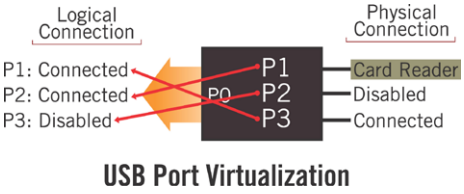

Figure 1: how PortMap can simplify the system architecture.

interface can be used to attach modules that provide additional features such as WiFi, Bluetooth and GPS connections. It is even possible to build custom firmware to control new applications attached through SDIO.

When designing USB systems, the USB hub controller portion of the device provides advanced features that simplify system design. This includes PortMap (figure 1) which provides flexible port mapping and disable sequences. The downstream ports can be reordered or disabled in any sequence to support multiple platform designs. When a port is disabled, the others are automatically reordered to match the USB host controller's port numbering scheme. This enables car makers to offer different feature sets without requiring a new design for each option.

Another advanced feature, PortSwap adds per port programmability to USB differential-pair pin locations. High-speed connections like the 480 Mbps data lines (D+ and D-) of USB 2.0 require that the signal paths for each data line be closely matched. This is done to prevent signal degradation and electromagnetic compatibility issues. PortSwap

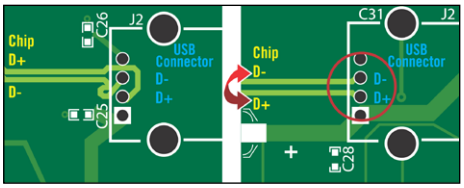

Figure 2: PortSwap prevents signal degradation and electromagnetic compatibility issues by swapping the data lines on a board layout.

(figure 2) allows the data lines to be swapped when the location of the connector pins is such that it would require different routings for each line or crossing of the lines.

SMSC's PHYBoost enables four programmable levels of USB drive strength in downstream port transceivers. PHYBoost attempts to restore signal integrity in complex system environments. Sometimes long cables or other system challenges result in a compromise of the eye diagram of the USB signal. Additional drive strength can help restore the eye diagram as can be seen in figure 3.

USB card reader

The card reader portion of the device includes an 8051 controller that can transfer data as fast as 35MB/s. This is faster than many host controllers can process information and faster than many SD cards can be read. The high performance of the card reader seeks to avoid a bottleneck when moving data between a memory card and a host device. The specifications for the SD interface allow some room for interpretation and optional features that can result in incompatibilities with different products. Hence SMSC has done extensive testing to support a large number of cards currently in the market. The current devices also support the use of an external ROM to create secure memory formats or add customized applications based on system requirements. Incompatibilities with cards from different manufacturers could result in warranty claims against a car maker. A service call is very expensive to them so it is important to avoid such calls even if a consumer brings in a device that he got for free somewhere.

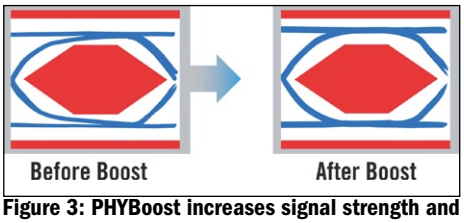

Figure 3: PHYBoost increases signal strength and help restore the eye diagram.

The combination hub/card reader function allows the placement of this device away from the main host controller to provide connectivity where it is needed. For example, the glove compartment or center console in a car could allow consumers to easily connect their devices without requiring long cables to the main head unit. ■

1GB and 2GB DDR3 memory modules

in rugged small outline unbuffered ECC DIMM

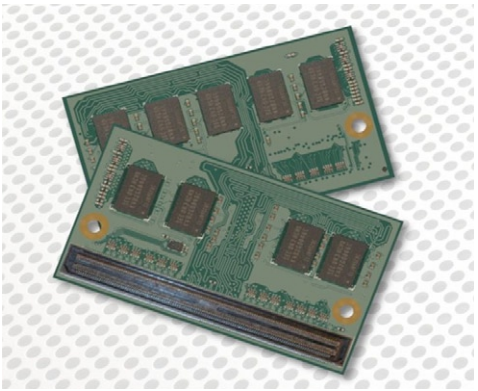

Swissbit's launched what the manufacturer claims to be the world's first DDR3 RS-UDIMM (Rugged Small Outline unbuffered ECC DIMM) in a 38x67.5x5mm form factor, for industrial and embedded system operating in tough environmental conditions. The DDR3 240-pin 72-bit RS-UDIMMs are electrically compatible to the JEDEC DDR3 72bit SODIMM and target the latest Intel based embedded platforms as well as other DDR3 ECC capable applications for speed grades of PC3-8500 and PC3-6400. The module connection between RS-UDIMM and system board is implemented by a 240-pin mezzanine connector, specially designed for use in rugged environments. Mounting holes and optional fixtures secure the module against shock and vibration. The 72bit wide module data interface provides ECC support to aid in the correction of single bit errors potentially caused by electrical or magnetic interference and/or background radiation. The RS-UDIMM is available in densities of 1GB and 2GB, both in commercial temperature grade 0 to +85°C (TCase) and industrial temperature grade (-40 to +95°C).

Swissbit

www.swissbit.com

Universal flash IP core

for non-volatile storage applications

Arasan Chip Systems has been developing a Universal Flash Storage (UFS) IP core, the next generation memory interface being finalized by JEDEC. The company has engaged with its strategic lead customers to help them productize this interface. The UFS standard is expected to scale in performance and features to span existing and emerging usage models for non-volatile memory.It is based on a modular layered protocol architecture. This architecture enables an efficient interface implementation with the ability to scale the interface to meet future performance requirements. The IP core solution for UFS spans software, controller IP, MIPI UniPro Link, and M-PHY IP. The UFS IP core is delivered as RTL IP for the controller, PHY IP, verification IP and a portable software stack.

Arasan

www.arasan.com

Printed rewritable memory array

128 bits combined with logic circuitry

Thin Film Electronics ASA together with PARC, a Xerox Company, have completed the design of a 128 bit Thinfilm addressable memory. The design combines Thinfilm's memory technology with transistor technology developed by PARC, and includes CMOS circuitry to significantly reduce the number of contact pads necessary to control the Thinfilm Memory. Thinfilm's current product is a contact-based 20 bit memory and memory controller for advanced interactive toys and games. Addressable Thinfilm memory products will allow integration to create fully printed systems, such as ID tags, sensor tags, and disposable price labels.

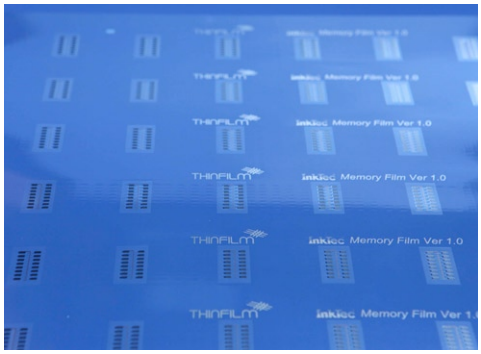

Thinfilm and PARC announced in October that they are working to provide next-generation memory tags using printed electronics, and is now developing prototypes of the addressable memory. These prototypes will be ready later this year. Transfer to production is expected in 2012. Using printing to manufacture electronic memory makes it possible to reduce the number of process steps, dramatically reduce manufacturing costs, as well as the environmental impact as compared to traditional semiconductor processes. Commercial applications of printed electronics include e-paper, electronic readers, and organic light emitting (OLED) displays. Sensors, batteries, and photovoltaic energy sources are also in development, and together with Thinfilm's memory technology they will open the door to new products and applications.

Thin Film Electronics

www.thinfilm.se

576-megabit low-latency DRAM

for network equipments

Renesas Electronics released a family of 576-megabit low-latency DRAM products (µPD48576109, µPD48576118, µPD48576209, µPD48576218, and µPD48576236) for network equipment.

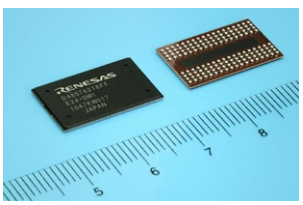

The devices offer doubled memory capacity, 25% improved random cycle performance for high-speed reading and writing of data, 33% faster operating frequency, and over 10% reduction in power consumption compared to the company's existing 288 Mb low-latency DRAM products.

The maximum operating frequency has been raised from 400MHz to 533MHz. To enable shorter development cycles, the 576Mb devices come in the same package as existing 288M low-latency DRAM products, measuring 11 x 18.5mm

Renesas Electronics

www.renesas.eu

Smart storage array

supports 5 million virtual machines

EMC Corporation announced new software capabilities for its Symmetrix VMAX storage systems, dramatically increasing performance and simplifying the way petabytes (1 petabyte = 1,000 terabytes) of information can be managed. Designed for mission-critical virtual data centers, the latest version of Symmetrix Enginuity software doubles system performance with no hardware upgrade required. It fully leverages Intel Xeon technology integrated into VMAX to deliver up to twice as many OLTP transactions and DSS queries, making the fastest storage system even faster and further increasing workload scale. The new EMC FAST VP (Fully Automated Storage Tiering with Virtual Pools) software optimizes Symmetrix VMAX for performance, utilization, and cost. Compared to single-tier systems, FAST VP delivers up to 40% more application performance at a 40% lower cost while requiring 87% fewer disks and 75% less power. The virtual server integration scales up to 5 million virtual machines on a single VMAX. In addition, new VMware API support provides 800% faster management and provisioning and 300% faster replication and mobility. Hardware-based encryption safeguards data-at-rest with no performance impact.

EMC Corporation

www.EMC.com

4Mbit asynchronous SRAM

with integrated error correction code

Integrated Silicon Solution announced the ECC based IS64WV25616EDBLL, a 4Mbit high speed asynchronous SRAM that is form-fit-function compatible with current industry standard devices without ECC.

The IS64WV25616EDBLL is organized as 256Kx16 and has an operating voltage range of 2.4V to 3.6V. The device is offered in 44-pin TSOPII and 48-Pin BGA packages. To provide enhanced reliability, the 44-pin TSOPII is a copper lead frame based product. IS64WV25616EDBLL is offered in both leaded and lead free options to address various customer requirements. The memory chip is based on 65nm technology and draws not more than 45mA operating current when operated in the range of -40 to 85°C. Typical operating current is 25mA. The device is offered with a 10ns access time, it is available for use in the -40 to 85°C and -40 to 125°C temperature ranges. The error detection scheme is based on an independent hamming code for each byte. The device detects and corrects one bit errors for each byte. This, it provides better reliability than parity code schemes which can only detect an error, but not correct it - a feature that makes it particularly attractive for many automotive applications, especially engine and transmission control. A 512Kx8 option offered in a 44-pin TSOPII will be available in April 2011 and a 1.8V option will be available in June 2011. The product is well suited for automotive, industrial, medical, Mil-Aero and telecom/networking applications, the vendor claims.

Integrated Silicon Solution Inc

www.issi.com

Testing multicore software

By Mike Bartley

HARDWARE VERIFICATION ENGINEERS have always faced the complexity of concurrent execution and temporal considerations when verifying hardware designs. However, silicon manufacturers are now moving to multicore designs (i.e. multiple CPU cores on a single chip) to achieve the relentless drive for improved performance at lower power now demanded by consumers. This pushes the responsibility for realizing that performance to the software community. Rather than simply relying on clock frequency and CPU design improvements to achieve their performance increase, software engineers will need to write their code to take advantage of the additional cores. Testing such software is fraught with potential issues and we will investigate just a few of them in this short article.

So why is testing concurrent software hard? The main issue is that concurrency introduces non-determinism so that running the same multi-core program twice is not guaranteed to

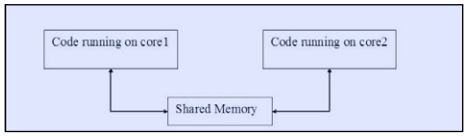

Figure 1: Threads executed two cores sharing memory.

give the same result each time due to different inter-leaving of threads of execution. Consider two pieces of code running on different CPUs sharing some common data (as depicted in figure 1).

The hardware will take care of the data coherency. That is, it will ensure that both cores have a consistent view of the shared memory. So, for example, if core1 has a locally cached memory location then the hardware will ensure that core2 knows about it when it tries to assess that location. However, that does not mean that the software behaves correctly. For example, consider the two threads that share the variable "num" in figure 2.

The two boxes demonstrate two different inter-leavings (steps) of those threads which lead to different values in "num". No hardware data coherency mechanism will protect you from this. You may feel that the thread should have been written as a "num++" but remember the underlying execution model of the machine code. The above threads will each get translated into three CPU instructions on the hardware: fetch to a register; increment that register; and then write the register value back to memory. This non-determinism can often mean that a test may fail on one run but then pass on the next - often referred to as a "heisenbug" (after the Heisenberg uncertainty principle). This makes debugging even more difficult.

```
static int num = 0;

thread1 () {
 int val = num; // step 1
 num = val + 1; // step 3
}

thread2 () {
 int val = num; // step 2
 num = val + 1; // step 4
}
```

```
static int num = 0;

thread1 () {
 int val = num; // step 1
 num = val + 1; // step 2
}

thread2 () {
 int val = num; // step 3
 num = val + 1; // step 4
}
```

Figure 2: Threads accessing shared code.

There are design and coding techniques to avoid the "race conditions" observed when running the code shown in figure 2. The main technique is to make operations on the data "atomic". This works by taking a "lock" on the data before operating on it. This means that no other thread can also work on the data in parallel and thus avoids the race conditions described above. Unfortunately however, this approach introduces its own new problems. For example, processes can deadlock if they are trying to both lock the same items of data.

Remembering that the code in figure 3 can potentially have different interleaving of execution so that process P could lock variable "M" and then process Q could lock variable "N". They are then waiting on each other to release their locks to continue execution and no progress can be made.

The conditions that create deadlock

Tasks that claim exclusive control of the resources they require ("mutual exclusion" condition). Tasks that hold resources already allocated to them while waiting for additional resources ("wait for" condition). Resources that cannot be forcibly removed from the tasks holding them until the resources are used to completion ("no preemption" condition). The existence of a circular chain of tasks, such that each task holds one or more resources that are being requested by the next task in the chain ("circular wait" condition).

Fortunately, there are also rules to avoid

Mike Bartley is founder and CEO of Test and Verification Solutions (TVS) - www.tandvsolns.co.uk
He can be reached at mike@tandvsolns.co.uk

Code for Process P	Code for Process Q
Lock(M)	Lock(N)
Lock(N)	Lock(M)
Critical Section	Critical Section
Unlock(N)	Unlock(M)
Unlock(M)	Unlock(N)

Figure 3: Potential for deadlock

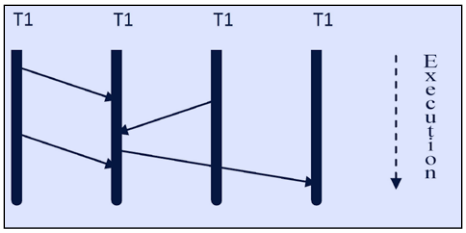

Figure 4: Threads message passing.

deadlock. Each task must request all its required resources at once and cannot proceed until all have been granted ("wait-for" condition denied). If a task holding certain resources is denied a further request, that task must release its original resources and, if necessary, request them again together with the additional resources ("no preemption" condition denied). A linear ordering of resource types should be imposed on all tasks.

But these rules rely on your programmers following them. Locking data in this way can also introduce other problems, for example when a process takes a lock but then throws an exception, does the programmer always remember to release the lock in the exception handler? If not then other processes can potentially hang. These race hazards not only make debugging difficult but they make testing much harder too. It means that if we execute the code once and it passes then it does not mean our job is necessarily done. We have now tested one particular interleaving of the threads and this interleaving worked. There may still be potential interleavings that fail and we simply haven't run them yet. Worse still, there is currently no way of identifying all of the potential interleavings and then measuring which ones we have tested. This means that our traditional models for measuring how well we have tested our code no longer work. So, for example, we may use code coverage techniques (such as line or branch coverage) and get figures of 100%. However, this does not tell us what percentage of the potential interleavings we have tested and so what confidence we should have in our testing. You may argue that this was always the case with code coverage – i.e. it is not a good measure of the thoroughness of your testing. However, whatever measures you have been using to measure the thoroughness of your testing (feature coverage, scenario analysis, exploratory testing, model-based testing) will fall foul of the same issue. Using locks to make access to shared data atomic is not the only technique programmers use to share data. There are other techniques such as message passing shown (see figure 4) where processes communicate with each other by passing messages rather than sharing memory.

This technique makes the sharing mechanism more explicit and allows us to build coverage models around the messages passed. However, all have various potential coding hazards which the tester needs to be aware of. Message passing also adds the benefit that we can start to build coverage models regarding the interleaving of messages passed between the various processes.

There are companies trying to improve the languages we use to write code and how it is executed on the underlying hardware. The language "XC" and hardware developed at XMOS in Bristol is a good example. ■

Designing a USB streaming video control system over WiFi

By Jörg Kaleita

DANIEL IS AN ENGINEER. He is also a proud father, who wants to keep an eye and ear on his small child. Unfortunately, his smart new house has three flours and he does not want to run up and down the stairs every time he hears something. Therefore, he needs a video control system urgently. Daniel is already successfully working with Altium Designer and Altium's prototyping platform NanoBoard 3000.

To keep his project as cheap as possible, Daniel wants to use his existing USB video class (UVC) webcam to support streaming video. If possible, he wants to display the video stream on his brand new iPhone, as he carries it in his pocket anyway. The transmission of the video stream should be realized via the WIFI network in his house. The primary part of an embedded system generated with Altium Designer is a soft CPU (TSK3000, 32-Bit RISC Processor), which comes free of charge with the software package. This includes all necessary IP for the interfaces of the NanoBoard 3000. Altium's OpenBus methodology allows Daniel to design his hardware on a very high abstraction level to program the used FPGA.

The soft CPU has separate peripheral and memory interfaces. Daniel connects a soft terminal to the peripheral interface to display status information. He can use the available TFT display instead just by changing the standard output in the Software Platform Builder described later in this article. Additionally, he connects an interrupt controlled USB-hub to his peripherals to connect the webcam as well as his USB WIFI stick to his NanoBoard.

Next, Daniel connects an SRAM to his memory interface to store the video stream. With a shared memory controller, he can use the same bus for both SRAM and USB. Finally, he configures his soft CPU to allow

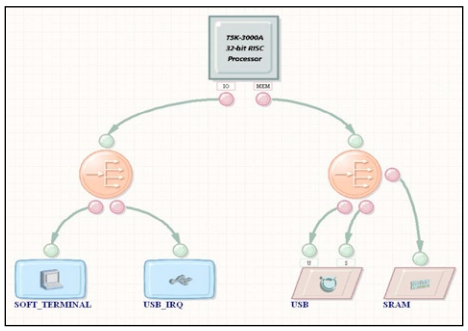

Figure 1: the OpenBus system

32k internal memory inside the FPGA for execution and assigns an interrupt to the used USB controller. Altium Designer will perform the necessary memory mapping automatically and can generate a c-header, if required. With that, Daniel can change the hardware mapping later without modifying the related software. The header allows the software to use the variables in the header, not the physical addresses. The embedded system Daniel generated is based on the well defined wishbone-bus allowing easy access to the growing opencores.org community. Figure 1 shows the complete OpenBus design generated with Altium Designer

Jörg Kaleita is technical account manager EMEA for Altium Europe - www.altium.com

representing Daniel's application.

The top-level schematics will afterwards instantiate the generated OpenBus system and connect the interfaces with the NanoBoard3000 library. Part of this library is the pin out of the used FPGA, therefore Daniel does not need to do it manually to program the device.

Daniel does not really want to simulate his functionality as this requires additional time. But he likes the concept of the so called virtual instruments. These include, for example, a logic analyzer, frequency generators, cross switches, etc. He can also design his own virtual instruments graphically, if needed. They are accessible without starting Altium Designer via a small application called Dashboard. The advantage of this concept over simulation is that Daniel can measure, read and write real signals from within his FPGA.

Due to these virtual instruments, the Altium USB-JTAG cable has two JTAG interfaces: a hard and a soft interface. With the hard JTAG, Daniel programs his FPGA inclusive virtual instruments. Once programmed, the virtual instruments are accessible via the soft JTAG to read out or write the requested data. The soft JTAG requires 4 additional IO-pins of his FPGA.

There are two main reasons why Daniel is using Altium Designer to develop his system: the platform is using the FPGA vendor tools too. Therefore he can do whatever the vendor tool can do and also use the same sources, if he wants to use another vendor's FPGA. While the vendor tools can design very easy and elegant complete embedded systems, they stop when it comes to the software part of the design. In this area, Daniel is using the full potential of Altium Designer. Beside the hardware, Altium Designer will also generate the necessary software beginning with the hardware mappers, the drivers and application software. Furthermore, it will even generate the needed services such as POSIX, storage service or keyboard context. Ideally, he does not even need to use an operating system as all services required for his application will be generated. In many cases he just needs one line of c-code: connect the output of one interface to the input of another. In any case, Daniel can fully concentrate on his application without the need to get the interfaces running as this will be done by Altium Designer's Software Platform Builder (SPB).

The software can be generated with the SPB by opening a new embedded project and linking it with the FPGA-project. This will instantiate the embedded project under the FPGA-project as well, as his FPGA-project can be instantiated later by a PCB-project. Then Daniel adds a software platform to his embedded project and imports the used wrappers out of the FPGA-project. For each wrapper, he grows the software up to the necessary layer, e.g. TCP-IP network adapter for USB WIFI. The WIFI uses the USB Host Wireless Networking Context, which also links to the interrupt controller. Additionally, he generates the USB Host Video context as beside his USB WIFI he is using an USB webcam on the USB Hub of his NanoBoard 3000. For the Virtual Terminal Instrument he grows the stack up to the serial device IO context. This will be used as the standard output within the automatically started POSIX services. If he wants to display a text to his standard output, he just needs to use 'printf ', and the POSIX will redirect this correctly.

As Daniel plans to always use the same router, he switches off the DHCP server in his TCP/IP networking service. In this case, he has to manually assign a valid IP address related to his router. In the wireless network context, he will include the SSID of the used router while using WSA_PSK encryption inclusive password. With that setting his NanoBoard 3000 will automatically register at the router when he powers up his system.

The hardware generated by Altium Designer delivers the video stream coming from the webcam through the software platform builder and provides the transmission over WIFI. Daniel can therefore concentrate on how to connect the two interfaces within his application.

Finally, the test time: Daniel connects the USB WIFI stick and the UVC webcam to the USB hub of the NanoBoard 3000, which is automatically connected to the router after power up. He starts the web browser on his iPhone and enters the address of the video stream shown in the virtual terminal. Now he can see the live video stream. He can rest happily in the knowledge that he will not have to walk all the stairs in his house unnecessarily in the future. ■

Designing USB functionality into embedded systems with no OS

By Vincent Dargaud

USB FLASH DRIVES are a popular, simple and inexpensive method of moving data from one PC to another. Their use in the embedded market has been limited, however, due to the requirement that a system must have USB host capability to communicate with a flash drive.

In the past, this usually meant that the system needed to be a PC. However, the introduction of microcontrollers with USB host capability means that embedded systems can more easily take advantage of this popular portable data storage medium.

There is still an implementation question for embedded designers, however: the USB interface in IT equipment is easily provided via a full-featured, large operating system (OS), either Windows or Linux. But many industrial systems contain either no OS, or a small, fast real-time OS (RTOS) with limited functionality. This article addresses the question, what is the best way to design host support for USB mass storage into embedded systems that do not use Windows or Linux?

Designers will have no trouble finding an MCU optimised for industrial applications and with a built-in USB interface. Typically, industrial MCUs contain no memory management unit, which means that they will not support a full-featured OS such as Windows or Linux. Most families of 16-bit and 32-bit industrial MCUs include some devices that offer an interface for either a USB 2.0 Host or USB 2.0 On-The-Go (OTG) where the device can operate in both host and device modes.

More complicated is the question of implementing the USB software stack (see figure 1) for either a host or OTG device. There are essentially two main choices to be made. Where will the software stack be sourced from? A commercial supplier, freeware provided by the MCU manufacturer, or open-source software?

Does the device require a complete USB implementation, or is a smaller version of the USB software stack adequate?

In making their choice, developers need to take into account questions such as the ease of integration, the robustness and stability of the chosen implementation, the required data rate and response time, the necessity of technical support and the implementation of bug fixes and updates.

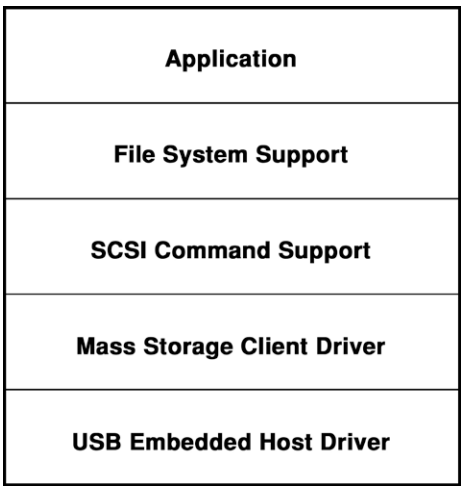

Fig. 1: the layers of the USB software stack.

Developers can choose from royalty-free open-source software, commercial software for which a licence fee must be paid, or freeware supplied by the manufacturer of the chosen MCU. Since the MCU freeware is both royalty-free and backed by technical support, it is often the best option. USB freeware might be limited in terms of performance or code size, but the restrictions will be acceptable for the majority of embedded applications.

Developers who are attracted by the freeware option will need to choose their MCU carefully: some vendors, such as Freescale Semiconductor and Microchip, offer extremely good USB freeware. Freescale's USB stack is implemented in the company's royalty-free MQX RTOS on ColdFire and Kinetis (ARM Cortex-M4) MCUs and Power Architecture microprocessors (see figure 2). Freescale also offers a 'bare metal' USB OTG stack, minimising the memory footprint required on the MCU.

MQX is pre-loaded on a ColdFire MCU on the CrossBow proof-of-concept development boards available from Future Electronics (EMEA) via www.my-boardclub.com.

Microchip also provides a complete freeware solution compatible with its PIC24 and PIC32 families of 16- and 32-bit MCUs. Microchip's stack does not need to be hosted on an RTOS, thus bringing this USB drive capability within the reach of lower-end applications.

Reduced functionality and code

A full USB stack implementation is appropriate for IT equipment, in which there are in practice no constraints on memory space and processor performance. Embedded systems, however, benefit in terms of cost, size and power from any measures that reduce processing and memory overhead – and there are options to achieve this when designing in support for USB flash drives.

One such technique is to identify the known set of USB peripheral devices that the embedded application will interface to, and preload only these known drivers in

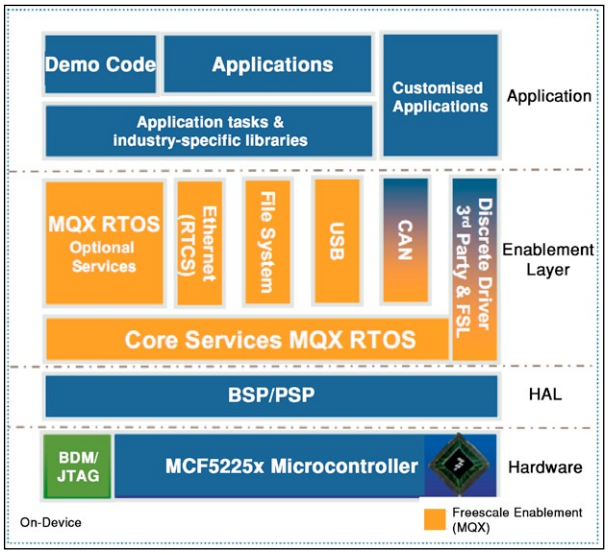

Fig. 2: Freescale's USB stack freeware.

Vincent Dargaud is field application engineer at Future Electronics (France) – www.futureelectronics.com - he can be contacted at info-EUR-Future@futureelectronics.com

the embedded host firmware. The advantage of this so called Embedded Host option over the normal Full Host is the elimination of the extra software and firmware required for the application to act as a full USB host.

A USB embedded host differs from a full USB host in several small but important aspects. The USB embedded host supports only specific peripheral devices and/or classes of devices. It supports only transfer types required by the supported devices and makes hub support optional. Power specification are more relaxed than for a full USB host. These restrictions allow an embedded host to be implemented on a device with fixed, limited memory.

Another memory-footprint reduction option is to implement dual role functionality in place of OTG. This is usable when an application must be able to operate as either host or device, but without switching dynamically between the two, a capability provided by the OTG protocol. Dual role USB eliminates the need for the OTG Host Negotiation Protocol (HNP), and support for the Session Request Protocol is optional, so there is a marked reduction in memory footprint from implementing dual role in place of OTG.

Dual role USB requires both Type A and Type B connectors. For example, a data logger can act as a peripheral device when downloading data to a PC, and act as an embedded host when transferring data to a USB flash drive. The data logger would determine its role as device or host depending on which cable was attached. If a cable was attached to the Type-A receptacle, the data logger would act as a host. If a cable was attached to the Type-B receptacle, the data logger would act as a device. The Type-A and Type-B receptacles must both operate concurrently unless they are only accessible one at a time.

File system support

If data are to be transferred between a PC and the embedded application via a USB flash drive, it is necessary to use a file system to organise and index these data. The most common file system used for this is Microsoft's FAT16 or FAT32 file system with (optional) support for long names. (FAT stands for File Allocation Table.)

For memories smaller than 2GB, the FAT16 standard can be used. The FAT32 standard should be used to address memories larger than 2GB. Both the Freescale MQX and Microchip freeware USB stacks provide file systems compatible with FAT32.

Power management

While memory footprint and feature-selection are important criteria in USB software implementation, on the hardware side power management and distribution are of high importance in the correct design of USB peripherals. Proper methods of designing USB peripheral power distribution are crucial to ensure full compliance with the USB specification, including compliance with electromagnetic interference and voltage-regulation requirements. USB specifications limit the device current to 100mA prior to enumeration. After enumeration, a high-powered device can request no more than 500mA.

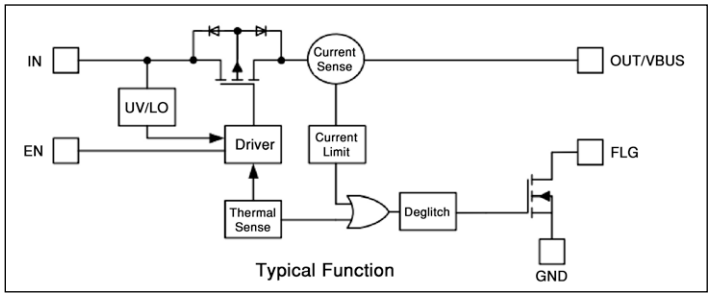

Fig. 3: typical USB power management circuit.

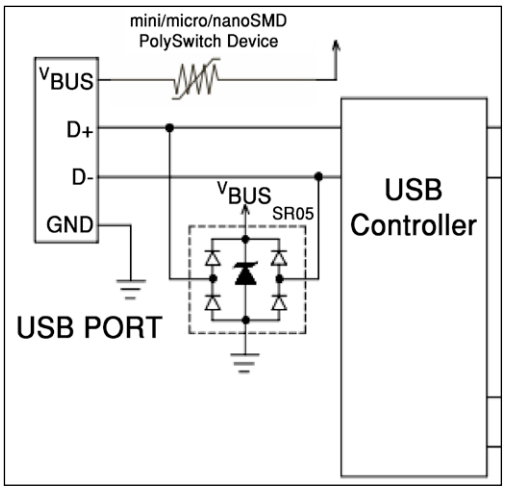

Fig. 4: circuit providing ESD protection for a single USB port.

The DC output voltage, measured at the board side of the device connector, must remain below 5.25V and above 4.75V under all continuous load conditions. Many IC manufacturers offer dedicated circuits to meet these specifications. The two main features required for effective host connection (see figure 3) are the VBUS voltage switching with low drop-out, and current limitation/protection.

If the application has a low power requirement, a thermal fuse such as Tyco's Polyswitch or Polyfuse from Littelfuse is adequate for current protection. But for most applications, a dedicated USB power switch, available from manufacturers such as Micrel, Diodes Inc, AnalogicTech and ON Semiconductor, will be more suitable.

The power circuit also needs to provide protection against electro-static discharge. Today, the internationally recognised ESD standard is IEC 61000-4-2. Implementation of ESD protection is well illustrated through the example of the SR05 low-capacitance Transient Voltage Suppression (TVS) diode array from Semtech. Figure 4 shows a circuit which protects one USB port. When the voltage on the data lines exceeds the bus voltage (plus one diode drop), the internal rectifiers are forward-biased, conducting the transient current away from the protected controller chip. The integrated TVS diode directs the surge to ground. The TVS diode also acts to suppress ESD strikes directly on the voltage bus. Both power and data pins are protected with a single device. ■

PC/104 meets Qseven: what's your combination?

By Matt Ferraro

EMBEDDED ELECTRONICS is all about applying commercially available components to a design in a highly custom fashion. OEM products vary so greatly in computational requirements that a "perfect" off the shelf solution rarely exists. Developing the perfect product can be very time consuming and risky; having the ability to quickly configure off the shelf components to get close to perfection quickly and economically is a great virtue.

Stackable modular computers using PC/104 have been in use for many years. Over time, PC/104 has evolved to encompass the latest computer interconnect technology using PCI Express. Evolutionary changes make the stackable modules a top choice for many embedded computing applications. The PC/104 ecosystem is very mature and offers a wide range of modules of all types - processing, I/O, FPGA, power supplies, and enclosures. The form factor is very popular in applications requiring small, rugged, and highly versatile computing elements.

The Qseven concept

A more recent entrant to the world of embedded computing is the computer-on-module (COM). The Qseven concept is an off-the-shelf, multi-vendor COM that integrates all the core components of a common PC and is mounted onto an application specific carrier board. Qseven modules have a standardized form factor of 70x70mm and have specified pin-outs using a low cost, high-speed MXM connector system with a standardized pin-out, regardless of the vendor. Qseven module functions can include, but are not limited to, graphics, sound, mass storage, network, and multiple USB ports. A single ruggedized MXM connector provides the carrier board interface to carry all the I/O signals to and from the Qseven module.

Unlike previous COM standards, Qseven's primary design best supports mobile and ultra-mobile applications. It defines fast serial differential interfaces such as PCI Express and Serial ATA but omits support for legacy interfaces like EIDE and PCI bus in order to provide support for today as well as future CPU's and chipsets. Qseven COMs are also some of the smallest full function COMs on the market making them very attractive for many embedded applications.

Qseven specifies PCI Express, USB 2.0, ExpressCard, high definition digital audio, Serial ATA, an LPC interface, a secure digital I/O interface, Gigabit Ethernet, DisplayPort, TDMS or SDVO, an LVDS display interface, and a CAN bus, all through a rugged MXM connector.

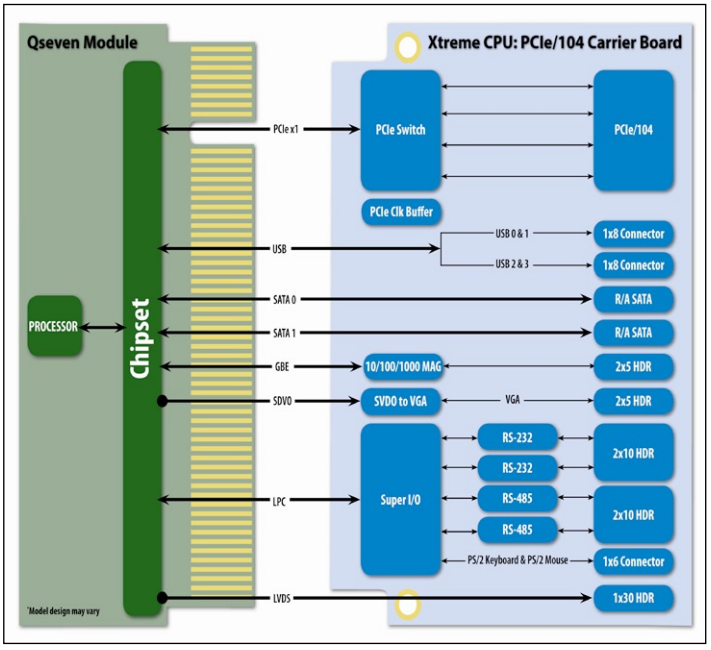

Figure 1: Typical Qseven COM to PCIe/104 interconnect.

Carrier board designers can use as little or as many of the I/O interfaces as deemed necessary. The carrier board can therefore provide all the interface connectors required to attach the system to the application specific peripherals.

In many applications, the number of units to be built or the large number of product variations does not warrant the cost of developing a highly integrated custom computing system. The time and expense to develop is cost prohibitive. Off-the-shelf solutions are frequently missing a key feature or have many features that are unnecessary but add to the cost of the module.

The integration of a highly modular and stackable computing platform with a very small COM form factor provides a level of customization that is well suited for many common applications. The Qseven modules are small enough to fit on the 90x96mm PCI-104 Express and PCIe/104 modules, yet leave enough space for the other necessary functions and connectors to be included on '104 modules. The combination of the two delivers a robust I/O platform and a processing module that can stay current with the latest the industry has to offer.

The Qseven specification defines pin-outs for four PCI Express x1 lanes for expansion capability PCI Express provides a high performance serial interconnect to a wide range of interfaces that are not provided on the COM. Table 1 lists the PC/104 Consortium form factors that match up well with Qseven modules. Each of these PC/104 Consortium boards has direct support for PCI Express. The small size of EBX and Epic form factors make them ideal candidates as host carriers for COMs. Having a mature technology supported by several vendors gives the designers I/O or processing options that can benefit them through the entire design process.

Scalability

Applications are scalable, which means once a product has been created there is the ability to diversify the product range with different performance class Qseven modules and I/O payloads through the PC/104 modules. As Qseven COM host carriers emerge on other form factors, you can reuse a favorite COM and its associated software in a new I/O

Matt Ferraro is a senior hardware engineer & project manager at Connect Tech - www.connecttech.com

configuration. When new COMs and carriers enter the market, you get a multiplying effect on the number of possible combinations. Simply unplug one module and replace it with another, no redesign is necessary.

Processor technology moves very quickly, with enhancements to existing processor families occurring two or three times a year. Compounding the problem are the numerous new processors introduced every year. Having the processing element, memory, and supporting chipsets on a COM lets you try out the newest technology and then move it into your product line at your convenience. Because PCI Express is the common denominator, you have many cross platform processor choices, including Intel Atom, Freescale i.MX, TI OMAP and NVIDIA Tegra, with more to come as the processor suppliers come to appreciate the time-to-market advantages of Qseven. You can offer instant upgrades to existing products and never be trapped in an obsolete design.

Rapid prototyping

Another more subtle advantage is the capability of rapid prototyping or proof of concept development with off-the-shelf PC/104 and Qseven modules. Being able to quickly prototype with various processor and I/O combinations lets you find the most suitable platform for your application. No matter how well you simulate a design, nothing beats being able to create a fully functional prototype.

Flexibility

Additional flexibility is gained by using one of many FPGA modules readily available through the PC/104 Consortium ecosystem. Designers can complete their application by adding FPGA or various I/O solutions from a well-established, robust and well supported PC/104 ecosystem. The programmable flexibility of FPGAs takes the modularity concept to the next level. Several FPGA modules are available from PC/104 Consortium members using Xilinx, Actel, and other FPGAs. System cost can be reduced in lower unit volume applications by using high volume components like Qseven COMs. COM suppliers can increase unit volumes and reduce part numbers generating savings that can be passed onto the integrators. Usually board or module products that snap together have a slightly higher cost associated with them, but with the Qseven COM - PCIe/104 combination, one can use a very cost competitive, higher unit volume module, offsetting the cost of integration with a high degree of design flexibility.

Figure 2: a Qseven module.

Form Factor	Type	Size (mm)	Bus	Advantages
PCI/104-Express	Stackable Module	90x96	PCI + PCI Express	Small size, adds the robust PCI bus and the high-speed PCI Express bus.
PCIe/104	Stackable Module	90x96	PCI Express	Small size, exclusive use of PCI Express bus.
EBX Express	Motherboard	146 x 203	ISA + PCI Express	Same advantages as EBX, plus it adds the high-speed PCI Express bus.
EPIC Express	Motherboard	115 x 165	ISA + PCI Express	Same advantages as EPIC, plus it adds the high-speed PCI Express bus.

Table 1: PC/104 Consortium form factors with PCI Express capability.

Selecting an OS

Operating system and drivers must support your chosen processor platform and I/O payload. If you plan to interchange either one, be sure that the software is capable of handling the changes or compatible software modules are available. Embedded Windows and Linux minimize this challenge, but you should be sure to get the latest drivers and BIOS from the COM supplier to avoid software integration issues. ■

Mini-ITX form factor motherboard

with support for 2nd generation Intel core

Emerson Network Power has announced its latest series of Mini-ITX form factor embedded motherboards will support the new 2nd generation Intel Core i7-2710QE and Intel Core i5-2510E processors. The MITX-CORE-800 series targets makers of intelligent kiosks, digital signage, medical clinical equipment and gaming machines. It will also offer improvements in graphics processing, PCI Express and on-board storage interfaces, says the company, as well as enhanced software support.

Featuring up to 8GB DDR3 memory, dual display capability from multiple physical display connections and Intel vPro technology, the new embedded motherboards should appeal to the rapidly growing market for connected consumer portals.

The MITX-CORE-800 series features a PICMG standard Extended Application Programming Interface (EAPI) to simplify the control of essential hardware functionality such as the backlight inverter and the watchdog timer. Emerson Network Power says it has also simplified the way in which developers can customise and configure the motherboards' BIOS to speed integration and accelerate time-to-revenue.

Emerson Network Power

www.emersonnetworkpower.com

System trace and profiling tool

offers visual insights in runtime behaviour

Quadros Systems has introduced RTXCview, a system trace and profiling tool for embedded systems. Using a vertically scrolling run-time viewer, RTXCview works as a companion to traditional debuggers, providing a visual insight and understanding of the runtime behaviour of embedded systems using the RTXC RTOS.

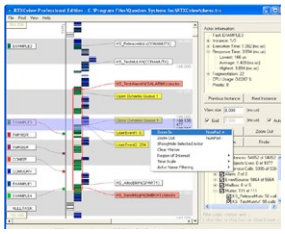

At many times in the course of developing an application, the developer needs to understand exactly what is happening inside the system. RTXCview allows the developer to study the detailed runtime behaviour of the system as it actually occurred.

In real production code settings, RTXCview addresses the questions of who, what, when and why by presenting the runtime information in a unique, highly efficient manner. The tool shows everything on a scalable timeline using the vertical axis, allowing the user to zoom in and out to get a different look to address the question of when. The unique presentation allows for zooming in on details, down to microseconds, but also gives a good overview when zooming out over several seconds. The question of who can be easily determined because tasks, threads and interrupt service routines are clearly shown in color-coded sections along the timeline.

Quadros Systems

www.quadros.com

Open NFC protocol stack for Android

promises consistent API and functionalities

Inside Secure launched a new version of the Open NFC protocol stack geared to the latest version of Google Android (aka Gingerbread). This makes Open NFC the first truly hardware-independent, open-source NFC protocol stack for this popular smartphone operating system.

Open NFC version 4.2 for Google Android 2.3 simplifies interoperability and provides the NFC ecosystem with a consistent NFC application programming interface (API) and functionality, offering chip vendors, smartphone manufacturers, wireless carriers and software developers a way to implement NFC functionality independently of the underlying NFC hardware as Gingerbread is adopted for use in a broad range of mobile products around the world. With its separate hardware abstraction software layer, adapting the Open NFC stack to new hardware is much simpler as the hardware dependencies are well layered and the overall code base to adapt is significantly smaller.

Open NFC supports several levels of functionality, from low-level RF control to high-level NFC Forum tag handling, peer-to-peer communications as well as Bluetooth and Wi-Fi pairing, interactions with single-wire protocol SIMs and other secure elements and compatibility with smart cards and RFID tags based on Felica, Mifare and ISO 14443 standards.

Open NFC 4.2 for Android 2.3/2.4 will be available for download on the INSIDE Secure website by the end of February as a free and open-source edition under the Apache License, Version 2.0.

Inside Secure

www.insidesecure.com

1.5GHz gaming platform

features the dual core 1.6GHz AMD Fusion

Quixant launched the QXi-200 gaming controller, the first such controller to use AMD's new Fusion technology, says the company. AMD Embedded G-Series APUs integrate onto a single chip one or two new high performance 64-bit x86 processor cores and cutting-edge AMD Radeon HD 6310 graphics. The QXi-200 is available in both single core 1.5GHz and dual core 1.6GHz variants. The AMD Radeon HD 6310 graphics engine on the APU features support for the latest DirectX 11, OpenGL 4.0 and OpenCL standards.

The UVD3 multimedia engine enables playback of both SD and HD video feeds simultaneously with advanced graphics.

The QXi-200 can accomplish this on two independent high resolution digital monitors. Another major benefit of the AMD Embedded G-Series platform is its very low power consumption. It is combined with Quixant's patented case design to provide a totally fanless system in the QXi-200. The platform operates from a single 12V supply, eliminating the extra expense and reliability issues of an ATX PSU. The gaming architecture of the QXi-200 is the same as on the QXi-100, enabling customers to easily migrate and take advantage of the extra features and performance the QXi-200 provides.

Quixant

www.quixant.com

Digital signage solutions

fanless embedded units based on AMD processors

Mdina Media has revealed its latest digital signage solutions, each based on an AMD processor and targeting low power digital signage applications.

The Fanless Embedded system measures 175x180x72mm, it is powered by an AMD dualcore 25W processor and carries an AMD 780E, SB710 chipset with integrated ATI Radeon HD 3200 graphics. The use of industrial strength SSD drives with SLC technology eliminates any moving parts.The company also unveiled a 107x154x60mm nano-PC system powered by an Embedded AMD Turion II Neo N54L processor (2.2GHz). It features the AMD 785E, SB 8X0 chipset with an integrated ATI Radeon HD4200 graphics. Based on the brand new power-optimized AMD Embedded G-Series Platform with 9W or 18W TDP, the Mini-ITX Embedded system (177x177x50mm) is powered by the AMD Fusion APU which integrates a x86 CPU and a discrete-level AMD Radeon GPU. The system supports DVI, HDMI, VGA video outputs and advanced interfaces such as 6Gb/s SATA, 2nd Generation PCI Express, and HD Audio. An integrated DDR3 memory controller running at up to 1066 MHz is standard. All three systems can adopt various configurations of SSD drives.

Mdina Media

www.mdinamedia.com

CompactPCI serial subrack system

supports data transfer rates up to 32Gbit/s

Schroff has announced a complete 4U, 9-slot subrack system based on the recently released CompactPCI Serial specification. The system takes full advantage of the specification's high-speed serial bus topology to deliver a data transfer rate of up to 32Gbit/s via PCIe, 10GbE, S-ATA, USB 2.0 or USB 3.0 interfaces. Particularly suitable for applications in the process, industrial control, defence and transportation sectors, the new CompactPCI Serial system comprises a subrack, backplane, fan tray and power supply.

The shielded 19in. aluminium chassis is 4U high and has a depth of 275mm. Featuring a perforated top cover and base plate, it is supplied with front handles and a set of guide rails to enable it to be populated with 3U-high boards. The 9-slot backplane - providing one system slot and eight peripheral slots - conforms to PICMG CompactPCI-S.0 and can be specified with or without rear I/O. Effective cooling of the system is ensured by a 1U hot-swap fan tray under the board cage. Operating from a wide input range of 100 to 240VAC, the built-in 300W ATX power supply features an IEC AC input socket, mains switch and dedicated fan. In addition, the CompactPCI Serial system includes plug-in units for DVD and hard-disk drives, as well as various 3U front panels.

Schroff

www.schroff.co.uk

Magneti Marelli joins Wind River

to create the Genivi-compliant infotainment platform

Automotive component supplier Magneti Marelli and embedded software vendor Wind River have announced a technological collaboration to create what they call the first Genivi-compliant in-vehicle infotainment (IVI) solution for the automotive industry. The first customer is BMW. Magneti Marelli will lead the development efforts, contributing its automotive know-how in integrating complex systems and technologies for the vehicle environment. Wind River provides an integrated, tested and validated IVI software platform based on the Genivi open source standard, as well as customization and consulting services.

The Genivi open source platform standard aims to provide automobile manufacturers and their suppliers a common underlying framework to simplify elements of the in-vehicle infotainment development process that have historically been duplicated across the industry.

Wind River

www.windriver.com

Windows7-compatible Micro ATX board

for Intel Core i7/i5/i3 64-bit processors

The latest addition to BVM's extensive family of SBCs is the MS-C72, a powerful platform with a wealth of features including I/O and mass storage. It conforms to the Micro ATX footprint and is fully compatible with Windows 7. Based on the Intel QM57 chipset, the MS-C72 supports the desktop 64-bit 32nm Intel Core i5/i3 processors with HD graphics and the 45nm i7/i5 without HD graphics, both in the LGA1156 socket. Up to 16GB of DDR3 1066/1333 system memory is available, and in the HD graphics versions, the graphics core on the processor die gives a two-chip system architecture that requires less power than alternative three-chip platforms. Intel turbo boost and hyper-threading technology optimise performance to match the system workload, increasing the efficiency of processor-intensive tasks.

BVM

www.bvmltd.co.uk

Credit card-sized computer on module

nanoETXexpress 2.0 specification released

The nanoETXexpress Industrial Group has published the nanoETXexpress 2.0 specification for credit card-sized COM Express ultra Computer-on-Modules. The new revision aims to address the requirements of new and highly compact applications based on SFF processors even more explicitly. It incorporates all the changes of the COM Express specification rev. 2.0 that affect the ultra standard. By continuing to follow the only manufacturer-independent standard for Computer-on-Modules, any carrier board designed for nanoETXexpress is consequently always developed in a way that is also COM Express compatible. This enables users to benefit from the as-needed scalability of COM Express Computer-on-Modules, from "basic" and "compact" to the "ultra" footprint, already designated by many as COM Express "ultra". The consortium of leading embedded computer manufacturers developed the nanoETXexpress specification in order to provide the market with a solid manufacturer-independent basis for a corresponding COM Express ultra form factor through the PCI Industrial Computer Manufacturers Group. In addition, the aim is to promote a uniform specification for Computer-on-Modules in order to offer users a single, universally usable Computer-on-Module standard for development and product maintenance. Among the members of the nanoETXexpress Industrial Group are Aaeon, ADLINK, Advantech, E.E.P.D., IBase, Toradex and Kontron.

nanoETXexpress

www.nanoetxexpress.org

Miniature desktop computer

130x95x15mm unit runs a Nvidia Tegra 2

CompuLab is introducing Trim-Slice, a miniature desktop computer powered by the Nvidia Tegra 2. With an all-metal 130x95x15mm housing, CompuLab describes Trim-Slice as its smallest and most energy-efficient computer to-date. Trim-Slice differs from most other ARM based solutions by being a commercially available open platform for software developers. The Tegra 2 integrates a 1 GHz dual-core ARM Cortex A9 and an ultra-low power GeForce GPU onto the same chip, making it one of the most powerful ARM based system-on-chip available today. The high performance, low-power, rich I/O and miniature rugged design are intended to position Trim-Slice as an attractive solution for a variety of applications - media players, IPTV, infotainment systems, signage, gaming or even as a desktop replacement to name a few. CompuLab will cooperate with ISVs that select Trim-Slice as a reference platform. The unit features an NVIDIA Tegra 2 Dual Core ARM Cortex A9 1GHz CPU with integrated ultra-low power GeForce GPU. It has 1GB DDR2-800 of storage, takes full size SD (SDHC), Micro SD (SDHC) and SATA SSD (up to 64GB). The miniature desktop can be networked via 1 GbE; WiFi 802.11n, it supports HDMI 1.3 full-HD and DVI (dual head), stereo line-out, line-in, 5.1 digital S/PDIF.

CompuLab

www.trimslice.com

Vehicle control system

features 2-wire CAN connection

Müller Industrie-Elektronik has developed a flexible vehicle control system for the functional requirements of commercial vehicles. The VCS vehicle control system consists of a VCS-TU control unit in small design, mountable in a common DIN radio slot in a vehicle. Alternatively there is an operating unit in a robust heavy-duty-case with wired capability of connecting the additional remote unit VSC-RU in handheld case. This remote unit in handheld case can be used as an independent controlling unit to control seven different functions. The two mounting operating units VCS-TU-DIN and VCS-TU-HD are customized programmable up to nineteen control commands and have push-button slide-in labels for user-specific identification of the command functions. Programming can be done via an external programmable logic control or via the software included, in combination with the Müller power unit. Setting and retrieving vehicle data is available through the display. The power unit module VCS-PU is available in three versions with different switching capacity in various design of cases. Options include vehicle case or heavy-duty-case in two sizes, provided with high protection class IP69k. An OEM-version is available with PCB in open-type assembly.

Müller Industrie-Elektronik GmbH

www.mueller-ie.com

Packet processing software

for XLP832 multicore communications processor

NetLogic Microsystems and 6WIND are expanding their collaboration to deliver high-performance solutions for next-generation telecom infrastructure, networking equipment and security appliances. Building on existing solutions, 6WIND will support the XLP832 multi-core communications processor, designed for next-generation networking applications such as security appliances, Layer 4 through Layer 7 switching, storage networking, 3G/4G wireless and small business networks. The 6WINDGate software delivers the high packet processing performance required in networking equipment for data center and cloud applications which must be capable of scaling to line rates of 160Gbps while providing support for intelligent networking, advanced security protocols and energy reduction. NetLogic Microsystems' XLP multi-core processor family features quad-issue, quad-threaded and superscalar out-of-order capabilities with scalability to 128 NXCPUs operating at 2.0GHz. Designed in TSMC's advanced 40nm process, the XLP cores are quad-threaded to effectively minimize bottlenecks and memory latencies that are inherent in network data-plane processing applications, and are equipped with a tri-level cache architecture with over 50 Mbytes of fully coherent on-chip cache. 6WINDGate is expected to enable customers to eliminate up to twelve months of development time, reaping the rewards of being early to market with cost-optimized products of higher performance.

NetLogic Microsystems

www.netlogicmicro.com

32-bit ARM Cortex microcontrollers

with on-chip full speed USB device connectivity

Toshiba Electronics Europe announced a range of 32-bit ARM Cortex-M3 microcontrollers including parts featuring on-chip full-speed (12Mbps) USB-device connectivity. The TMPM366 devices are supplied in LQFP100 packaging, the TMPM366FDFG, TMPM366FYFG and TMPM366FWFG combine on-board Flash program memory with comparatively high levels of on-board RAM. Respective Flash/RAM capacities are 512Kbytes/64Kbytes, 256Kbytes/48KBytes and 128Kbytes/32Kbytes.

Toshiba Electronics Europe

www.toshiba-europe.com

Six new parts in the PIC MCU family

with Ethernet, CAN and USB connectivity

Microchip launched a six-member family of 32-bit PIC32MX5/6/7 microcontrollers that provide integrated, CAN, Ethernet USB and serial connectivity peripherals with new, more cost-effective memory options. The parts draw 0.5 mA/MHz of active current, flash memory endurance is 20k read/write cycles. By maintaining common pin-outs, the PIC32 portfolio provides designers a good balance of memory and cost for their high-performance applications as well as a seamless migration path for scalability and flexibility. The raw performance of the MIPS32 M4K core has been maximised to achieve best-in-class performance of 1.56 DMIPS/MHz. The family provides 32 Kbytes of RAM and up to 140 Kbytes of flash. Each of the six new microcontrollers is available in five different pin-compatible packages: 100-pin TQFP 12x12mm, TQFP 14x14mm and BGA packages, as well as 64-pin TQFP and QFN packages.

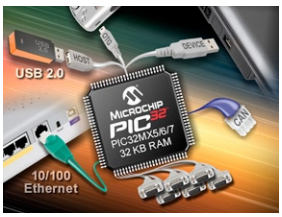

Microchip

www.microchip.com

3.5” embedded board

based on the Intel Core i7 processors

Based on Intel’s 32nm process technology and featuring a two-chip platform that integrates the Northbridge chip with the CPU, the ECM-QM57 from Data Modul delivers enhanced performance, energy efficiency, manageability, security functions and smoother visual experiences. Using the low voltage Intel Core i7-620LE or the ultra-low voltage Intel Core i7-620UE processor, it is equipped with the latest Intel QM57 Express chipset. It provides 5 x USB 2.0, 2 x COM (one of them is switchable to RS232/422/485), 16 GPIOs, 2 x SATA, HD Audio, Dual Gigabit Ethernet and one Compact Flash Socket. One Mini PCI Express Socket is optionally available via Daughter Board.

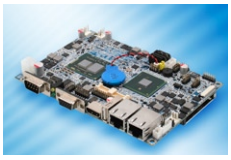

Data Modul

www.data-modul.com

Audio processing chip

boosts the performance of small low-cost speakers

STMicroelectronics' STA321MP audio processing chip connects directly to the latest miniature microphones and provides a new way of boosting performance from small, low-cost, or even damaged speakers. The STA321M features a built-in interface for a MEMS digital microphone as well as a standard microphone input. The audio processor has a second MEMS input that can be connected to a tiny sensor for detecting speaker movements. This allows designers to apply active-sound shaping using ST's patented techniques to optimize speaker performance. In active-sound shaping, the processor compensates automatically for distortion or damage to the speaker enabling high-quality audio performance from over-driven speakers or from those with torn or restricted-movement cones.

As a third-generation Sound-Terminal IC, the STA321MP delivers advanced features and performance. Support for designers includes complete development kits for MEMS-based systems, as well as the APWorkbench graphical development environment. APWorkbench helps designers configure Sound Terminal products intuitively, and includes features for setting up active sound shaping.

STMicroelectronics

www.st.com

High radiant intensity infrared emitter

delivers up to 55mW of optical power, 10ns switching

Vishay Intertechnology has broadened its optoelectronics portfolio with the release of an 850nm infrared emitter featuring a parabolic lens for a narrow ± 3° angle of half intensity. Based on a surface emitter chip technology, the VSLY5850 offers radiant intensity of 600 mW/sr at a 100mA drive current, high optical power to 55mW, and fast 10ns switching times.

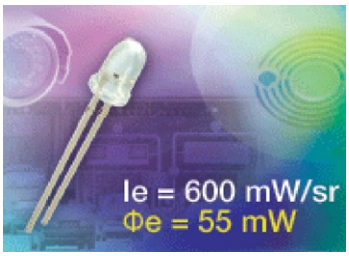

The surface emitter technology used in the VSLY5850 represents a unique die construction in which all the light generated inside the semiconductor is emitted through the top surface of the chip. This greatly reduces side emissions out of the device's 5 mm plastic package, providing the user with a narrow, well-directed emission beam without disturbing fractions towards the sides.

The emitter's high radiant intensity allows significant intensities to be achieved at low drive currents, reducing power consumption by up to three times when compared to the next available intensity class of infrared emitters. The VSLY5850 is optimized for IR illumination in CMOS cameras, fire alarm systems, and smoke detectors. The emitter operates from - 40 to + 85 °C and is suitable for high pulse current operation.

Vishay Intertechnology

www.vishay.com

Inductive sensor integrates MCU

offers an adaptive detection range up to 29mm

Eaton's Electrical Sector is launching a new iProx inductive sensor series that integrates an MCU. Thanks to the Smart Sense technology, the devices can be adapted to each application. The cylindrical sensors offer switching distances of up to 29mm, twice the switching distance of other inductive sensors of the same size, according to the manufacturer.

A large number of additional functions can be activated using optional programming tools. The ProxView software for Microsoft Windows and Windows Mobile devices can be used to program any sensor for a wide range of applications so that it is also possible to adapt and iProx sensors as 100 % replacements for other inductive sensors. The iProx product range includes several models as 3-wire or 4-wire DC devices, different thread diameters from M12 to M30, and a wide range of connection options. All iProx sensors can detect metal targets reliably. They are designed with a stainless steel cylinder, a 300 ° visible dual-colour LED, and corrosion resistant front plates. The connection configuration on DC-iProx models is automatic, so that it detects NPN and PNP terminals automatically and switches the sensor accordingly (NPN = sinking : PNP = sourcing). Sensing range, band detection background object detection (metal), delay and speed detection can all be configured, thanks to the microprocessor-based Smart-Sense technology. A teach function simplifies the programming of the sensor in its application environment. The sensors can also be set so that disturbing metallic objects can be ignored in the foreground or in the background. Enhanced time control functions such as delays and speed detection can also be enabled without the use of external controllers.

Eaton's Electrical Sector

www.eaton.com

Contactless gearbox sensor

predicts shifting direction

Automotive electronics supplier Paragon AG has developed a new sensor for gearbox applications. Unlike available products, the sensor detects already during the start of a shifting process which gear will be put in.

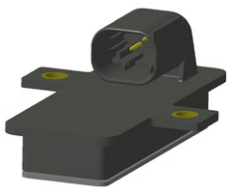

The contactless sensor is designed for manual transmissions and can also be used to detect the neutral position. The device employs a novel inductive measurement principle for which a patent application has been submitted. The sensor enables the design of comfortable gear switches, the company promised, which make the gear shift process smooth and fast at the same time. In addition, separate sensors for reverse gear become redundant.

Paragon AG

www.paragon-online.de

Tiny GPS receiver

connects to both GPS and GLONASS satellites

ST-Ericsson has launched the CG1950, what the company claims to be the smallest receiver able to see both GPS and GLONASS positioning satellites, enabling mobile devices to deliver much faster and much more reliable location-based services.

Requiring just two external components, the CG1950 is the first GPS/GLONASS receiver to be built in 45 nanometer silicon, enabling manufacturers to produce sleek, low-cost devices capable of supporting highly-accurate navigation, mobile social networking, augmented reality and other location-based services. The low-power device is designed to be integrated into mobile phones, camcorders, cameras and other mobile devices. The combination of GPS and GLONASS, will enable devices to retrieve positioning data from more than 50 satellites by the end of 2011.

The receiver features -163 dBm acquisition sensitivity and -165 dBm tracking sensitivity with embedded LNA (single-ended antenna), autonomous (cold start) time to first fix (TTFF) of less than 38 s, and A-GPS (hot start) TTFF below 1 s. Its low power consumption enables continuous GNSS (Global Navigation Satellite System) tracking for more than 30 hours with a standard 850 mAh battery.

Set to be in commercial handsets by the third-quarter of 2011, the CG1950 will be available both as a standalone component and integrated into ST-Ericsson complete cellular platforms.

ST-Ericsson

www.stericsson.com

Ultracapacitor-based backup module

for uninterruptible power supply systems

Maxwell Technologies has introduced a 56-V ultracapacitor module designed specifically to address the short-term ride-through and bridge power requirements of uninterruptible power supply (UPS) systems for mission-critical installations. Features include maintenance-free operation, an estimated 14-year life to ensure low cost of ownership, 3U and 4U rack-mount form factors for easy integration into standard equipment racks, and the ability to connect in series for systems requiring up to 750 V.

The modules are available in 4-kW, 5-kW, 6.5-kW and 10-kW versions to easily configure for UPS system requirements and feature a rugged construction that meets IBC Zone 4 earthquake resistance standard. Further, the modules use green technology with no heavy metals or toxic substances requiring special recycling.

Unlike batteries, which produce and store energy by means of a chemical reaction, Maxwell's BOOSTCAP ultracapacitor products store energy in an electric field. This electrostatic energy storage mechanism enables ultracapacitors to charge and discharge in as little as fractions of a second, perform normally over a broad temperature range (-40 to +65°C), operate reliably through one million or more charge/discharge cycles and resist shock, vibration and overcharging.

The company offers ultracapacitor cells ranging in capacitance from 5 to 3,000 farads and multi-cell modules ranging from 16 to 125 V.

Maxwell Technologies

www.maxwell.com

Power MOSFET family for DC-DC

improves switching applications efficiency up to 2%

International Rectifier has introduced a family of DirectFETplus power MOSFETs featuring IR's new generation of silicon that sets a new standard in efficiency for 12V input synchronous buck applications including next-generation servers, desktops, and notebooks.

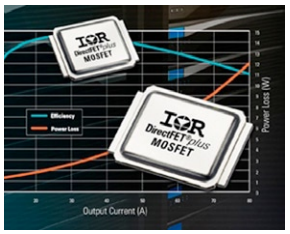

The first two DirectFETplus devices in the new family, the IRF6811 and IRF6894, reduce on-state resistance (Rds_{on}) and gate charge (Q_g) compared to previous generation devices to improve efficiency up to two percent. In addition, the devices offer ultra low gate resistance (R_g) enabling further efficiency improvement by minimizing switching losses in DC-DC converters. The IRF6811 control MOSFET is available in a Small Can while the IRF6894 synchronous MOSFET is offered in a Medium Can. The 25V DirectFETplus pair combines industry leading Rds_{on} and R_g, combined with low charge to minimize conduction and switching losses. The IRF6894 also features a monolithically integrated Schottky that reduces losses associated with body diode conduction and reverse recovery. The new DirectFETplus MOSFETs are footprint compatible with previous generation devices.

International Rectifier

www.irf.com

Single-layer capacitive multi-touch MCU

does not require cross-over isolation points

Integrated Device Technology announced a touch screen controller IC optimized for its proprietary single-layer multi-touch projected capacitive touch screen technology. The touch screen technology and controller solution offers a true single-layer solution, not requiring cross-over isolation points that other solutions require, yet still maintaining full multi-touch capability. The IDT LDS7000 and LDS7001 multi-touch, full resolution touch screen controller ICs feature up to 30 and 35 sensor channels, respectively. Both controllers operate at a fast 8ms data rate for quick response to touch inputs to improve the user's experience and enable advanced applications that require fast response times. The IDT LDS7000 family's proprietary algorithms eliminate undesired multi touch ghosting, providing accurate individual X and Y coordinates in dual touch mode, a critical requirement for customers using host-interpreted custom gestures to differentiate their end product. The devices' analog front-end design provides a high level of noise rejection performance that negates the need for a separate touch screen shield layer for most applications, lowering the overall solution cost even further. The IDT LDS7000 and LDS7001 devices offer native support for a 10-bit resolution (1024 x 1024) with the ability to scale the touch values of the X and Y axes.

Integrated Device Technology

www.idt.com

IP suites for Lattice's ECP3 FPGAs

with ready-made building blocks for high-speed data transfer

Lattice Semiconductor has announced five IP suites for the LatticeECP3 FPGA family, comprising PCI Express, Ethernet Networking, Digital Signal Processing, Video & Display, and Value. The suites offer ready-made building blocks for high-speed data transfer, Ethernet networking, high speed memory interfaces, digital signal processing and video pixel processing. Lattice is offering a limited quantity promotion on a complete set of design tools to encourage customers to build their next generation systems with LatticeECP3 FPGAs, which includes five IP Suites, the Lattice Diamond Design Software Subscription License, and the LatticeECP3 PCI Express Development Kit. All Lattice IP can be fully evaluated prior to the purchase. In the free evaluation mode, customers can fully configure an IP, integrate it in their designs, perform full verification, and even run it in hardware for a limited time. Purchase of an annual node-locked IP Suite license enables the member IP to operate in hardware for an unlimited time.

Lattice Semiconductor

www.latticesemi.com/IPSuites

Analog ECG front-end subsystem

integrates five ADC channels for diagnostic-quality data

ADI's first high-accuracy ECG (electrocardiogram) analog front end (AFE) includes pacemaker pulse detection and respiration measurement for battery and line-powered ECG applications. These measurements enable the accurate analysis of numerous heart conditions, including birth defects, arrhythmias, problems with heart valves and lack of blood flow to the heart muscle. The ADAS1000 simplifies the design of a five-electrode ECG system by significantly reducing the signal chain bill of materials from up to 50 components down to a single chip plus a few discrete components. The device can be configured to optimize noise performance, power, or data rate, making it suitable for home, ambulatory, and clinical ECG systems. Future products in the series will provide optimized features and functionality to support traditional and emerging ECG system applications. A DC-coupled channel implementation offers simplified input switching, increased versatility, reduced power and distinct post processing advantages. The chip operates the five ECG electrode measurements from as little as 19mW and any unused channels or features can be conveniently disabled to further minimize power to as low as 11mW for one lead. The device's low noise performance (10μV pk-pk over 0.05Hz to 150Hz) supports end equipment regulatory standards, and different data frame rates including 2, 16 and 128kHz are available to ensure ultimate ease in data capture. The ADAS1000 comes in a 10x10mm 64-lead LQFP package.

Analog Devices

www.analog.com/healthcare

Radio energy harvesting kits

win one of two kits for either wireless sensors or battery charging

This month, Powercast is giving away two of its radio energy harvesting kits, one geared towards powering wireless sensors (P2110-EVAL-01) and the other (P2110-EVAL-02) more specifically designed for battery charging applications. The P2110-EVAL-01 combines Powercast's 915 MHz RF energy harvesting system and Microchip's eXtreme low power PIC MCU to eliminate batteries in a wide range of applications such as wireless sensors. This kit includes a 3W Powercaster transmitter as the power source, two P2110 Powerharvester receiver evaluation boards, two custom-designed wireless sensor boards, the XLP 16-bit development board pre-loaded with jointly developed application software, an IEEE 802.15.4 transceiver and other accessories. It helps demonstrate and develop smart-energy, wireless-sensor applications that are reliably powered by radio waves up to 40 feet away from the emitter. The P2110-EVAL-02 kit features near-loss-less energy storage and efficient charging via the Thinergy Micro-Energy Cell (MEC), a solid-state, rechargeable thin-film micro-battery from Infinite Power Solutions. It also includes the Powercaster and the P2110 board, a custom-designed battery-charging board, a 1mAh Thinergy MEC evaluation card and other accessories.

Powercast

Check the reader offer online at **www.electronics-eetimes.com**
www.powercastco.com

Bio-metrics for NFC

provides two-factor authentication to mobile applications

Inside Secure has announced biometric matching capabilities for its SecuRead NFC component. The system-in-package device enables manufacturers of NFC-enabled devices to provide two-factor security and greater privacy for a variety of mobile applications. The component uses fingerprint identification software from Neurotechnology running on the SecuRead secure element and debuts on the TazCard, the NFC electronic wallet from TazTag SA. The PC will first be used to capture the TazCard user's fingerprint as part of the enrollment process. The fingerprint minutiae points will be extracted by the Neurotechnology software on the PC and transmitted to the TazCard, where it will be securely stored in the SecuRead. When a purchase is attempted the SecuRead will then use its onboard Neurotechnology software to compare new fingerprint data with the trusted data previously stored during enrollment and which never leaves the secure element. If the two match the payment transaction will be allowed to proceed. The same process can be used for other types of applications such as secured access control or ID. The TazCard incorporates both the SecuRead module and the VaultIC security module to protect access to the personal data and applications. VaultIC provides secure storage of keys, certificates and user data while dramatically reducing or eliminating the need for custom development. Privacy is ensured because the original fingerprint data is never exposed again.

Inside Secure

www.insidesecure.com

16-channel GPS simulator

supports advanced GPS receiver signal testing

Spectracom announced the Pendulum GSG-55 GPS receiver test instrument capable to simulate Satellite-Based Augmentation Systems (SBAS). Navigation systems that use SBAS can improve the accuracy and reliability provided by the GPS satellite signals alone, enabling critical applications such as aircraft navigation, and surveying and mapping. SBAS simulation (support for Europe's EGNOS and North America's WAAS) is a new feature in the GSG-55. It is also able to generate white noise, making it possible to test receiver sensitivity under different signal-to-noise ratios. The GSG-55 builds on the Pendulum GSG-54 eight-channel simulator including accurate testing of GPS timing receivers in a lightweight bench-top chassis.

Spectracom

www.spectracomcorp.com

Sealed circuit headers for transport applications

high-density connection systems target powertrains

Molex extends its CMC family with the introduction of a 154-circuit header in compliant-pin mount and 32- and 112-circuit solder-mount headers. Designed to perform in high-conductivity applications and in harsh environments, CMC is an industry standard interface used in automotive and transportation powertrain applications, including ECUs (engine control units), automatic gearboxes, suspension controllers and electric parking brakes. Customers using Molex 154-circuit, three pocket, right-angle header with compliant-pin terminals will benefit from using a solder-free process, as opposed to conventional press-fit type mounting processes, improving productivity and reducing assembly costs.The 154-circuit header features 120x0.635mm signal pins, 24x1.50mm terminals and 10x2.80mm power pins, supporting high currents in applications such as large ECUs. The sophisticated and robust design of the integral guiding plate on Molex 154-circuit header protects the 5 terminal rows and ensures positioning whilst withstanding high-vibration levels up to 10G. This makes the header suitable for heavy-duty applications such as those found on agricultural equipment.The 32-circuit, single pocket, right-angle header features solder-mount terminals and suites small powertrain applications as well as hydraulic and electronic suspension controllers.

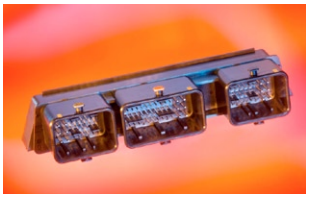

Molex

www.connector.com

Touch screen solution offers 884 nodes

for large multi-touch screens up to 11,6 inches

Cypress' latest single-chip TrueTouch solution for tablets offers more sensing channels than any other single chip solution for large multi-touch screens up to 11.6 inches. Additional sensing channels are said to be essential for greater accuracy, linearity, support for smaller sized fingers and the ability to resolve multiple fingers close together. The new CY8CTMA884 family offers 60 sensing I/O channels with support for up to 884 nodes on the screen. The CY8CTMA884 offers 10-finger touch support, low power consumption, grip and palm rejection, high accuracy, and fast scan times. All major tablet operating systems will be supported in this family of devices, and Cypress says it works closely with all the leading ITO sensor vendors, providing designers with flexibility, optimal performance and quality.

Cypress Semiconductor

www.cypress.com/go/TrueTouch

Industrial connector PCB adaptor

enables direct board mounting

Harting launched a PCB adaptor which allows direct board mounting of its Han Q 12/0 industrial connector to device boards. In addition to the actual PCB adaptor, special versions of the Han Q 12/0 connector's male and female inserts for the PCB adaptor are being introduced. The PCB adaptor can be installed independently of the connector, allowing it to be pre-mounted according to the pinhole pattern on the board, as with other components. The male or female insert for the PCB adaptor is installed after the board has been installed in the housing by being mounted on to the PCB and fixed in place on the bulkhead-mounted housing with the fixing screw. Because of the separation of the connector from the PCB adaptor, all 12 contacts and the PE conductor can be used for connection to the PCB. The adaptor meets the requirements of the DIN EN 61984 standard, with electrical ratings of 7.5 A and 250 V rated and a 4 kV impulse voltage at pollution degree 3. The new PCB adaptor model complements Harting's existing range, which includes adaptors for the Han DD, Han E, Han-Modular and Han Q series of connectors.

Harting

www.harting.com

Meter shunt resistor

offers direct PCB mounting, low resistance

Vishay Intertechnology has introduced a new Power Metal Strip meter shunt resistor that is the industry's first with sense leads designed to be soldered directly to the PCB, eliminating the need for costly flexible leads. The WSMS2908 combines a 3 W power capability in the 2908 size package with low resistance values down to 100 µΩ.

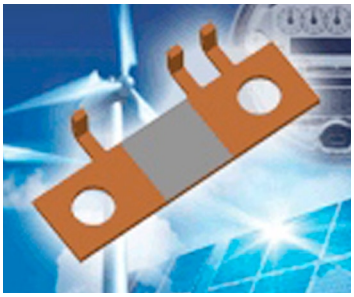

The WSMS2908 meter shunt resistor features a proprietary processing technique that produces extremely low resistance values of 100 µΩ, 250 µΩ, 300 µΩ, 430 µΩ, and 500 µΩ. These values allow for increased accuracy in current meter shunt applications for industrial and consumer single- or multi-phase energy meters. With a tolerance of 5.0 %, the new resistor provides power companies with more accurate data to determine customer usage and to adjust billing terms. The WSMS2908 features a 5-terminal connection design and an all-welded construction that contributes to its superior electrical performance. The resistor offers low inductance values below 5 nH and a low thermal EMF of less than 3 µV/°C. The device is lead (Pb)-free and RoHS-compliant.

Vishay

www.vishay.com

Magnetic cylinder sensors

determine exact pneumatic piston position

Baumer's new generation of magnetic cylinder sensors has been designed for detecting the exact position of pistons in pneumatic cylinders. The new cylinder sensors available for T and C slots have been enhanced in performance and reliability. For instance, the sensitivity was optimized to produce a more precise and reproducible switching characteristic. Every sensor is tested upon its sensitivity prior to delivery. Additionally, upgraded materials are used to make the sensors more robust. The cables, for example, now consist of PUR, ensuring the trouble-free use of the cylinder sensors in oily environments. The increased robustness shows well especially with the salt water climate resistance of the sensors. The magnetic sensing technique, with no moving parts, enables maximum sensor product life by keeping wear down to a minimum. The MZTK 06P1013 line of sensors for T slots can be quickly installed as no cylinder disassembly is required, even if both front faces are equipped with installation panels. Thanks to the wide selection of mounting accessories, the sensor can be conveniently attached to all common cylinder types. All sensors are available in PNP or NPN as well as cable or flylead connector versions.

Baumer

www.baumer.com

4-port PoE power source manager

works completely autonomously with no host MCU

Microsemi Corporation has unveiled an Auto Mode 4-port Power-over-Ethernet (PoE) Power Source Equipment (PSE) Manager targeted at small and medium business (SMB) and small office/home office (SOHO) PoE applications. The PD69104A device is the most integrated and energy efficient 4-port PoE solution in the market, claims the IC vendor. The device works completely autonomously with no need for a host processor or MCU-based system management, allowing switch and router vendors to add PoE capabilities to their systems with minimal to no software development required. Although targeted at small systems, the PD69104A has an easy-to-use UART and I²C host interfaces that enable it to be used in larger systems where system-wide power management is performed by the host processor. The PD69104A features built-in 0.3 ohm MOSFETs and an internal DC/DC regulator. The device offers integrated dynamic power management for up to four ports and detects all pre-standard PoE devices.

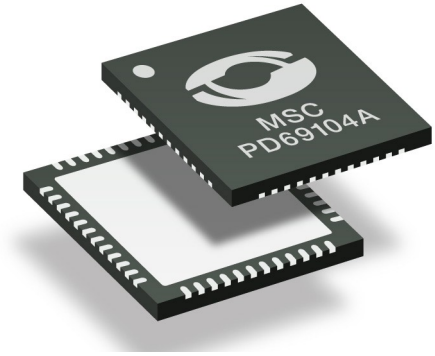

Microsemi

www.microsemi.com

DISTRIBUTION CORNER

Softcore evaluation SDK

for Altera's Nios II embedded MPU

Arrow Electronics has announced its BeMicro software development kit (SDK), enabling engineers to evaluate Altera's Nios II embedded processor and help simplify their embedded designs. The BeMicro SDK comes in a plug-and-play USB stick form factor and includes an FPGA-based board, software development tools and hands-on labs. The kit is fully supported with an open-source Eclipse-based integrated design environment, templates and board support package.

Arrow Electronics

www.arrow.com/bemicrosdk

Mouser signs Arduino

for open-source hardware

Mouser Electronics has teamed with Arduino to form a worldwide distribution partnership and resell open-source physical computing platform based on flexible, easy-to-use hardware, and software. This global distribution agreement gives Mouser a single-board computing solution with a 100,000-user-strong design community. In addition, it offers a cross- development platform centered on the Atmel AVR 8-bit microcontroller, especially well suited for expanded functionality by simply adding one of the various "shields" to the main Arduino board.

Mouser Electronics

www.mouser.com

RFMW expands TriQuint

product lines in stock for EMEA

RF and microwave components specialized distributor RFMW signed a distribution agreement with TriQuint Semiconductor to bring the company's products to Europe, the Middle East, and Africa. This agreement expands and builds upon the successful relationship that the two companies have had in North America since 2005 and in Israel since 2010. RFMW established sales offices in the U.K., Germany, and Italy at the beginning of 2011, and the TriQuint product line helps ensure a successful start in Europe. The company's goal is to expand TriQuint's design in opportunities.

RFMW

www.rfmw.com

Midé Technology worldwide

with Digi-Key disti agreement

Digi-Key Corporation and Midé Technology announced they have entered into an agreement in which Digi-Key will distribute Midé Technology products to customers worldwide. Midé's Volture piezoelectric energy harvesters are in stock and available for purchase on Digi-Key's global websites. To complement its energy harvesters, Digi-Key also stocks Midé's Slam Stick portable, rechargeable vibration data logger.

The Slam Stick helps customers identify if their application is suitable for energy harvesting.

Digi-Key

www.digikey.com

RS Components eTech mag iPad app

eTech magazine iPad app

Electronics and maintenance products distributor RS Components is continuing its push to bring innovative services to its customers with the launch of eTech for iPad. This iPad app provides a new way to display product and technology information, extending and enhancing the magazine content with additional detail, images, animation and video, making it a truly interactive customer experience. The app is available free of charge for download worldwide from the Apple store. eTech for iPad is an industry first and has been developed to go way beyond the basic page-turner digital magazine, taking full advantage of the device's off-line viewing capabilities. Written specifically for design engineers, technical managers and technicians working in the field of Electronics, eTech is an online and printed publication distributed globally. RS works with electronics design engineers and experts to provide information on new technologies for a wide variety of applications. eTech for iPad is free to download and subscribers will receive a notification for each issue.

RS Components

www.rs-components.com

Managing open source licensing for semiconductors

By Kamal Hassin

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

Increasingly, developers supplement custom coding with open source software, enabling rapid development at a reduced cost with a high degree of flexibility, security and stability. Open source opens up new opportunities that did not exist even a few years ago. As with every significant opportunity comes the need for responsible practices to ensure fairness and sustainability.

There are a number of approaches to license management, ranging from doing nothing to fully automated real-time scanning of software to detect and report license obligations. The cost of managing software license obligations is analogous to managing defects in the development process. It is well understood that the earlier a defect is identified and corrected, the less expensive it is, and the same is true for licensing obligations. Likewise, the earlier a development organization identifies "licensing bugs" that attract unacceptable license obligations, the less expensive it is to adjust the software to achieve licensing compliance.

A complete approach to assuring open source license compliance will generally include three major aspects: the definition of a licensing policy which must be met by all software projects; the auditing of all project software to detect any third party source code including open source that is unacceptable based on the licensing policy; and corrective processes to ensure that all released software conforms to the licensing policy.

The creation of the licensing policy is an important step as it forms the basis upon which decisions regarding acceptance of open source software will be made. The licensing policy should be defined in accordance with both the business goals of the organization as well as its engineering processes, and generally requires the involvement of business and engineering managers, as well as legal counsel. Software auditing is required to ensure production code conforms to the licensing policy, although the audit implementation can take a number of forms. Audits can range from ad hoc developer training and post-development cycle auditing to proactive automated approaches such as periodic and real time auditing.

Several approaches are available to address license compliance at different points in the development process.

Do nothing: this option ignores the compliance issue and carries the lowest up-front cost, but imposes the highest business risks and largest corrective costs as a product moves closer to launch.

Developer training: some companies consider developer training and project planning is sufficient enough. This however, can be an expensive labour-intensive option given the increasing diversity of software licenses, the high cost of developer training, and the constant churn within the development environment. With this option, compliance depends solely on busy developers and is prone to human error.

Post-development license audits: auditing late in the project lifecycle does not impact the development workflow and can be implemented manually or using automated software tools. This option does, however, lead to more expensive rework due to additional system retesting cycles.

Periodic assessment: licensing analysis during development allows for corrections along the way if license violations are detected. This type of analysis can be automated and tends to be less expensive than post-development assessment since changes and re-tests are always easier to undertake earlier rather than later in the cycle.

"The overall goal is to minimize the time and cost of correcting the final software release"

Real-time preventive assistance: the most proactive approach is to audit software in real time at the developer workstation. The development process is not disturbed and the cost of corrections is minimal as there is no impact to system integration and testing. This process can be automated and generally requires very little developer training.

Regardless of the type of auditing approach, the overall goal is to minimize the time and cost of correcting the final software release so that it meets all functional, quality and license compliance requirements. Each organization must consider their approach, balancing the short term cost of developer training and tools versus the potential longer term cost of post-release legal problems.

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

Kamal Hassin is director, R&D and product management at Protecode - www.protecode.com
He can be reached at khassin@protecode.com

EE Times europe

now available on iPAD

EE

free iPAD App

Stroke of Genius

Now Shipping the NEW Cobalt Series of Virtex-6 Modules!

Unleash your inner genius with Pentek's new Cobalt™ Virtex-6 XMC modules. Built to optimize high-performance, wideband, communications, SIGINT, radar, beamforming and telemetry applications, the high-speed modules include:

- A/Ds with sampling ranges from 10 MHz to 1 GHz
- 800 MHz D/As, DUCs & Multiband DDCs
- Dedicated memory for each I/O stream
- Intelligent chaining DMA engines
- Secondary serial gigabit interface
- Multichannel, multiboard synchronization
- ReadyFlow® Board Support Libraries
- GateFlow® FPGA Design Kit & Installed IP
- VPX, XMC, PCI, PCIe, cPCI, VME/VXS, rugged XMC

Now Available in VPX!

COTS & Rugged Versions

With extensive resources and new enhancements in DSP engines, block RAM, logic and clocking technology, the Virtex-6 family represents the industry's most advanced FPGAs.

Call 201-818-5900 or go to www.pentek.com/go/eetcobalt for your FREE *Putting FPGAs to Work in Software Radio Handbook*, technical datasheets and price quotations.

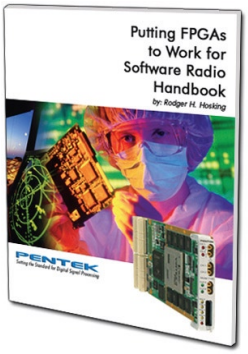

Cobalt

PENTEK
Setting the Standard for Digital Signal Processing

Pentek, Inc., One Park Way, Upper Saddle River, NJ 07458 Phone: 201.818.5900 Fax: 201.818.5904 e-mail:info@pentek.com www.pentek.com