October 2017
Volume 123
Issue 1976
£5.90

www.electronicsworld.co.uk

Electronics WORLD

THE ESSENTIAL ELECTRONICS ENGINEERING MAGAZINE

Multi-Protocol Industrial Ethernet Made Easy with RZ/N1

from Renesas Electronics

SPECIAL REPORT
EMBEDDED DESIGN:
- Interfaces
- Certification
- Programmable Logic
- VR and AR

ALSO IN THIS ISSUE:
- Elastomer contactors
- Battery pack manufacture

Trend
The 21st century farm requires drones

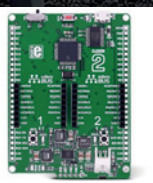

MCU column
Clicker-2 is better than one

Latest Products
Next generation 25W, 28V GaN HEMT delivers more gain

Realise your vision

We have all the parts you need to solve your technical challenges and bring your design projects to life.

Never get stuck again. Our extensive range of semiconductors, passives and interconnects has got you covered. With thousands of brands you know and trust, you can count on us for the best solution.

Find the right parts at

uk.rs-online.com

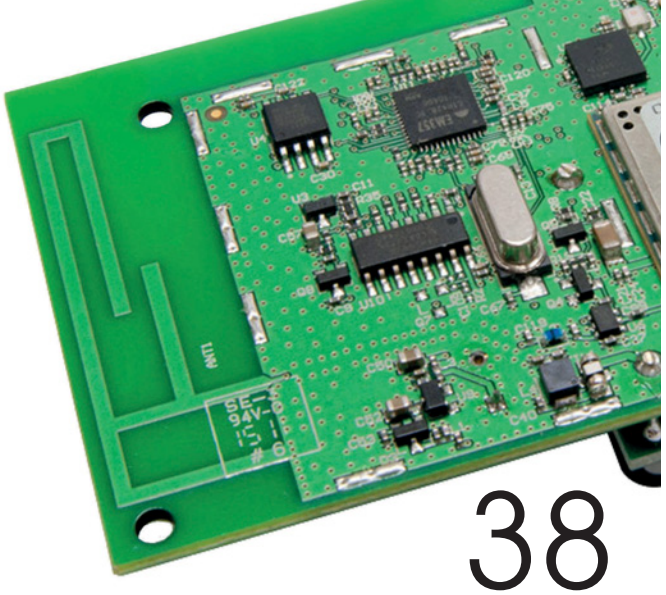

38

Cover supplied by RENESAS ELECTRONICS
More on pages 8-9

REGULARS

24

FEATURES

45

Don’t Leave Radar Performance Up In The Air.

Your systems keep passengers safe, trust us to do the rest. Wolfspeed, A Cree Company, has the broadest portfolio of S-band radar GaN power amplifier solutions.

Wolfspeed™ | A CREE COMPANY

Request a sample of the first 50 V MMIC power amplifiers specifically optimized for S-Band radar.

GO.WOLFSPEED.COM/RF-SAMPLE

THE 21ST CENTURY FARM REQUIRES DRONES

Agricultural practices need to continue to develop and become more efficient, if the findings of recent research reports are anything to go by. PwC states that agricultural consumption will increase by nearly 70% by 2050 due to the world's growing population – projected to hit nine billion people in the same year. Similarly, the World Bank predicts that we'll need to produce 50% more food by 2050 should global population continue to rise at its current pace.

Fortunately, technology is constantly advancing, its developments helping agriculture, too. Today, two important technologies of interest to agriculture are drones and autonomous vehicles. Drones have become a very lucrative market, with global revenues expected to increase by 34% to over £4.8m in 2017. US technology research firm Gartner has also predicted that drone production will jump by 39% this year compared to 2016.

Drones offer numerous benefits in agriculture:

- **Easy planting:** they can take the stress out of planting and looking after produce, thanks to new systems that reduce the costs of planting by as much as 85%. Drones can shoot pods of seeds and inject nutrients into the soil.
- **Easy irrigation:** to avoid wasting water around a farm, drones can be fitted with remote sensing equipment, either multispectral, hyperspectral or thermal. The technology will easily and quickly identify the driest sections of a field and help farmers allocate water resources more economically.
- **Spraying and monitoring crops:** Drones can effectively scan the ground and spray the correct amount of liquid needed for even coverage, yet reduce the amount of chemicals that leech into the groundwater.

With crop monitoring, time-series animations through the drones will display the exact development of a crop, and detail any inefficiencies with production. These insights were previously gained by satellite imagery, sadly only available on a limited basis, whereas monitoring through drones can be used whenever needed and at any time.

> Today, two important technologies of interest to agriculture are drones and autonomous vehicles

Autonomous Vehicles

As with drones, the market for autonomous vehicles is looking very bright. In fact, a comprehensive report by Business Insider Intelligence forecasts that eventually there will be some 10 million cars on the road, either semi-autonomous or fully autonomous. From a more general perspective, management consulting firm Bain has estimated that the global opportunity for assistive and autonomous technologies for the business-to-business market will be between $22bn and $26bn a year by 2025.

Steps have already been made to show how autonomous vehicles can be used in agriculture. For instance, a team of agricultural engineers from Harper Adams University in Shropshire have created a tractor that can drill, seed and spray land by remote steering. The same team – Johnathan Gill, Kit Franklin and Martin Abell – is now developing automated combine harvesters. Such vehicles will create a fleet in the hands of agricultural analysts, looking after several farming robots and monitoring the development of crops.

Meanwhile, in the Burgundy region of France, inventor Christophe Millot has succesfully built a vine-pruning robot. Developed to counteract a shortage in farm labour, the latest-generation model of the four-wheeled gadget includes six cameras, two arms and a tablet computer inside. The machine learns as it goes about its task of, say, trimming the vine or grass around it with a cut every five seconds.

So, farming drones and autonomous farm vehicles are certainly the future in agriculture, thanks to constant improvements to modern technologies.

By William Barne, Director at Lycetts, insurance brokers (www.lycetts.co.uk)

EDITOR:
Svetlana Josifovska
Tel: +44 (0)1732 883392
Email: svetlanaj@sjpbusinessmedia.com

SALES:
James Corner
Tel: +44 (0)20 7933 8999 | Mobile: +45 93 86 42 65
Email: jamesc@electronicsworld.co.uk

DESIGN: Tania King

PUBLISHER: Wayne Darroch

ISSN: 1365-4675

PRINTER: Buxton Press Ltd

SJP business media

2nd Floor,
52-54 Gracechurch Street,
London, EC3V 0EH

SUBSCRIPTIONS:
Subscription rates:
UK - 1 year digital only £53.00+VAT
UK - 1 year print and digital sub £68.00
UK - 2 year print and digital sub £109.00
UK - 3 year print and digital sub £143.00
International - 1 year digital only £53.00
International - 1 year print and digital sub £164.00
International - 2 year print and digital sub £290.00
International - 3 year print and digital sub £409.00
Tel/Fax +44 (0)1635 879361/868594
Email: electronicsworld@circdata.com
www.electronicsworld.co.uk/subscribe

Follow us on Twitter
@electrowo

Join us on LinkedIn

RENESAS AND COCORO ENABLE CARS TO READ DRIVER'S EMOTIONS AND RESPOND ACCORDINGLY

Renesas Electronics has developed a kit that determines driver's emotions to optimally manage the vehicle and its speed. The development kit is based around the Renesas R-Car system-on-a-chip (SoC) and cocoro SB's Emotion Engine (see box), an artificial sensibility and intelligence technology.

The cocoro SB's Emotion Engine relies on the speech of the driver to recognise the emotional state, such as confidence or uncertainty, and the Renesas R-Car SoC then displays this on a new user interface. The technology is based on artificial intelligence (AI) machine learning, which will enable the car to learn from conversations with the driver.

The driver's emotional state, facial expression and eyesight direction are combined with his or her vital signs to improve the car/driver interface, bringing closer the era of self-driving vehicles. For example, if the car recognises the driver's uneasy emotional state, even if he or she has verbally accepted the switchover to hands-free autonomous driving, the car can question the driver. Furthermore, understanding the driver's emotions enables the car to control its speed accordingly, especially in harsh conditions such as at night or in rain, snow, fog or icy conditions.

Renesas plans to release the development kit later this year.

The Emotion Engine's map used to determine driver's emotional state

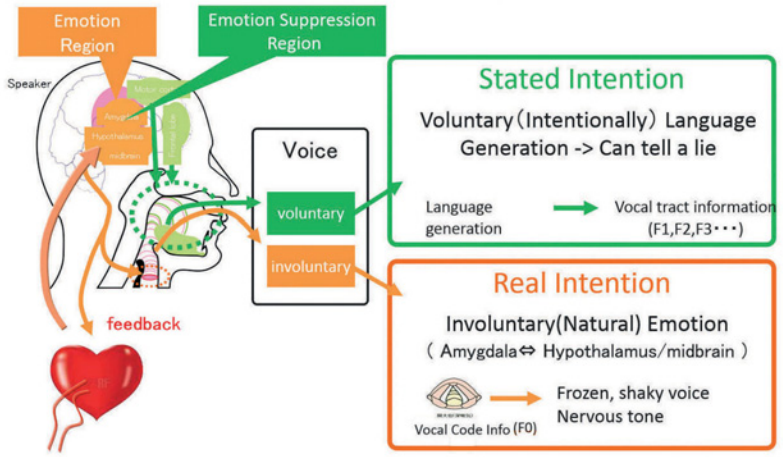

COCORO SB'S EMOTION ENGINE

Cocoro SB's Emotion Engine is an artificial sensibility and intelligence technology, consisting of two parts. The first is voice emotion recognition that analyses voice signals, whereas the second is emotion-generation that replicates emotions by forming a virtual hormone balance derived from various connected sensors.

The engine determines speaker emotion by charting a map, highlighted with colours, such as yellow for happy or red for sad, and at different degrees.

STUDY PROVES VIABILITY OF QUANTUM SATELLITE COMMUNICATIONS

Researchers in Canada have demonstrated the viability of a secure quantum communication via moving satellites. To ensure the tests were a valid proof of concept, the team at the Institute for Quantum Computing (IQC) and Department of Physics at the University of Waterloo, Ontario, designed their prototype receiver to match the size and operating environment of a micro satellite.

Free space links have been shown to work over ground at varying distances, both in stationary and moving tests. But despite losses due to geometric effects scaling quadratically with distance, atmospheric absorption and turbulence and the need for a clear line of sight limit terrestrial free-space transmissions to a few hundred kilometres. Satellite-based systems can extend quantum communication to a global scale.

"Quantum key distribution (QKD) establishes cryptographic keys between two distant parties in a way that is cryptanalytically unbreakable. Ground-based QKD systems use optical fibre links and are limited to distances of a few hundred kilometres due to absorption losses, which get significantly worse as the distance increases," said study lead Christopher Pugh.

To test their system, the team used a Twin Otter aircraft to pass their ground transmitting station at varying distances, achieving a quantum signal link and a secret key extraction for half of them.

"This is an extremely important step that took almost eight years of preparation. It finally demonstrates our technology is viable," said principle investigator Professor Thomas Jennewein. "We achieved optical links at angular rates similar to those of low-Earth-orbit satellites, and for some passes of the aircraft over the ground station, links were established within ten seconds of position data transmission. We saw link times of a few minutes and received quantum bit error rates typically between three and five precent, generating secure keys."

"We have proved the concept, and our results provide a blueprint for future satellite missions to build upon, just in time for a quantum satellite mission by the Canadian government," said Jennewein, referring to Waterloo's Quantum Encryption and Science Satellite (QEYSSat) initiative.

Professor Thomas Jennewein: "This is an extremely important step that took almost eight years of preparation"

RZ/N1 – MULTI-PROTOCOL INDUSTRIAL ETHERNET MADE EASY

By Ognen Basarovski, Renesas Electronics Europe

With the latest RZ/N1 family, Renesas aims to replace the existing multi-chip solutions in industrial controllers, industrial switches and operator terminals with a single, low-power-consumption SoC (System on Chip). Cast on a single die, an ARM® Cortex®-A7 application core and an ARM® Cortex®-M based real-time communication engine work hand in hand to provide real-time, deterministic communication. The innovative abstraction layer allows easy implementation and exchange of multiple industrial Ethernet protocol stacks.

R-IN Communication Engine – innovative support for multiple protocols

The automation industry is moving into a new era often referred as Smart Factory or Industrie 4.0. Future production facilities will be based upon cyber physical systems whose key trait is real time and deterministic communication. In order to facilitate this trend, Renesas developed and brought to market the R-IN Communication Engine.

The story of the R-IN Engine started in 2012 when Renesas introduced the R-IN32M3 series of industrial Ethernet communication SoCs with multi-protocol support. It includes a real-time capable switch, an ARM® Cortex®-M3 CPU as well as 4 accelerators that take part of the load off the CPU:

- The HW-RTOS supports the software execution of a real-time OS by handling some portion of tasks' processing, service calls and task scheduling in hardware.
- The CheckSum accelerator automatically calculates "on-the fly" the 4-byte checksum placed at the end of Ethernet frames.
- With the Header EnDec accelerator, the CPU has a fast and direct read and write access to all frame header information without latency.
- Finally, the Buffer management accelerator automatically controls the buffer allocation and release functions for the Ethernet processing in the hardware.

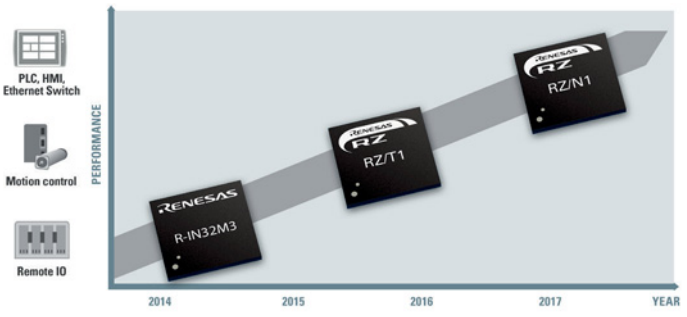

RZ/N1 – 3rd Generation R-IN Engine

Since the above functions are directly executed in silicon, there is no involvement of the Cortex®-M3 CPU in their processing. This saves CPU performance and power dissipation during high Ethernet traffic operations, which enables low power operation (~300mW). As an example, the R-IN Engine with accelerators has only 30% CPU load at UDP communication rates of 95 Mbps. This leaves enough performance bandwidth for a simple IO application on the same core. Furthermore, the accelerators improve the overall system timing with their deterministic and low-jitter processing speed. This is paramount for high-speed real time networks, as it enables better quality in an isochronous network environment.

Both the initial market success as well as the positive feedback from customers resulted in several R-IN Engine implementations. Three years after the R-IN32M3 was released, the motion control SoC RZ/T1 was introduced to the market and finally in 2017, the RZ/N1 family. The RZ/N1 is a logical extension of the Renesas industrial automation portfolio and gives a scalable and software compatible platform for industrial automation applications.

R-IN Engine Family – Performance Meets Power Efficiency

With the introduction of the new RZ/N1 family, Renesas brings more performance to typical industrial communication applications. Both the application and the communication block are integrated on a single chip, which makes it easier to achieve the real-time performance and determinism.

The key focus of RZ/N1 is on demanding applications in factory automation. With its' multi-core architecture, it offers enough performance for controllers, HMI solutions or industrial switches with industrial Ethernet support. Thanks to its sophisticated embedded 5-port switch, RZ/N1 brings cutting-edge performance in industrial networking, enabling features required for a TSN end-point like traffic shaping, time synchronization and frame classification/prioritization.

The RZ/N1 family consists of 3 products Both RZ/N1D and RZ/N1S have one application block and one communication block integrated on the same die. These blocks are connected via a communication API (Application Programming Interface) based on GOAL (Generic Open Abstraction Layer). GOAL provides an API for the application development, making therefore both the core-2-core communication and the R-IN Engine transparent for the application developer.

The application block of the RZ/N1D SoC features a dual ARM® Cortex®-A7 and is suitable for applications with higher performance requirements. The product is available in two packages, 400 BGA and 324 BGA. The application block of the RZ/N1S SoC features a single ARM® Cortex®-A7 and is available in 324 BGA and 196 BGA packages. The communication block of both RZ/N1D and RZ/N1S is nearly identical and depending on the product option, offers 3 Ethernet ports in the smaller package or 5 Ethernet ports in

the larger package. The RZ/N1L has no dedicated application block, as it only features the Cortex®-M3 CPU. The idea is to offer a simple communication device for customers who want to bring network functionality to legacy devices. The RZ/N1L is available in 196 BGA package.

Innovative Software – all protocols under one umbrella

Only a few years ago, it was sufficient for a device in an automation network to measure the data and forward it to the controller via Fieldbus protocol. This was usually handled in the hardware. As the industrial Ethernet technology started taking over the communication on the factory floor, the embedded firmware got more complex. Not only that the Ethernet based protocols are now handled in software, also new hardware components such as the embedded Ethernet switches need to be supported by the software. This increased system complexity severely impacts the time-to-market requirements. Because of the growing pressure to develop more functionality in less time, the software has to be reused across different hardware platforms.

There is a variety of industrial Ethernet technologies on the market today. Technologies like EtherCAT, EtherNet/IP, PROFINET and several others have a long history, and significant functional differences between them. Some of them even require dedicated hardware support, which is why it is quite challenging to achieve a high cohesion. It is no surprise if for example, an application written for PROFINET needs to be completely rewritten for an EtherCAT implementation. Furthermore, with some hardware platforms available on the market, different protocols run on completely different ecosystems, which make the change between protocols very time, money and resource consuming.

The GOAL middleware for the RZ/N1 family successfully overcomes these challenges. It provides a software framework for development of low coupled software design. The GOAL API give interface to all components of the system and it is stable regardless of the underlying hardware. The interface to the underlying system itself is done by the GOAL Target API which is fully transparent to the programmer.

Key benefit of the GOAL platform is the easy exchange of the Ethernet communication technologies with minimal impact on the application software. All the communication stacks (protocols) that are ported over GOAL share the same abstraction layer and they are connected with the application software over a unified software interface. There is of course, a certain dependence of the application software on the protocol stacks, so it has to be amended for each different communication protocol. However, this modification is minimal and quite straightforward. Another benefit of the GOAL is that applications written on top of the GOAL API can be used across different hardware components! Whether it is the dual A7 or the single A7, or even the M3 alone, the GOAL simplifies the migration from SoC to SoC.

Besides the Industrial Communication part, the GOAL is equipped with a TCP/IP stack independent web server and

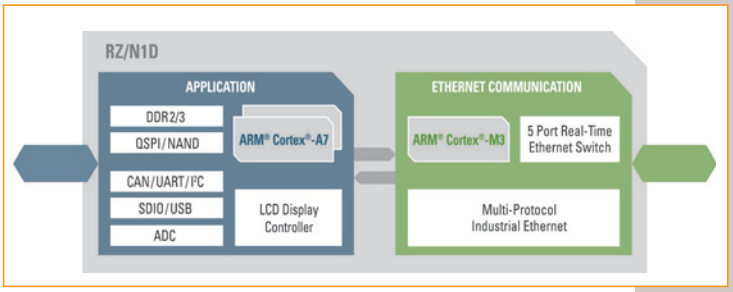

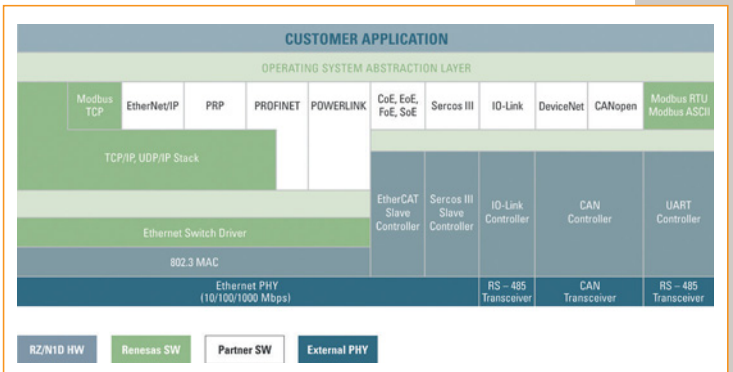

OSAL/GOAL - Easy migration between protocols and SoC

comes with supporting network protocols such as SNMP and RSTP. This makes GOAL the ideal basis for network aware devices such as managed switches, routers and gateways. On top of the TCP/IP stack GOAL provides an independent stable API for TCP/IP based applications with support for a variety of TCP/IP stacks regardless whether the stacks support the BSD socket interface or not.

Companies with limited know-how in industrial Ethernet can simply regard the R-IN Engine within the RZ/N1 as an industrial Ethernet black-box. Together with our business partners like Cannon Automata, Net Module, port GmbH, and TMG TE who are market leaders in the field of industrial communication software, we have already ported numerous protocol stacks to the R-IN Engine. The GOAL API is available directly on the Cortex®-A7 side, where the application software resides.

The rest of the RZ/N1 ecosystem depends on the product version. The RZ/N1D application block runs under the Linux OS or ThreadX. The RZ/N1S, on the other hand supports ThreadX and VxWorks. The R-IN Engine in the communication block of all three devices runs with uItron libraries.

A solution kit based on RZ/N1D and GOAL is already available and can be ordered directly at the Renesas web shop or at one of the distribution partners. The RZ/N1D CPU Board (RZ/N1D-DB) provides access to the basic SoC interfaces. For evaluation of the full SoC functionality, an optional expansion board CONNECT IT! RZ/N1 Expansion board (RZ/N1-EB) can be used. ●

The Solution Kit includes an extensive software package. A variety of sample applications are also delivered together with the GOAL SW package explaining the basic functionalities of the run time system. For more information please visit the RZ/N1 web page:: ***www.renesas.com/rzn***

Clicker-2 is better than one

BY **LUCIO DI JASIO**, MCU8 BUSINESS DEVELOPMENT MANAGER AT MICROCHIP TECHNOLOGY

It's been barely a year since I first wrote about my growing appetite for mikroBUS slots. I am referring to those little connectors where we can insert daughter boards (or shields) for rapid prototyping on the latest generation of PIC evaluation boards. With 300+ Click boards to choose from, how could I be satisfied using only one at a time?

Last year, when there were only 200 Click boards in the catalogue, I'd discovered the Clicker-2 board series. As the name implies, there are two mikroBUS slots, and many models offer a choice of microcontrollers, from 8-, to 16- and 32-bit architectures. Unfortunately, I realised that the only 8-bit PIC model was based on the PIC18F87J50 family (which is some 10 years old) that offered many traditional peripherals and an integrated USB port but none of the more modern Core Independent Peripherals (CIP) I was interested in. More importantly, it was too old to be supported by the latest rapid development tools such as MPLAB Code Configurator (MCC), an addiction of mine. In an article then, I described how I could circumvent this limitation by leveraging the similarities with the more-recent PIC18F45K50 models.

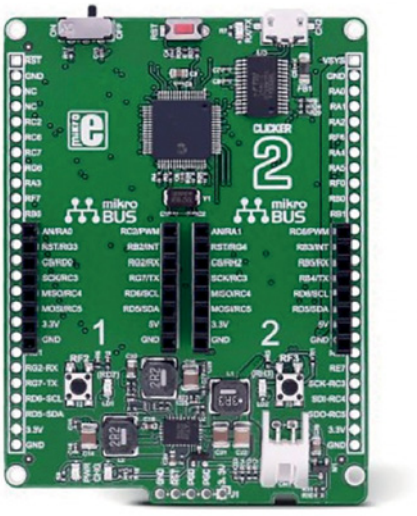

Figure 1: MikroElektronika Clicker-2 board

This year, it seems both Microchip and MikroElektronika heard my prayers: The Curiosity series of development boards now has several new models with two mikroBUS connectors for PIC24, PIC32MM and PIC32MX. A new Curiosity High Pin-Count (HPC) has also been added to support the many 28- and 40-pin models of the 8-bit PIC microcontrollers.

MikroElektronika refreshed the Clicker-2 series and launched a new PIC18 board dedicated to the 'K40 series.

For those of you who don't closely follow the PIC saga, the PIC18 K40 family is a relatively recent addition to the portfolio that marks the adoption of all CIPs (developed exclusively on the PIC16 families over the past seven years) by the PIC18 architecture. This puts the two main 8-bit PIC architectures on an unprecedented path to convergence. You should compare the PIC16F188xx and the PIC18xxK40 models side by side to get a sense of the level of compatibility achieved.

Clickers' Curiosity

The MikroElektronika Clicker-2 boards differ from Microchip Curiosity boards in many ways. While most of the Curiosity boards provide a socket for DIP packages to accommodate many similarly-packaged PIC models, Clicker-2 boards have a fixed (often surface-mounted) target processor.

Clicker-2 boards are generally smaller and offer a Li-ion battery connector and an on-board USB charger. This makes Clicker-2 boards more suitable for actual applications and product prototypes than for pure code development and debugging. In fact, looking at Figure 1, you will notice the two outer rows of contacts that are compatible with the sister Mikromedia family of boards (featuring TFT displays) and available for further expansion with more shields, mikroBUS slot expansion boards and even robotics development platforms.

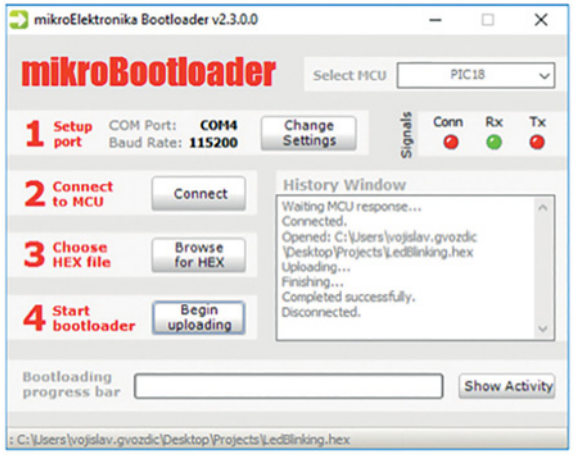

Figure 2: mikroBoot loader

But the biggest difference in the developer experience is perhaps related to the programming and debugging interface. Microchip Curiosity boards offer the equivalent of a PICKIT 3 programmer and debugger on board. Whereas MikroElektronika Clicker-2 boards are designed to connect to the company's own mikroProg programmer/debugger; in its absence, they can use a bootloader via a virtual serial port.

I will confess that the virtual serial port is an excellent feature (sadly, still missing on the Curiosity side), but perhaps not sufficient to balance the lack of fully-integrated programmer/debugger.

Arduino users might perhaps be of a different opinion, but in my experience a bootloader (no matter how well-designed; see Figure 2) is always a little more intrusive and less effective than a real programmer.

Drag 'n' Drop

The only other environment where I have been happy to trade a real programmer with a (virtual) serial port is with the MPLAB Xpress evaluation boards. Here, though I got not only a serial connection, but also a Mass Storage USB connection (MSD), which practically means the entire act of programming the target PIC microcontroller is reduced to a simple drag-and-drop gesture. When Xpress boards are connected to a personal computer they essentially appear to the host operating system (Windows, Mac or Linux) as a hard drive. Copying any properly-formatted .hex file to the (virtual) drive triggers a programming sequence; see Figure 3.

This is great when paired with a Cloud toolchain (as with MPLAB XPRESS IDE), so that without installing any special bridge or driver (MSD interfaces are universal), it is possible to access the target safely even from a browser. But, as many discovered, this can be effective also from a more traditional desktop IDE, such as MPLAB X, once a simple added "build step" is configured in the project configuration; see Figure 4.

The XPRESS-Loader is fast, much faster than traditional PIC programming tools, and requires no manual intervention, as opposed to most bootloaders; just hit "Build" and off you go.

Porting XPRESS To The Clicker-2

The magic of the XPRESS programmer is due to a small USB device, another PIC microcontroller in fact, mounted next to the target microcontroller. In the first incarnation, this was a PIC18 of the 'K50 series, chosen because of the large amount of Flash and RAM (32K/2K) in its relatively small (28-pin QFN) package, but also the integration of the USB peripheral with Active Clock Tuning (ACT). This is a feature that allows the chip to use the internal oscillator (saving precious space and cost) and yet achieve the tight clock tolerances required by full-speed USB.

In more recent incarnations, the XPRESS loader code has been optimised further to fit the PIC16F1454 device (14KB/1K) in a 16-pin UQFN package (4x4).

The source code for the XPRESS-loader has been published on GitHub under a friendly Apache2 (open source) license, and is available to all for use beyond pure development tool purposes. Any developer might want to simply place an XPRESS loader on his own prototype/product board in place of a plane USB/serial interface chip while possibly reducing cost and board space. In fact, that was exactly my proposal to MikroElektronika: let's "upgrade" the latest Clicker-2 for the PIC18 K40 by replacing the standard issue (FTDI) serial interface with a modified XPRESS loader!

As often is the case, their reply was rather enthusiastic and we quickly got to hacking away new board connections and a new firmware branch for the little PIC16F1454.

Details, Details...

There are, of course, many small details to be considered when intervening in an existing product that's part of a larger series and has (a bit of) its own legacy.

This starts with presenting a backward-compatible virtual serial port (USB-CDC), which incidentally, despite being one of the

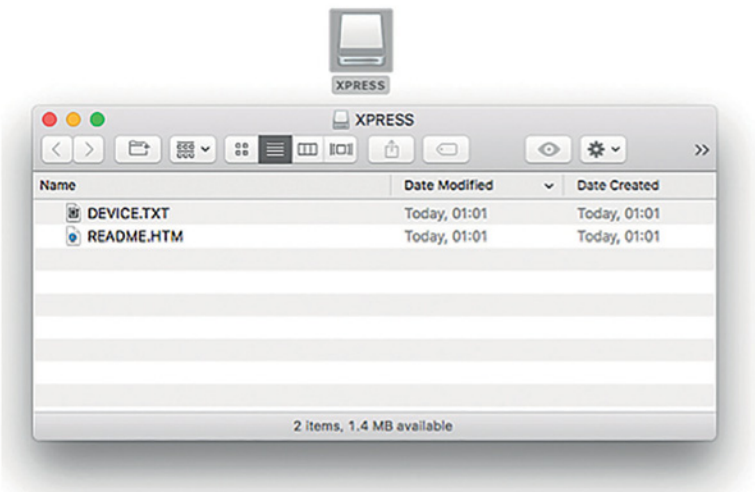

Figure 3: A virtual drive

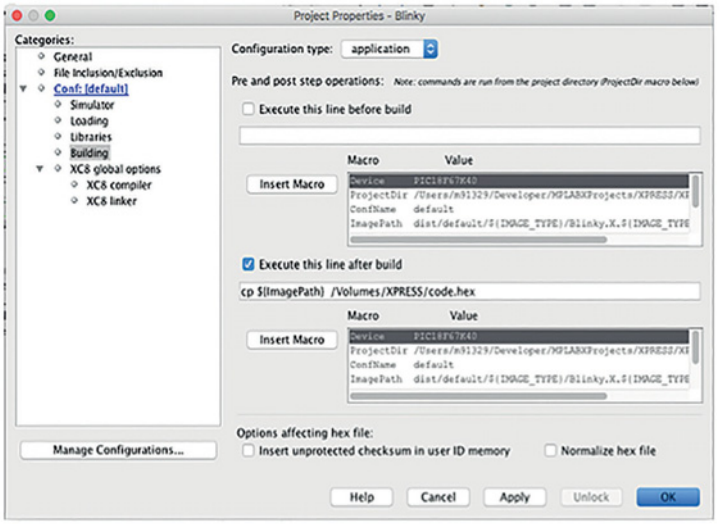

Figure 4: MPLAB X build settings

most common interfaces, is not available in Windows without installing a (signed) device driver. Mac and Linux users are spared such a nuisance.

Microchip development boards use a common driver (connected to a standard company VID/PID), but MikroElektronika boards have their own. Luckily, this is an easy parameter to configure, leaving this decision to the very last minute.

The other USB interface, USB-MSD, is more universal but comes with its own set of (minor) limitations. It is meant to be truly PC-centric; so much so, in fact, that the storage device has virtually no way to communicate changes back to the host PC other than demanding to be "ejected". So, the reset button of the Clicker-2 board must be repurposed to act as a target reset, but also an eject button for the XPRESS virtual drive.

Further, although you might have not realised this before, there are two alternative ways of programming an 8-bit PIC microcontroller: high-voltage mode (legacy, 9-12V) and low-voltage mode (LVP, 3-5V only). This refers to the sequence of steps required to bring the target device into programming mode so the non-volatile memory and configuration bits can be erased and reprogrammed. PICkit3 and other universal programmers typically fall back to the old legacy mode.

But LVP-mode is clearly the most convenient, and is used by the XPRESS loader to avoid having to add a boost circuit (for space and cost reasons) on board. In fact, you will find in Listing 1 an example LVP-mode enter sequence extracted directly from the XPRESS-loader source code. Here you can see how the simple manipulation of three GPIOs: ICSP_CLK, ICSP_DAT and ICSP_slaveReset is all it takes.

```
void ICSP_init(void )
{
 __delay_us(1);
 ICSP_ANSEL_DAT = DIGITAL_PIN;
 ICSP_TRIS_DAT = INPUT_PIN;
 ICSP_CLK = 0;
 ICSP_ANSEL_CLK = DIGITAL_PIN;
 ICSP_TRIS_CLK = OUTPUT_PIN;
}

void ICSP_sendCmd(uint8_t b)
{
 uint8_t i;
 ICSP_TRIS_DAT = OUTPUT_PIN;
 for( i=0; i<8; i++ ){ // 8-bit commands
 if ( (b & 0x80) > 0 )
 ICSP_DAT = 1; // Msb first
 else
 ICSP_DAT = 0;
 ICSP_CLK = 1;
 b <<= 1; // shift left
 __delay_us(1);
 ICSP_CLK = 0;
 __delay_us(1);
 }
 __delay_us(1);
}

void ICSP_signature(void){
 ICSP_slaveReset(); // MCLR output => Vil (GND)
 __delay_ms(10);
 ICSP_sendCmd('M');
 ICSP_sendCmd('C');
 ICSP_sendCmd('H');
 ICSP_sendCmd('P');
 __delay_ms(5);
}
```

Listing 1: Low voltage programming with ICSP

As a side effect, the direct connection of the target reset pin (ICSP_slaveReset) to an XPRESS loader GPIO (PIC16F1454), means the high-voltage (legacy) programming mode won't be possible anymore as the GPIO (clamping) diodes (on the Xpress device) would limit the voltage applied to Vdd.

Mind you, this would be the case only when the user applies an external programmer/debugger to the board, effectively bypassing the XPRESS programmer, but it is a legitimate case and it must be supported.

In the proposed implementation, we should avoid adding external components (as done on the Xpress evaluation board, for example), but instead simply advise the user to configure the programmer/debugger of choice to use LVP as well. ●

www.congatec.com
info@congatec.com
Phone: +49 (991) 2700-0

congatec

The Server on a Module.

conga-B7XD

First COM Express Type 7 full compliant module.

Server-On-Module.

- Intel® Xeon® D CPUs with up to 16 cores & 24 MB cache
- 32x PCI Express lanes, 2x 10 GBit Ethernet
- Smallest server board ever - 125 x 95 mm²
- Optimized for demanding real time applications
- Ideal for rugged micro servers for industrial environment

intel IoT Solutions Alliance

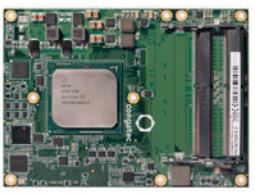

Embedded in your **success**.

Applying the ESP8266 processor in IoT applications

BY **DR DOGAN IBRAHIM**, PROFESSOR AT NEAR EAST UNIVERSITY, CYPRUS

The Internet of Things (IoT) is a commonly encountered term today; it refers to devices communicating over the Internet with users and each other. One example is the intelligent fridge that alerts when a product has perished or needs replacing.

The IoT concept is not knew – it dates back several decades; however, the technology was not ready then. Nowadays, IoT applications are touted in as diverse fields as home and factory automation, health care, transportation, environmental studies, farming and many more.

IoT Applications

An IoT device typically contains wireless hardware and a communications application with which users interact with mobile phones, tablets or laptop computers.

The success of an IoT application depends on many hardware-related factors, such as size, cost, power consumption, security, robustness and reliability. Being portable and largely battery-operated, in IoT devices power consumption is perhaps most important, since they are expected to operate non-stop and consume as little power as possible.

Several communication technologies can be used for IoT applications, the most common of which are Wi-Fi, Bluetooth, RF and ZigBee.

Figure 1: The ESP-01 development board

● Wi-Fi

Wi-Fi is very secure, reliable and readily available in homes, factories and offices. In addition, most portable devices, such as mobile phones, tablets and laptops, are equipped with Wi-Fi.

Wi-Fi's main disadvantage however is that, generally, it is power demanding, making long-term operation with batteries not possible.

● Bluetooth

Bluetooth is also secure, but not as widely available as Wi-Fi, and its power consumption is also high – although there are currently low-energy Bluetooth (BLE) devices available, developed primarily for the IoT market.

On the flipside, Bluetooth cannot be accessed outside its coverage area.

● RF

Using RF in IoT applications is not very common for several reasons: first, RF consumes a lot of power; second, the data is not encrypted and it depends on the developer to use a suitable encryption algorithm; and, third, RF-based devices can't be accessed from a mobile device.

New MCUs

Several chip manufacturers are currently developing low-cost, low-power microcontrollers with embedded Wi-Fi and/or Bluetooth capabilities.

Examples include the recently-developed, highly popular and inexpensive ESP8266 and ESP32 microcontrollers, both from the Shanghai-based manufacturer Espressif Systems. Both processors consume very little power, with ESP8266 containing a built-in Wi-Fi module, and ESP32 both, Wi-Fi and Bluetooth modules.

ESP8266 and ESP32 are complete microcontroller systems with input-output ports that can be used in two formats: as standalone microcontrollers with communication capabilities, or as host processors where only their communication capabilities

are used by larger microcontroller systems.

In this article, we'll briefly discuss the ESP8266 microcontroller's details and give an example of its use in a standalone Wi-Fi-based application.

The ESP8266 Microcontroller

The ESP8266 is a low-cost microcontroller chip with Wi-Fi capability and a full TCP/IP stack that can be used with or without the MCU to access a Wi-Fi network. The ESP8266 can easily be interfaced to a development system such as the Arduino, where it can provide Wi-Fi capability to the host.

The ESP8266 mounts on a small PCB, available as a small development board. Many ESP8266-based development boards are available, normally named as ESP-xx. Figure 1 shows the first-version ESP-01 development board.

The ESP8266's basic features include:

- 32-bit RISC CPU (Tensilica Xtensa LX106);
- Operation at 80-160MHz;
- 64kB program memory;
- 96kB data memory;
- External QSPI flash memory (up to 16MB);
- IEEE 802.11 b/g/n Wi-Fi support;
- WEP or WPA/WPA2 authentication;
- 1 x 10-bit analogue-to-digital converter (ADC);
- Up to 16 GPIOs;
- Shared SDIO, SPI and I2C interface;
- Integrated TCP/IP protocol stack;
- UART support;
- 3.3V operation with less than 10mW power consumption.

The development of a project using the basic ESP-xx development boards requires a serial USB-to-TTL interface module to provide USB capability to the board. In addition, the ESP-xx boards offer only 2 GPIOs, and an external +3.3V power supply must be connected to the board.

Some manufacturers offer complete ESP8266-based development boards with USB interfaces and built-in +3.3V regulators. These include the NodeMCU (LoLin), SparkFun EDP8266 Thing, Olimex ESP8266-EVB, Adafruit Huzzah ESP8266, HiLetgo Mini NodeMCU, In-Circuit ESP-ADC, and others.

NodeMCU

NodeMCU (LoLin) is one of the most commonly used ESP8266 development boards because of its full functionality, low cost and ease of use (see Figure 2). It incorporates an on-board USB-to-TTL adapter, Micro USB socket, many GPIOs, large program and data memories, PWM, I2C, ADC, +3.3V voltage regulator and an UART interface.

NodeMCU (LoLin) is shipped with no firmware, which is left to the developer to install, based on the software to be used for programming the chip. The options are:

- AT command processor;
- Lua;

Figure 2: The NodeMCU (LoLin) development board

- C++ (via Arduino IDE);
- microPython.

The AT command processor was developed to send commands to modems and similar communication equipment for configuration, and send and receive data. Firmware can be uploaded to the ESP8266 chip to accept AT commands.

Lua is a powerful, efficient, lightweight, high level, free, scripting language which supports object-orientated, procedural, as well as data-driven programming. Lua is an interpretive language, making it ideal for rapid learning and prototyping.

C++ is a powerful object-orientated high-level language used mainly in large applications on laptops and desktop computers. Programming the ESP8266 microcontroller using C++ is through the Arduino IDE, where the firmware can be uploaded to the chip.

MicroPython is a subset of the powerful high-level language Python, used in almost all universities and technical colleges around the world in introductory programming courses. Python supports a large library of functions – one of its biggest advantages that differentiate it from other programming languages.

Python is a very popular programming language, available on most microcontroller systems. MicroPython firmware must be loaded to the ESP8266 chip before the chip can be programmed with this language.

Example

A project is given here to show how the NodeMCU (LoLin) development board can be used in a Wi-Fi application to send UDP packets to a remote server. Although both the UDP and TCP protocols are supported, the example is based only on UDP.

Figure 3 shows the system setup, where the NodeMCU (LoLin)

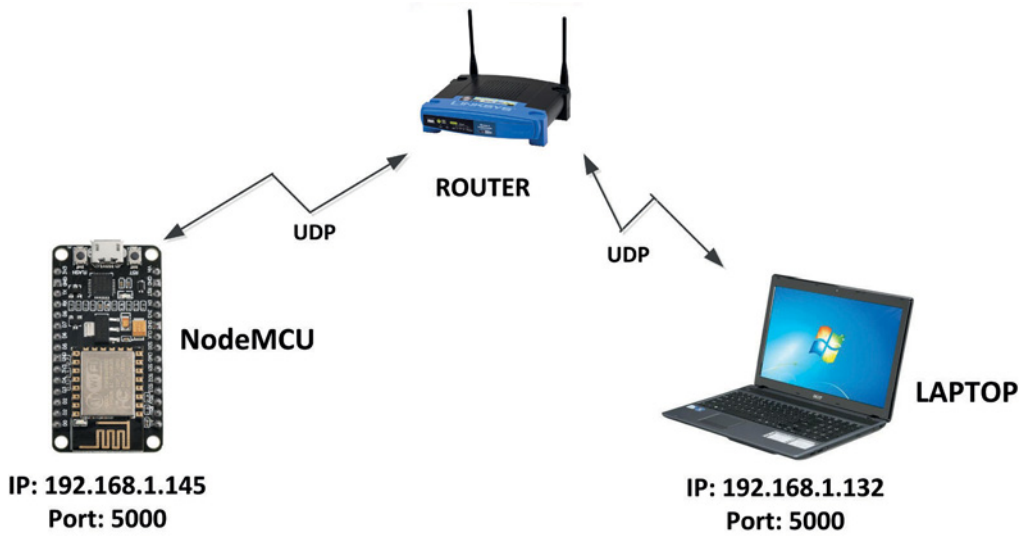

Figure 3: System setup

sends the HELLO message from the NodeMCU to a server over port 2000, using the UDP protocol.

UDP is a connectionless protocol, so there's no need to make a connection to the destination node before a packet can be sent. The communication between a server and client is as follows (note that both server and client can send and receive data packets from each other):

Server

- Define the client IP address and port number;
- Create a socket;
- Construct the data packet;
- Bind the socket to the local port;
- Receive data from the client;
- Send data to the client;
- Close the socket.

Client

- Define the server IP address and port number;
- Create a socket;
- Send data to the server;
- Receive data from the server;
- Close the socket.

As shown in Figure 3, the communication goes through the router. The following IP addresses and port number are used by server and client:

Server (NodeMCU) IP address: 192.168.1.145 port: 5000
Client (PC) IP address: 192.168.1.132 port: 5000

In this example, the NodeMCU program is developed using the microPython programming language. The complete program is shown in Listing 1.

At the beginning of the program, the socket library is imported and the port number defined. The message to be sent is stored in a string variable called Message. A socket is created with the name sock, and sock.bind statement is used to bind the socket to the local port. Then, data is received from the client using statement sock.recvfrom and displayed on the screen. The message is then sent to the client using the statement sock.sendto. The program terminates after closing the socket with the sock.close statement.

In this example, the AGG Com Port Data Emulator software package (www.aggsoft.com) was used on a laptop to send and receive UDP packets. This software is available free of charge with a time-limited license.

```
#-------------------------------------------------------------------------
------------------
#                    UDP PROGRAM FOR THE ESP8266
#
============================
#
# This is a MicroPython UDP program for the ESP8266 processor. The
# program sends the message "Hello From the NodeMCU" to a client.
# The IP addresses and ports of the server and client are:
#
# Server IP address: 192.168.1.145 Port: 5000
# Client IP address: 192.168.1.132 Port: 5000
#
# The data received from the client is displayed on the
```

```
screen
#
#
# Program:                udp.py
# Date:             July, 2017
# Programmer:   Dogan Ibrahim
#-------------------------------------------------------------------------
----------------
import socket

UDP = ("", 5000)
Message = "Hello From the NodeMCU"
                #
                # Create a socket and bind the socket
to local port
                #
sock = socket.socket(socket.AF_INET, socket.SOCK_
DGRAM)
sock.bind(UDP)
# Receive data from the client and display it
#
data,addr = sock.recvfrom(1024)
print (data)
#
# Send data to the client
#
sock.sendto(Message, addr)
#
# Close the socket
#
sock.close()
```

Listing 1: microPyton program

This example shows how easy it is to establish Wi-Fi communication and send and receive data packets with the help of the NodeMCU development board that costs only a few pounds.

It is likely that, because of their low power requirements, the ESP8266-type development boards will be used in many future IoT applications. ●

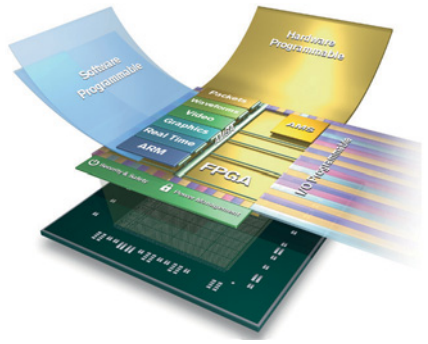

Cloud computing acceleration for high-performance image recognition

BY **GILES PECKHAM** AND **ADAM TAYLOR** OF XILINX

Advanced computer vision applications are becoming increasingly pervasive and are used to enable autonomous-driving modes of today's cars, as well as in augmented reality, security surveillance systems, healthcare, industrial inspection equipment, robotics and more. Users' expectations are rising, as familiarity inevitably brings demands for higher performance, such as faster response times, greater accuracy, or recognition of extra objects or features.

Recognition and classification of images uses deep machine learning techniques such as convolutional neural networks. Before they can be deployed, these networks need to be trained for the application.

The previous article in this series described tools for building high-performing embedded application-processing engines, capable of running deep neural networks trained in the Cloud for use in an autonomous edge device. This approach is suitable for systems like self-driving vehicles, where low latency is critically important and a reliable, high-bandwidth connection to the Cloud may not always be available.

Other applications, such as security surveillance or medical imaging, may be less demanding of outright speed and instead require complex analyses, extremely high accuracy to help specialists make informed decisions and the ability to allow multiple parties to access the results. In such situations, where compute-intensive tasks and flexible storage and access policies are required, the image-processing application can be more effectively and economically hosted in the Cloud.

Hardware-Accelerated Cloud Computing

Compared with the embedded vision systems discussed previously, hosting the application-processing algorithms in the Cloud presents a different set of challenges. It is here that compute-intensive deep machine learning, data analytics and image processing are implemented. Often the applications are required to stream processed data in near real-time and without dropouts.

Cloud data centres are increasingly unable to fulfil the demands of today's most intensive processing workloads using conventional CPU-based processing alone. Some have adopted FPGA-accelerated computing to achieve the throughput needed for these workloads and others, such as complex data analytics, H.265 encoding and SQL functions.

Historically, an FPGA has been teamed with a CPU to provide acceleration, but a new model is emerging based on arrays of FPGAs such as Xilinx Virtex UltraScale+ devices. These arrays deliver extremely high peak compute capability, with the added advantage of rapid runtime reconfigurability to repeatedly re-optimise for subsequent workloads.

Stack Streamlines FPGA Development

To fully harness the capabilities provided by programmable logic, an ecosystem is needed that enables development using current industry-standard frameworks and libraries. The Xilinx Reconfigurable Acceleration Stack (RAS) answers this need, by streamlining FPGA-hardware creation, application development and integration.

Hyperscale data centres can use these tools to jump-start development: several major operators are currently working with Xilinx to boost performance and service agility by introducing FPGA acceleration in their server farms, making extreme high-performance compute capacity available to customers as a web service.

Like the reVISION stack for embedded vision development, described in the previous article in this series, the RAS leverages High Level Synthesis (HLS) for efficient development of programmable logic in C/C++/OpenCL and System C. This HLS capability is then combined with library support for industry-standard frameworks and libraries such as OpenCV, OpenVX, Caffe, FFmpeg and SQL, creating an ecosystem that can be extended in the future to add support for new frameworks and standards as they are introduced.

Also like the reVISION stack, the RAS is organised in

three distinct layers to address hardware, application and provisioning challenges. The lowest layer of the stack, the platform layer, is concerned with the hardware platform comprising the selected FPGA or SoC upon which the remainder of the stack is to be implemented. The RAS includes a single slot PCIe half-length, full-height, development board and a reference design, created specifically to support machine learning and other computationally-intensive applications like video transcoding and data analytics.

The second level of the RAS is the application layer. This uses the Vivado Design Suite and SDAccel development environment, leveraging HLS to implement the application. SDAccel contains an architecturally optimising compiler for FPGA acceleration, which enables up to 25 times better performance per watt compared with typical processing platforms made up of conventional x86 server CPUs and/or graphics processing units (GPUs). The environment is featured to deliver CPU/GPU-like development and runtime experiences by ensuring easy application optimisation, providing CPU/GPU-like on-demand loadable compute units, maintaining consistency throughout program transitions and application execution, and handling the sharing of FPGA accelerators across multiple applications.

For machine learning applications, DNN (deep neural networking) and GEMM (general matrix multiplication) libraries are available on the Caffe framework, as shown in Figure 1. Libraries for other frameworks such as deep-learning TensorFlow, Torch and Theano are expected to be added later. It is worth noting at this point that the scope of RAS is not limited to machine vision or deep learning: as Figure 1 shows, other libraries are included that support MPEG processing using FFmpeg as well as data movers and compute kernels for data analytics on the SQL framework.

The third level of the RAS is the provisioning layer, which uses OpenStack to enable integration within the data centre. OpenStack is a free, open-source software platform made up of multiple components for managing and controlling resources such as processing, storage and networking equipment from multiple vendors.

Performance Boost With Power Savings

By using the RAS to streamline the development of Cloud-class FPGA-based computing, a significant increase in computing capability can be achieved, compared with processing on conventional CPUs. Image-processing algorithms can be accelerated by as much as 40 times, while deep machine learning can be up to 11 times faster. In addition, hardware requirements are reduced, which lowers power consumption, resulting in a dramatic increase in performance per watt. Moreover, the FPGA-based engine has the important advantage of being reconfigurable and so can be quickly and repeatedly re-optimised for different types of algorithms as they are called to be executed.

Conclusion

Automatic image analysis and object-recognition applications can benefit from the increased performance and reduced power consumption of highly optimised, reconfigurable FPGA-based processing engines. Whether the application is to run on an embedded system or in the Cloud, using an acceleration stack enables developers to overcome design and integration challenges, reduce time to market and maximise overall performance. ●

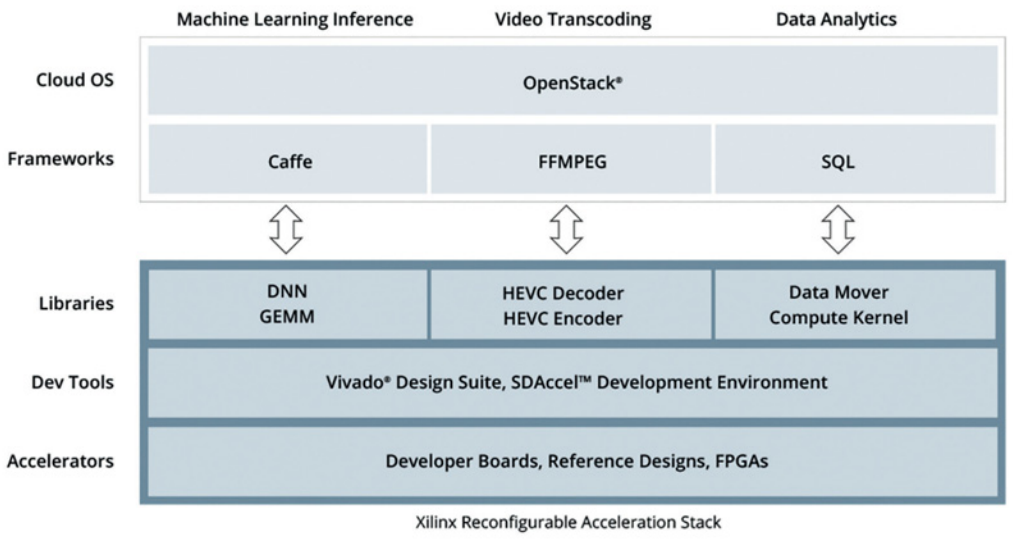

Figure 1: RAS addresses hardware development, application development and integration

Four ways to improve oscillator phase-noise performance

BY **TOMMY REED**, VP OF TECHNOLOGY AT BLILEY TECHNOLOGIES

If you've been following our articles, you probably have a great understanding of phase noise. So far, we've covered its basics, how to measure it using a crystal oscillator and even how to improve it. Now we'll go a step further, by assessing how to deal with it in different environments, including high levels of vibrations.

Bad Vibrations

Acceleration sensitivity is a crystal oscillator's inherent sensitivity to external forces applied in any direction. Quartz oscillators provide the heartbeat of the electronics we've all come to rely on. When voltage is applied to quartz, it begins to vibrate; but if vibration is applied, the quartz will produce a voltage. This voltage shows up as phase noise, and that's a problem; see Figure 1.

The magnitude of this phase noise or frequency change is directly proportional to the amount of applied force or acceleration – the higher the force, the greater the frequency instability, the greater the noise.

Frequency instability due to a crystal's acceleration sensitivity impacts many aspects of oscillator performance, including ADEV (short term stability), phase noise performance and RMS phase jitter.

Digital communication and RF systems' performance can be impacted by this vibration-induced phase noise, as an error that manifests itself as increase in BER (bit error rate). For example, in Doppler radar systems this shows as a notable signal degradation due to vibration-induced energy.

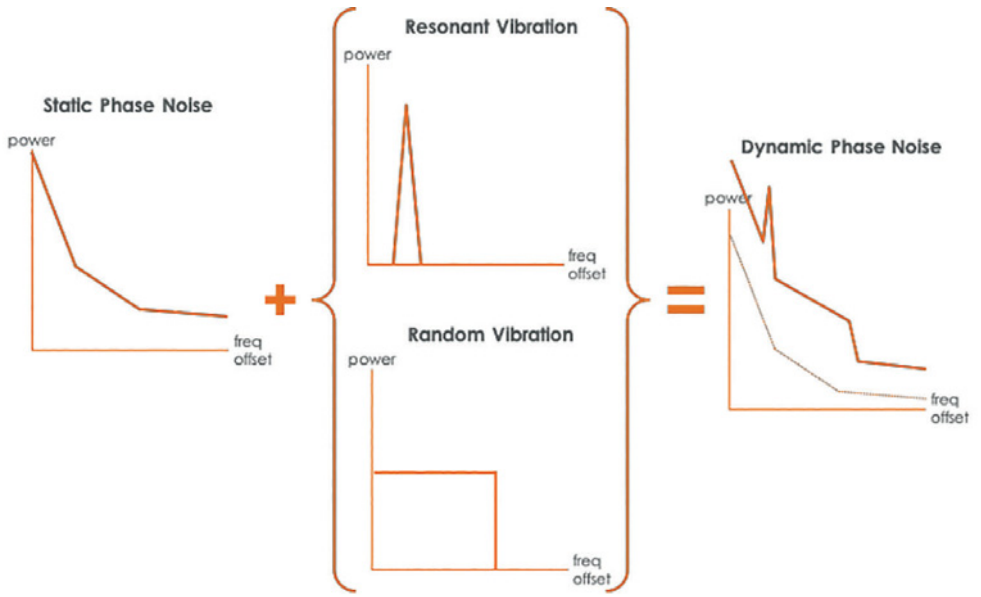

Figure 1: Vibration-induced phase noise

So, What Can Be Done?

All quartz crystals exhibit some level of inherent sensitivity to vibration. There are many different options to minimise the effect of acceleration sensitivity on customer systems. Here are a few options quartz crystal manufactures can apply:

- Proper resonator selection (for example: type of crystal cut used);
- Vibration screening on the resonator level (screen for best g-sensitivity);
- Passive isolation (mechanical dampening);
- Electronic compensation.

1. Resonator Selection

Proper crystal selection is by far the simplest approach to achieving better performance in dynamic conditions.

The cut of the crystal can make a large impact on system performance as a whole. AT cut crystals are widely used in many types of reference clocks. Research has shown that AT cut crystals can perform almost as good as SC cut crystals; however, SC crystals' overall gamma vector is still far better than that of AT cut crystals.

The crystal mounting structure also plays an important role on how the crystal performs under vibration. Various crystal packages employ different mounting structures; however, we have found that HC35 (TO-5) and HC37 (TO-8) packages with a 4-point mounting structure offer the best vibration resistance.

Securing the crystal in four places drastically changes the mount resonance of the package, reducing its overall impact on performance. Many vendors today still use 2-point mounts in HC45-style holders, along with a 2-point cantilever-type mount found in most surface-mount VCXO and TCXO packages.

2. Vibration Screening

After the crystal has been selected, the manufacturer then screens it for acceleration sensitivity on all axes, which helps reduce the cost of the final product.

At Bliley, we screen by inducing either a sinusoidal or random force with a known magnitude. We then measure the degradation using a phase-noise analyser. If a sinusoidal waveform at a given frequency is used, we measure the spur induced by the single tone. From that data we can calculate the resonator's acceleration sensitivity.

If a random spectrum is applied to the crystal, we can essentially measure the phase noise across the vibration bandwidth and calculate the g-sensitivity based on the given PSD (power spectral density).

3. Mechanical Isolation

Passive isolation can have a profound impact on the acceleration sensitivity of a reference clock. If the input vibration is reduced, so too will be the degradation.

Passive isolation does have its disadvantage that systems can get large, depending on what's being isolated.

The resonant frequency of the platform also can be a big concern. Typical mounts can have a transmissibility of 3.5-4, which means the force input into the system will be amplified. Care should be taken to not exceed the system's maximum displacement, or isolator damage can occur.

This problem leads to low natural-frequency isolation schemes. In theory a sub-1Hz isolation structure is ideal, but in reality these systems are very large due to the dampening mechanism and sway space needed. This is where electronic compensation comes in.

4. Electronic Compensation

By employing electronic compensation we can minimise the size of the passive isolation system while still reaping the benefits of the passive mount; see Figure 2. Electronic compensation also minimises the effect of the isolating structure's resonant frequency.

Electronic compensation offers superior results and can achieve levels of vibration insensitivity two orders of magnitude better than standard crystals.

Typical compensation schemes we use can cover vibration out to approximately 500Hz, while the isolation structure starts dampening at approximately 120Hz. This wide crossover greatly improves the performance of the crystal and, essentially, provides customers with a vibration-immune reference signal.

Figure 3 shows the theoretical performance benchmark of the options discussed here, generated in reference to a representative operational vibration specification.

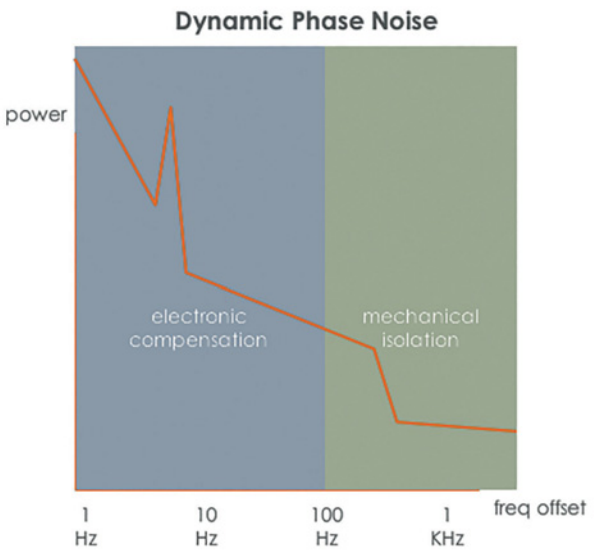

Figure 2: Phase noise, compensation and isolation

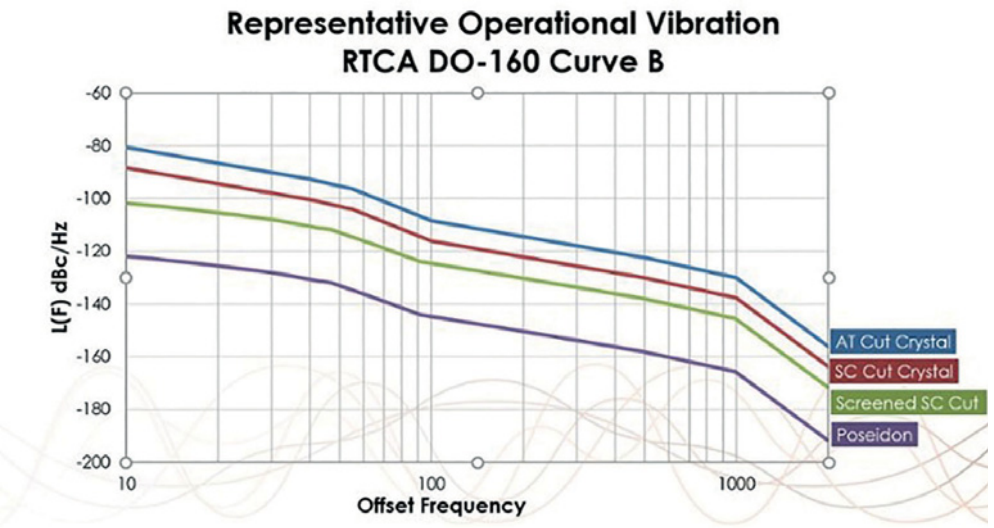

Figure 3: Theoretical performance benchmarks

Congratulations!

You made it! You're no longer just a young grasshopper when it comes to phase noise. Have fun applying your new phase-noise-ninja skills in your engineering life (and maybe showing off a bit).

Thank you to everyone that followed all our phase noise articles. We hope you learned a tonne. ●

Seamless Authentication and Encryption

CEC1702 Hardware Cryptography-Enabled Microcontroller

The CEC1702 MCU provides easy-to-use, seamless authentication and encryption capabilities for connected applications.

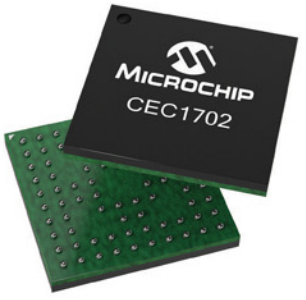

Product Features

- Low power
- Powerful, programmable 32-bit microcontroller
- Easy-to-use hardware cryptographic suite including encryption, authentication, private/public key capabilities
- Customer programming

Differentiators

- Security supervisor
- Can replace existing microcontroller
- Add authentication with little additional cost
- Can use as a standalone microcontroller or as a security coprocessor

microchip DIRECT
www.microchipdirect.com

MICROCHIP

www.microchip.com/CEC1702

VIRTUAL REALITY: NOT JUST FOR FUN

MOST OF US ARE FAMILIAR WITH THE FUN FACTOR VIRTUAL OR AUGMENTED REALITY BRINGS TO ENTERTAINMENT; BUT, HOW DO THESE ALTERNATIVE REALITIES WORK FOR INDUSTRY? **JASON HILL**, EXECUTIVE PARTNER AT REPLY, EXPLAINS

Virtual reality (VR) is a natural evolution of the media, which over the years have become very close to the consumer – from cinema, to television, even the smartphone, to a point that users can 'physically' enter the content. VR provides a far more visceral experience than other media, however, where user engagement is so strong that it's difficult to differentiate it from the real world, since movement and exploration are allowed, creating a complete immersion within the realm. On the other hand, with augmented reality (AR), the real world becomes physical support for the virtual content added on top of it.

> VR facilitates active hands-on learning, proven far more effective than conventional tools

The rapid adoption of these technologies is further supported by market analysis from several firms. In 2017, the consumer market of VR is valued at $4.8bn, making it a growing market for hardware and software. The analysis firm Superdata expects the revenues from XR software development to surpass those of hardware in 2020 with a global value of $20bn.

The numbers are also growing fast for the enterprise market. Companies are increasingly interested in using these technologies in their business processes, not just for marketing purposes. Nasdaq GlobeNewswire valued this market at $600m in 2016, projecting growth to $9bn in 2021.

Industry Support

The benefits are obvious for the consumer, but how do enterprises benefit from VR, AR and mixed reality?

Decades of experimental use and the more recent progress of digital technology have made it possible to create a favourable environment for widespread use of virtual and augmented reality technologies in the professional sectors. Today these technologies represent a new asset in support of decision-making processes, product design, design review processes, staff training and on-site maintenance.

The use of virtual and augmented reality in the enterprise environment can bring considerable advantages to a company, not only in terms of time and cost optimisation, but also in improvements of employee working conditions.

Here are five main fields in which VR can bring considerable advantages:

1. Training And Security

VR allows companies to train people for complex and dangerous procedures in an effective way. Industrial machinery and automated production systems are becoming increasingly more sophisticated and complex, requiring suitable preparation. This typically consists of specialist training courses (in many cases conducted on

actual systems), in a wide range of situations that can't easily be reproduced. All this significantly drives up the costs.

By using VR, however, a person can be immersed in a simulated environment and given the required preparation and training through specific operational scenarios – anywhere and in complete safety. Advanced communication technologies also enable remote team work, allowing several users to collaborate on tasks in the virtual world.

VR facilitates active hands-on learning, proven far more effective than conventional tools. In addition, being analytical in nature, the VR system helps identify defects and correct operational processes.

2. Complex Products Presentation

It's not always possible to present a prototype in an effective way, especially for complex machinery. VR allows a simple and effective presentation of products, both prototypes and finished ones, through 'viewers'. The internal mechanism of the machine is recreated in a virtual environment, so the producer can show how it works down to the smallest details, dismantling and reassembling different pieces at the customer's request. This function is also valuable for trade fairs, where often machinery can enter a pavilion but for safety reasons can't be turned on.

3. Design Review

Design review is one of the milestones in product development, where the requirements of a project are evaluated to verify the outcomes of previous activities, and identify issues, before committing to further work. The moment of design review, a time when all stakeholders evaluate and comment on a project, is crucial, since it generally enables the next stages of design and production.

VR technologies allow the product to be viewed during design, to real scale and based on natural interaction, even before a prototype is created. This optimises the follow-up stages and anticipates future problems, such as ergonomics, interferences and accessibility, reducing the impact on product development costs. In addition, collaborative VR tools simultaneously immerse all stakeholders in the same environment to evaluate the project together.

VR's benefits in everyday applications

4. In-Field Assistance

Technical assistance in the field is a critical factor for carrying out operations involving installation, maintenance, fine-tuning or repair of machines and systems in an industrial environment. In fact, these operations often require companies to send highly qualified staff to support local technicians, resulting in considerable waste of resources in terms of time and cost.

Thanks to the latest AR technologies these processes can be radically streamlined; using specific mobile devices or holographic viewers allows users in the field to superimpose digital information, such as infographics, images, videos, 3D models, diagrams, manuals and operating parameters, onto the real context. This means that, upon arriving in the area requiring intervention, technicians can autonomously and in real time consult in an augmented-reality environment all the basic technical information needed to perform the planned activities. In addition, instruments such as HoloLens allow quick and effective communication with headquarters: if the complexity of the operation requires intervention of a specialist, the local technician can request support through a remote assistance session in augmented reality. Here, an Internet-connected user can follow the operator in the field first-hand and guide them

VR brings benefits in industrial and medical applications and training

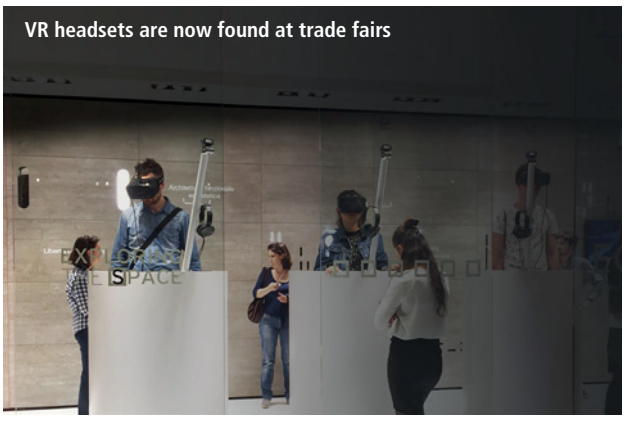

VR headsets are now found at trade fairs

vocally and with information given in AR. All operations can also be documented and archived through videos and images captured in real time during these intervention sessions.

Thanks to new methods of AR-enabled support in the field, technical assistance can be streamlined, with benefits including reduced intervention time, logistics and travel costs.

5. Customer Engagement

These media facilitate and enhance customer engagement, thanks to a new kind of experience that improves the human senses. Virtual and augmented reality can be used for marketing and educational purposes, analysing which content engages the customer most to reach the business goals. These instruments are valuable in this field because they take the client to a completely new virtual world.

Example

Reply collaborated with a leading company in energy and automation technologies to apply VR in training. The firm wanted to improve its training for employees dealing with repair and maintenance of electrical machinery. The training needed to be independent of actual equipment and in a safe environment to teach the employees how to handle potentially dangerous equipment.

Reply developed an e-learning platform, dedicated to the training of personnel operating the machinery. It consisted of two interactive modules: one online and based on the interactive Unity 3D engine technology, and the other featuring a virtual environment thanks to the HTC Vive VR headset. In both modules trainees could experience the exact maintenance operations used on actual systems. With the interactive 3D simulation, users gained an accurate and effective understanding of the tasks in complete safety.

Potentially dangerous tasks are fully simulated in the virtual environment, made possible through web-based VR software developed by Reply, in which all information is analytically collected to examine the person's potential qualifications for the role. The VR system leaves a numerical trail of everything implemented in the training session, providing benefits in terms of training and research and development. The software detects the difficulties experienced by multiple individuals, providing a report that shows the relevant department the changes needed to improve the tool's usability. ●

REPLY'S WAY

Every modern technology takes time to become habitual, and VR is in its ramp-up phase. Reply believes in this technology and has worked and invested in it since 2011. It developed tools, laboratories, dedicated infrastructure and expertise in the virtual and augmented reality realm, with the goal of developing innovative services, capable of offering high added value with sustainable costs/benefits for the companies involved.

Today, Reply's offer of VR services ranges from digital entertainment to professional applications; in fact, it uses the entire range of currently-available VR and AR technologies to develop projects designed to support enterprise core processes. Transverse apps have been produced over the last five years on several hardware visualisation platforms, from large-format, immersive systems such as the multi-wall CAVE (Cave Automatic Virtual Environment), to portable stereoscopic and holographic projection systems, 3D TVs, tablets and smartphones, up to the latest immersive visors.

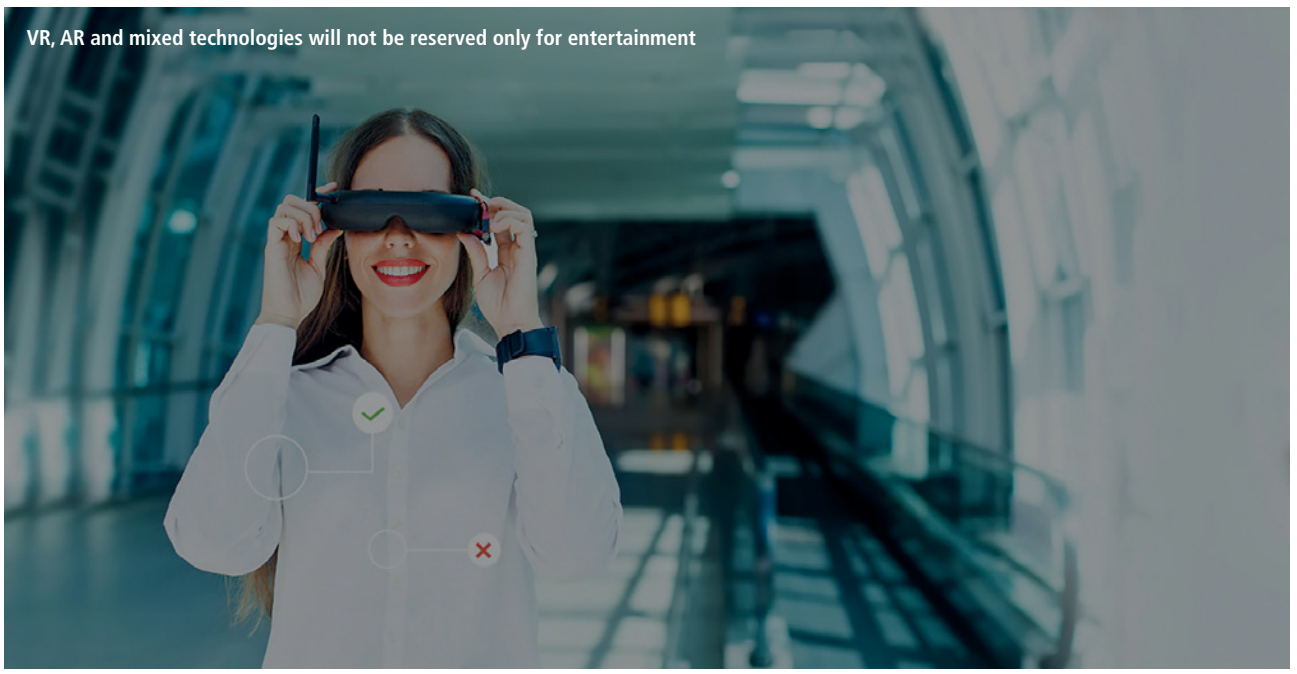

VR, AR and mixed technologies will not be reserved only for entertainment

OLSON®

CELEBRATING 55 YEARS
OF BRITISH MANUFACTURING

POWER DISTRIBUTION UNITS

YOU CAN NOW PURCHASE OUR PRODUCTS ONLINE

olsondirect.co.uk

SINCE 1961

Manufactured in the UK

OLSON®

+44 (0)20 8905 7273 | sales@olson.co.uk | www.olsondirect.co.uk

CHOOSING THE RIGHT INTERFACE FOR THE APPLICATION

MARC STACKLER FROM TELEDYNE E2V EXPLAINS THE DIFFERENCES BETWEEN PARALLEL AND SERIAL INTERFACES AND THE APPLICATIONS THEY BEST SERVE

High-speed data converters are nowadays found almost anywhere, ranging from communications (ground and satellite-based), high-energy physics (accelerators) and defence (electronic warfare, radar jamming) to industry (mobile phone test lines, tank container monitoring), test and measurement (spectrum analysers, mass spectrometers), earth observation (synthetic aperture radar), and more. There's clear demand for data converters that can deliver improved system performance and enhanced capabilities.

In turn, however, this increasingly involves choosing either a parallel or a serial interface. For example, there's currently a trend for using a full software defined radio (SDR) system, which requires large bandwidth and has an interface between an FPGA (Field Programmable Gate Array) and a data converter (a result of the SDR architecture being directly linked to the data converter's sampling speed).

> Parallel interfaces are now less common than serial interfaces, which have increased in popularity due to their improved bandwidth and speed

There are two ways to do this interfacing at high speed – LVDS (Low-Voltage Differential Signalling) parallel interface, or a serial interface. Each has pros and cons, depending on the application it serves, and they work in different ways. To date, parallel interfaces have been the traditional method, with serial interfaces appearing in the past ten years.

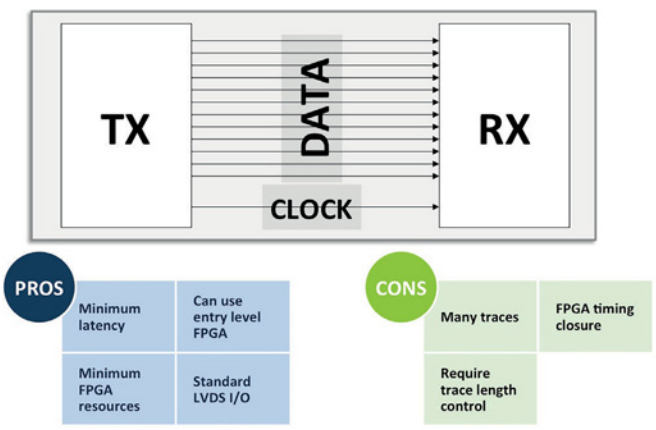

Figure 1: Pros and cons of a parallel interface

Parallel Vs Serial

Parallel interfaces are simple to implement; they work by using several lanes to transmit data simultaneously, with an additional lane to transmit the clock between transmitter and receiver. This is straightforward when it comes to PCB and firmware design, as each bit of the sample has its own path, transmitted at each digital clock cycle. Since the receiver has access to these clock signals, it easily recovers the data.

This works very well for low data rates, but at higher rates, a greater number of lanes are needed, complicating the interface design. Importantly, at speeds above 1GHz, parameters with negligible effect at lower speeds begin to limit performance. This is why serial interfaces began appearing ten years ago, and why they are now the preferred option in many applications.

Serial interfaces have a much simpler layout, with one lane. In a parallel interface, to transmit 10Gbps for example, at 1Gbps, 11 links are needed – 10 lanes between components A and B, plus an extra lane for the clock. In a 10Gbps serial interface however, only one lane is required to transmit 10Gbps, also saving on PCB space.

In a serial interface, each link has a high-speed transceiver, including a serialiser on the transmitter side and a deserialiser at the receiver end. Here, the clock does not need to be sent on a separate path, since it is extracted from the data during reception, via the Clock and Data Recovery system (CDR). This is what allows serial interfaces to work at much faster speeds than parallel interfaces. For example today, parallel interfaces in FPGAs are mostly limited to 1.6Gbps, whereas a high-speed serial transceiver in a serial interface can exceed 32Gbps.

Simplifying Data Processing

Efficiency is key for a serial interface, to ensure no data is lost during transmission. For this, a suitable protocol is needed and Teledyne e2v developed the open-source protocol, called

ESIstream, to simplify the serial data processing.

ESIstream improves latency and lessens the resource requirement, with small data overhead. It works using two-stage encoding and a two-overhead-bit process; the encoding is built with a scrambling process, followed by a disparity process. Without a data encoding/decoding protocol, there will be transmission Bit Error Rate (BER) and de-synchronisation, due partly to the CDR and partly to the AC coupling interface.

In a serial interface, the function of CDR is to recover the clock from the data stream on the receiver side – since it is not transmitted separately. This means the recovered clock has been subject to the same timing effects as the data – up to the CDR stage; hence, these effects cancel each other out, allowing for much higher speeds.

CDR enables the data rate per serial link to be increased, saving PCB space and simplifying layout. For CDR to do this, it must see 'edges' in the data to identify the clock. This is also the reason a long series of successive 1s and 0s must be avoided in a serial interface.

However, CDR can cause BER and synchronisation loss, as mentioned above. BER and synchronisation loss also occur because the interface between transceiver and receiver must be AC coupled; at high speeds, even a small difference in the common-mode can contribute to BER degradation. The AC coupling means that if the data stream is not DC-balanced, the AC coupling capacitor will become loaded and corrupt the data seen by the receiver.

The ESIstream protocol prevents both, long series of 0s and 1s and unbalanced DC transmission, thus improving BER and preventing synchronisation loss. This protocol is important since, if BER and synchronisation loss do occur, data is lost and the links need to be restarted, losing even more data in the process.

One of the benefits of ESIstream is that it does the expected function of a serial interface but more efficiently; it thus requires fewer resources within the FPGA than other, less-efficient protocols. Although the resources needed by the protocol are generally small enough not to matter, it can complicate the closing of the timing within the FPGA when the application requires more resources or high-speed digital design. The protocol is also used to digitally align multiple serial lanes, which means there's no need to precisely match trace lengths for an optimum serial interface.

The Right Interface For The Application

When it comes to choosing the right interface to the application, parallel interfaces have several benefits: they are easy to use at low data rates, they are cost-effective and have short latency. Parallel interfaces also use minimum FPGA resources, so entry-level FPGAs can be used. Their short latency is because, digitally, parallel interfaces are much simpler – they do not need encoding/decoding or transceiver stages. This short latency is an important advantage when it comes to applications such as radar jamming in electronic warfare, where a few nanoseconds can mean the difference between being spotted or remaining invisible to enemy radar systems.

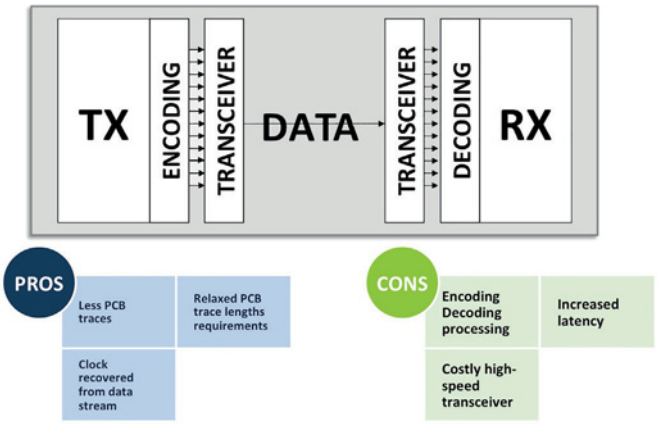

Figure 2: Pros and cons of a serial interface

Anti-collision radar is another application where the parallel interface is ideal, since detecting an object quickly is vital. A serial interface would be less effective for anti-collision radar since with a longer latency it takes longer to detect the object, leaving less reaction time. As a result, if the application requires a short latency and/or low data transmission rate, parallel interfaces are better suited.

However, parallel interfaces are now less common than serial interfaces, which increased in popularity due to their improved bandwidth and speed. They are not limited in terms of data rates per lane as with parallel interfaces, and the clock is recovered from the data stream, relaxing PCB and trace length matching requirements. As a result, more applications are using serial interfaces, and the trend is likely to increase further with rising developments in optical and modulation capabilities.

It's worth noting that serial interfaces can be expensive – for applications around 12Gbps a serial interface is cheaper, but over 12Gbps tuning is required to correct for very small effects. This can increase the cost, but overall a serial interface is still easier to use than a parallel one.

Both parallel and serial interfaces could be considered for low-cost or low-speed applications, depending on factors such as development time, re-use of already-developed subsystems and other requirements. When deciding which interface to choose, it's worth doing the research and asking the experts, to make sure you make the right choice.

Technology behind the interfaces will continue to evolve. Equally, as key advances in modulated serial interfaces continue and we approach the next step in high-speed data transmission, more problems will arise that will need solving. ●

HOP ON THE AVALON BUS: A SIMPLIFIED FPGA INTERFACE

BY **NOE QUINTERO** FROM LINEAR TECHNOLOGY, NOW PART OF ANALOG DEVICES

Many modern FPGA designs use some form of embedded processor for control. A typical solution involves a soft embedded processor such as NIOS; another solution is to use a SoC that includes a built-in hard processor. Figure 1 shows a typical Altera FPGA system that contains the processor and a mix of peripherals connected via the Avalon Memory-Mapped (MM) bus. These processors greatly simplify the end application, but they require strong programing background and knowledge of complicated toolchains, which can hinder debug – especially if a hardware engineer needs a simple way to read and write to the peripherals, without involving the software engineer.

This design idea uses Altera's SPI Slave to Avalon MM Bridge to simplify using the Avalon bus. There are two advantages for using this technique: it does not compromise the original system design, and the bridge can co-exist with the embedded processor.

For the system shown in Figure 1, the SPI-Avalon MM bridge would allow the engineer to directly control the frequency of the

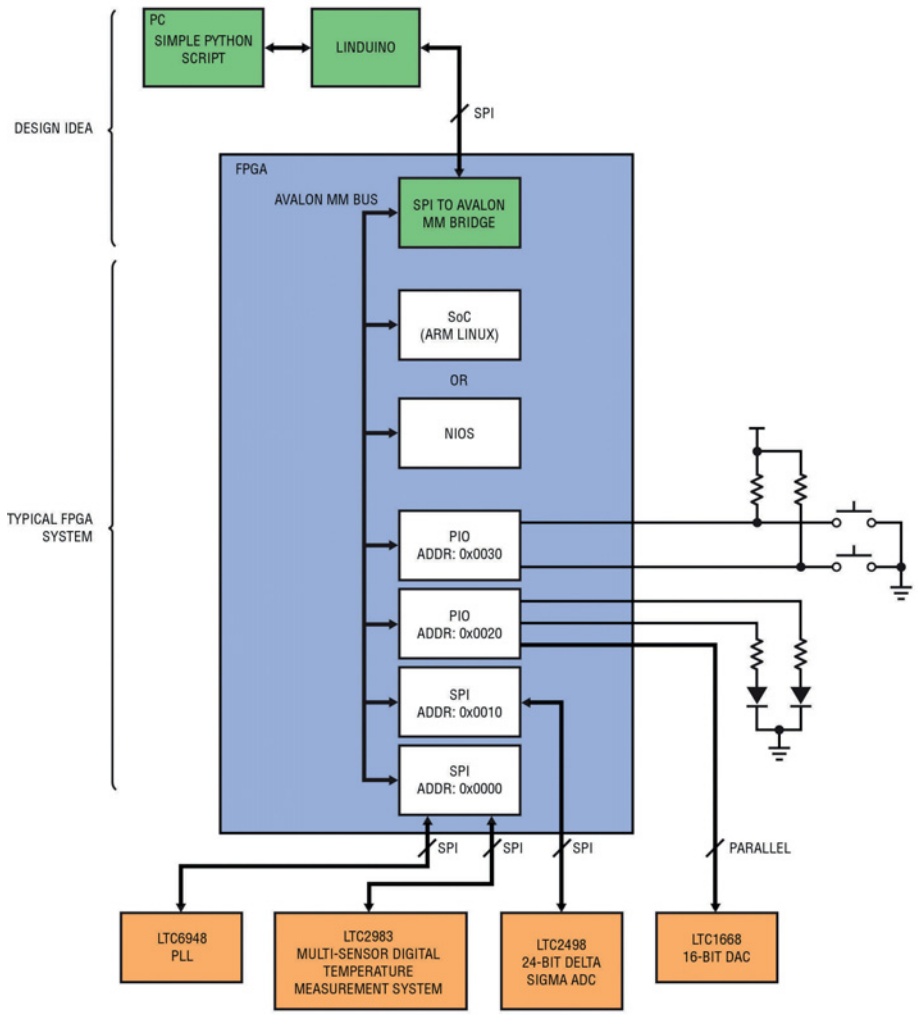

Figure 1: Typical Altera FPGA system connected via Avalon Memory-Mapped (MM) bus

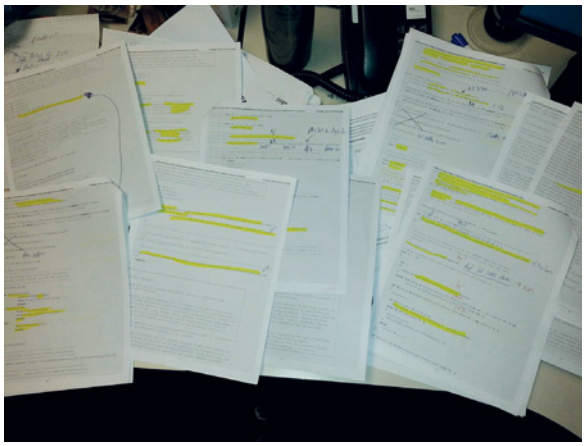

AvalonMM_packets.py

```
# Creates the packet for a read and write
def create_packet(trans_type, size, address, data = None):
    # Setup header from address
    header = [trans_type, 0, (size>>8)& 0xFF, size & 0xFF,(address>>24) & 0xFF
            (address>>16) & 0xFF, (address>>8) & 0xFF, address & 0xFF]

    temp_header = []

    for x in range(0, 8):
        if int(header[x]) == CONST_CHANNEL:
            temp_header.append(CONST_ESC)
            temp_header.append(int(CONST_CHANNEL) ^ 0x20)
        elif int(header[x]) == CONST_ESC:
            temp_header.append(CONST_ESC)
            temp_header.append(int(CONST_ESC) ^ 0x20)
        elif int(header[x]) == CONST_SOP:
            temp_header.append(CONST_ESC)
            temp_header.append(int(CONST_SOP) ^ 0x20)
        elif int(header[x]) == CONST_EOP:
            temp_header.append(CONST_ESC)
            temp_header.append(int(CONST_EOP) ^ 0x20)
        else:
            temp_header.append(header[x])

    header = temp_header
    if data is None:
        pass # Do nothing
    else:
        # Add the data to the header
        for x in range(0, size):
            # Insert ESC bytes to all cases and XOR data
            if(int(data[x]) == CONST_ESC):
                header.append(CONST_ESC)
                header.append(int(data[x]) ^ 0x20)
            elif(int(data[x]) == CONST_SOP):
                header.append(CONST_ESC)
                header.append(int(data[x]) ^ 0x20)
            elif(int(data[x]) == CONST_EOP):
                header.append(CONST_ESC)
                header.append(int(data[x]) ^ 0x20)
            elif(int(data[x]) == CONST_CHANNEL):
                header.append(CONST_ESC)
                header.append(int(data[x]) ^ 0x20)
            else:
                header.append(int(data[x]) & 0xFF)

    # Build the packet
    packet = [CONST_SOP, CONST_CHANNEL, 0 ]

    # Add the header
    packet += header

    # Insert the EOP byte
    if packet[len(packet)-2] == CONST_ESC:
        packet.insert(len(packet)-2, CONST_EOP)
    else:
        packet.insert(len(packet)-1, CONST_EOP)
```

Figure 2: Highlighter + example code + reverse engineer = Python script

LTC6948 Fractional-N PLL, set the LTC1668 DAC voltage, read a voltage from the LTC2498, or read temperatures from the LTC2983, just like the processor.

Altera provides a reference design for the SPI-Avalon MM bridge. Unfortunately, the documentation is sparse at best, and uses a NIOS processor as the SPI master. This effectively defeats the purpose of the SPI bridge, as the NIOS processor can interface directly to the Avalon MM bus.

A practical SPI master is Linear Technology's Linduino microcontroller, which is an Arduino clone with extra features to interface to LT demo boards. One extra feature is a level-shifted SPI port. This level-shifting function is especially helpful when interfacing to FPGA I/O banks with voltages as low as 1.2V. The Linduino firmware can be used to accept commands through a virtual COM port and translate the commands to SPI transactions.

After reverse-engineering the Altera example design (Figure 2), a Python library was developed to create packets the bridge would accept, which are translated into Linduino commands. A simple Python script allows the hardware engineer to have complete control of the project without having to reinvent the interfacing protocols.

> A simple Python script allows the hardware engineer to have complete control of the project without having to reinvent the interfacing protocols

An example Python script to control the frequency of a digital pattern generator for an LTC1668 DAC is provided in the LinearLabTools Python folder available at http://www.linear.com/solutions/linearlabtools. Figure 3 shows the demo setup.

Figure 4 is a block diagram of the FPGA system. Note that the numerically-controlled oscillator (NCO) can be managed by the shift register or the PIO core; the shift register is included for debug, as it allows direct control of the NCO. Setting the GPIO line logic level high enables the SPI-Avalon MM bridge, which in turn controls a 32-bit PIO port over the Avalon MM bus; the PIO output then controls the NCO frequency.

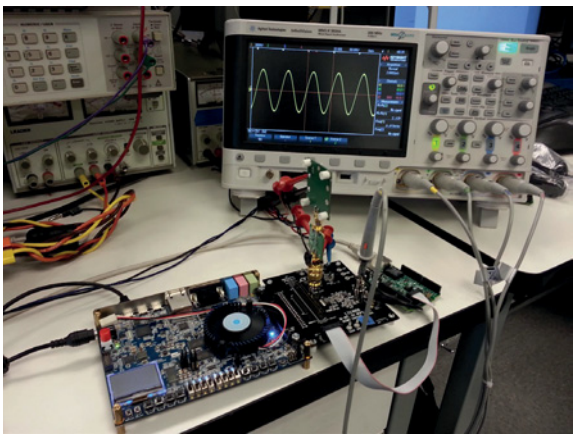

Figure 3: DC2459 in action

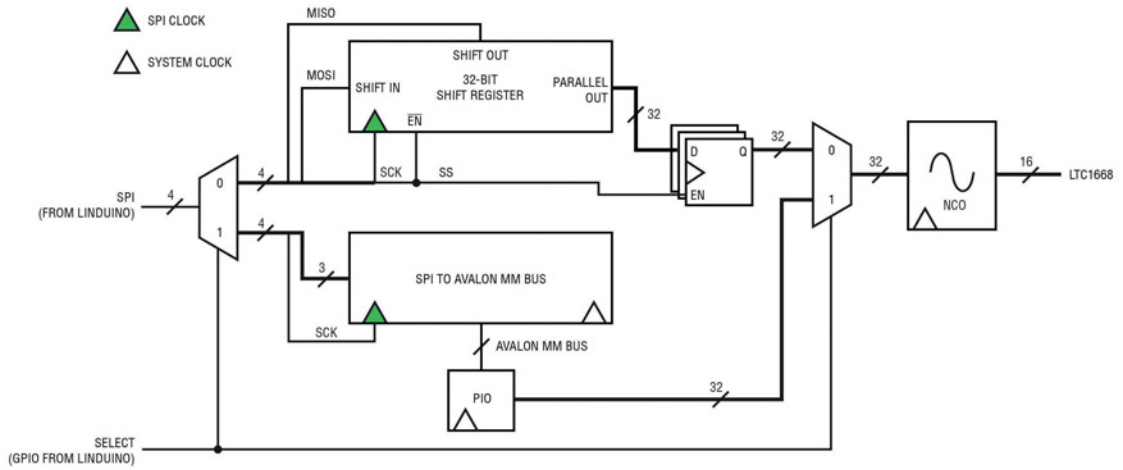

Figure 4: DC2459A FPGA system block diagram

Additional Avalon Peripherals

With the most basic system operation, additional Avalon peripheral IP cores can be connected to the Avalon MM bus. To design the system, Altera provides an integration tool called Qsys, which provides a GUI to connect the IPs to each other. Qsys is then used to translate the GUI system to HDL Verilog; see Figure 5.

Lastly, the system will then need to be added to the top level for implementation. The addresses of the IP are fully configurable. In the case of the example shown, the PIO is set to a base address location of 0x0.

Once the design is implemented in the FPGA, the Python library in the LinearLabTools contains two functions to interface to the design:

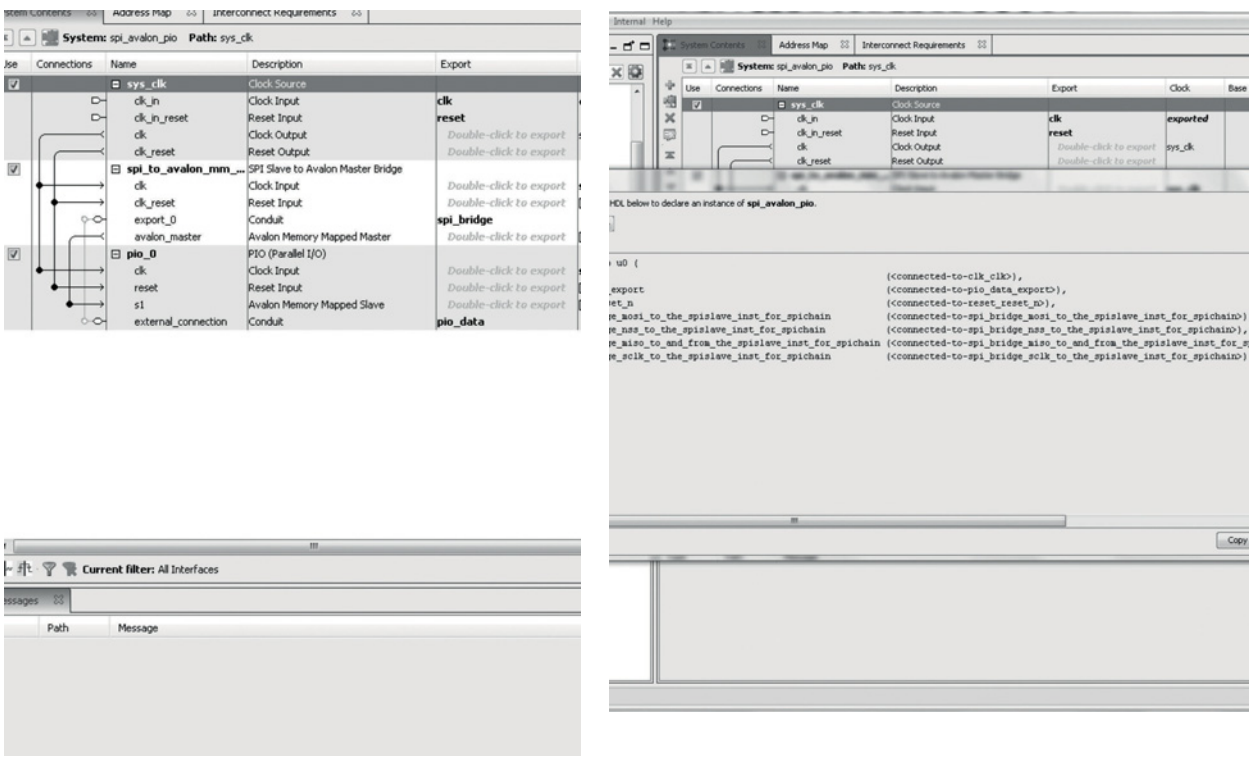

Figure 5: Qsys GUI

transaction_write(dc2026, base, write_size, data)
transaction_read(dc2026, base, read_size)

The first argument to these functions is the Linduino serial port instance; the second is the peripheral's address on the Avalon bus. The functions accept and return lists of bytes respectively. These two functions are written to allow flexibility when writing and reading to IP. To set the NCO for the provided example, the transaction_write function is all that's needed. Equation 1 is used to determine the tuning word:

$$tuning\ word = \frac{desired\ frequency}{system\ clock\ frequency} \times 2^{32} \quad (1)$$

To set the NCO to 1kHz with a 50MSPS sample rate, the tuning word value is 85899. This is 0x00014F8B in hexadecimal, which is passed as a list of four bytes. The Python code to set the DAC to 1kHz is then:

transaction_write(linduino_serial_instance, 0, 0, [0x0,0x01,0x4F, 0x8B])
Note: the base address of the PIO is zero from the logic design.

A simple Python script shown in Figure 6 is provided to demonstrate the interface of the FPGA design and Python script. It contains a simple text interface to configure the NCO.

An important note is that the Avalon SPI bridge uses SPI mode 3, painfully determined to be the correct mode by trial and error, and verified by analysing the NIOS processors SPI interface in the Altera's example.

Control Ability

To conclude, this example project demonstrates the ability to control a system without touching the embedded processor at all, allowing the hardware engineer to progress with a project without involving the software team. The benefits of this method can be unobtrusively added to the FPGA without affecting the original design. The hardware engineer can then focus on where needed the most – the hardware. ●

> This example project demonstrates the ability to control a system without touching the embedded processor at all

Code and Board

Python code is available at:
http://www.linear.com/solutions/linearlabtools

HDL for the LTC1668 demo board is included in the DC2459 design files, available at http://www.linear.com/demo/DC2459

```
    The views and conclusions contained in the software and documentation are
    those of the authors and should not be interpreted as representing official
    policies, either expressed or implied, of Linear Technology Corp.

    Description:
        The purpose of this module is to use the DC2026C as a Avalon MM Bus
        interface to set the DC2459A frequency out.
"""

########################################################################
# Libraries
########################################################################

import sys
sys.path.append('../../utils')
import connect_to_linduino as duino
import ltc_spi_avalon as avalon


if __name__ == "__main__":
    linduino = duino.Linduino() # Find the Linduino
    linduino.port.write("M3")   # Set to SPI Mode 3
    linduino.port.write('G')    # Set the GPIO HIGH
    try:
        print "Command Summary"
        print "    1-Send raw code"
        print "    2-Set frequency"
        print "    3-Exit program"
        user_input = input("Enter a command: ")
        while(user_input != 3):
            if(user_input == 1):
               code = int(input( "Enter raw 32 bit code: "))

               # Send the data to the avalon MM bus addr 0
               avalon.transaction_write(linduino, 0, 4,
                                [code & 0xFF ,(code>>8) & 0xFF ,
                                (code>>16) & 0xFF, (code>>24) & 0xFF])
            elif (user_input == 2):
               freq = float(input("Enter desired frequency(Hz): "))
               float_code = freq/50000000*(2**32-1)
               code = int(float_code)

               # Send the data to the Avalon MM bus addr 0
               avalon.transaction_write(linduino, 0, 4,
                                [code & 0xFF ,(code>>8) & 0xFF ,
                                (code>>16) & 0xFF, (code>>24) & 0xFF])
            else:
                print "**** Invalid Command ****"
            print "Command Summary"
            print "    1-Send raw code"
            print "    2-Set frequency"
            print "    3-Exit program"
            user_input = input("Enter a command: ")
    finally:
        linduino.close() # Close the port
```

```
Looking for COM ports ...
Available ports: [(23, 'COM24')]

Looking for Linduino ...
    Found Linduino!!!!
Command Summary
    1-Send raw code
    2-Set frequency
    3-Exit program
Enter a command:
```

Figure 6: Python Avalon bus example

PLC ARCHITECTURE IN THE INDUSTRY 4.0 SETTING: CHALLENGES AND TRENDS

BY **NILADRI ROY**, SENIOR STRATEGIC MARKETING MANAGER, INDUSTRIAL APPLICATIONS, INTEL PSG

Since their introduction, programmable logic controllers (PLCs) have evolved from simple input-output controllers to complete processor-based systems that execute complex control algorithms. PLCs have undergone significant form-factor changes: from industrial PCs and programmable automation controllers (PAC) in PC-like formats to compact enclosures and mini PLCs.

The scope of PLC functionality has also evolved. In addition to discrete control functions, PLCs have integrated functionality such as human machine interfaces (HMIs), motion control, real-time industrial Ethernet and data communication gateways.

The demand for additional features, precision and connectivity on the factory floor has driven this increasing integration. It has been sustained by PLC component cost reductions and the availability of higher-performance processing engines. In general, the evolution of PLC functionality has followed the industrial automation demand curves of features, performance and lower power. Following that trajectory, the demands of Industry 4.0 and IoT to a large measure will drive future PLC architectures.

> In general, the evolution of PLC functionality has followed the industrial automation demand curves of features, performance and lower power

Historical Role Of Motion PLCs

To appreciate the effect of market drivers on PLC architecture, it is useful to examine how motion PLCs drove performance requirements and, therefore, architecture. Over the last 10 to 15 years, a quiet revolution took place in factory automation: the number of controlled axes per machine increased significantly; see Figure 1 for the changes in CPU load.

With centralised control, where one CPU is in charge, all feedback loops, all axes (and other discrete I/O) share the processing power. This architecture limits performance and flexibility. An increased number of axes required faster update cycles and led to the deployment of faster, more expensive, power-hungry processors. But as energy costs became a significant concern and factory space became increasingly expensive, manufacturers moved away from large footprint, multi-processor architectures and back to centralised control. While this move solved the space constraints and associated costs, newer technological approaches were required to address the performance bottlenecks of single processors that had led to a distributed architecture in the first place.

Manufacturers attempted to address the performance issues with a novel approach: combining processors and FPGAs. The processor performed standard control functions (comprising primarily of gateway and discrete I/O control, but also HMI functionality), while the FPGA performed the rest, specifically motion control; see some examples in Table 1.

This architecture solved the performance problem but still suffered from increased power consumption of the faster processors. Additionally, it came at a much higher cost because of the high-performance processors and FPGAs.

The typical 5-year lifetime of processors also introduced additional product lifecycle headaches. Typical industrial equipment life cycles (seven to ten years) are incompatible with five-year processor lifetimes. This incompatibility led to obsolescence that contributed to delays in new algorithm development. Engineering teams scrambled to redesign existing equipment when a processor was at its end of life.

The hybrid approach also led to limited protection of manufacturers' software investments. There was no guarantee that a replacement processor would come from the same manufacturer. Therefore, roadmap scalability became a significant problem. Approximately 80% of manufacturers

Product	Architecture (Processor and FPGA)
PLC 1	NEC ASSP (MIPS) and FPGA
PLC 2	Renesas (V850) and FPGA
PLC 3	ST Spear and FPGA
PLC 4	Intel x86 and FPGA
PLC 5	TI Sitara and FPGA

Table 1: Hybrid processor and FPGA example architectures in real-world PLCs

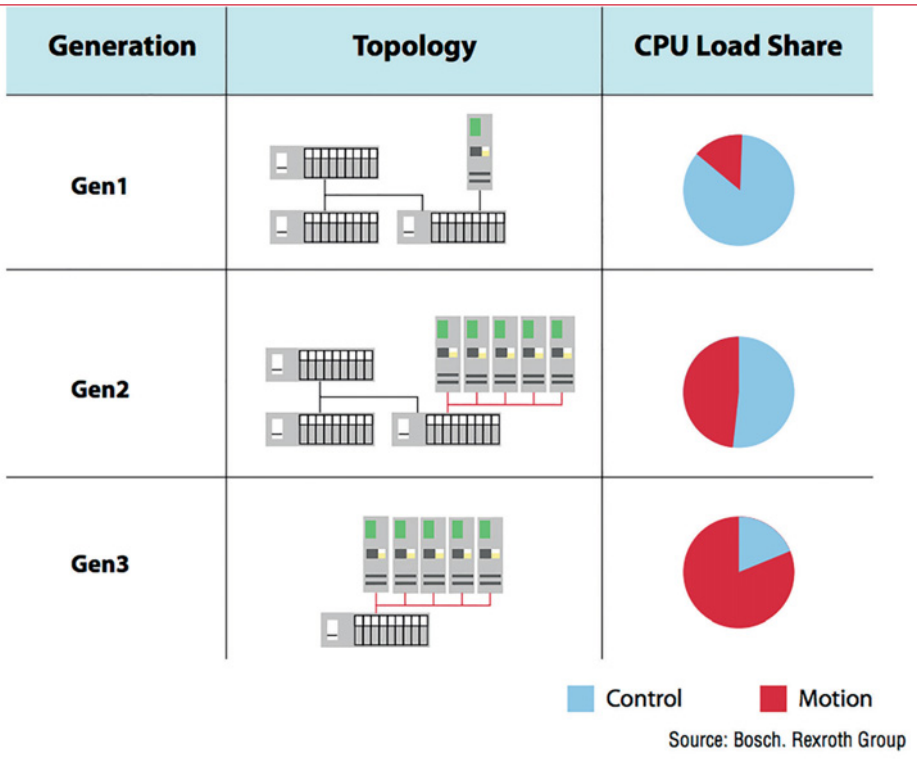

Figure 1. Control and motion CPU load share cloud-based manufacturing operations

encountered developmental delays and only about 30% of projects achieved their original forecast volumes due to delayed time to market.

Major PLC manufacturers experience issues such as longer design cycles, even more critical time to market and equipment needing to provide more deterministic I/O information. Other factors include increased pressure on cost, the demand for scalable high-performance applications and software development continuing to drive product differentiation.

Industry 4.0 Impact On PLC Architecture

The current manufacturing automation environment of Industry 4.0 demands high-performance PLCs enabled with secure enterprise connectivity and HMI. Today, multiple international Industry 4.0 initiatives rely on cyber-physical systems to implement the promise of smart manufacturing, leveraging connected systems for machine-to-machine (M2M) and enterprise interaction.

Making PLCs ready for Industry 4.0 is fraught with new challenges, requiring ground-up PLC re-design, including:

- High-performance control – Smart manufacturing environments require PLCs to process instructions, service interrupts and support integrated HMI at speeds faster than ever before. This has led to the use of more powerful processors with higher MIPS and multiple cores, resulting in higher cost and power consumption penalties.
- Connectivity – Deterministic M2M connectivity between disparate machines requires support for multiple industrial Ethernet protocols – including newly emerging standards-based deterministic Ethernet such as IEEE 802.1 TSN – within a single PLC system. Enterprise connectivity demands application interoperability frameworks such as OPC-UA.
- Secure communications – PLCs connected outside the factory network and to the enterprise are vulnerable to cyber attacks, making security a significant concern.
- Cross-platform interoperability – Choosing the wrong processor or ASSP can be an expensive error. Functional interoperability between diverse systems requires standardised operating systems running on non-proprietary processor cores.
- Future proofing – With an ever-evolving connectivity and interoperability environments, changes in market requirements are more frequent, leading to software and hardware updates.

Furthermore, pre-Industry 4.0 challenges remain, including scalability, functional safety, lower power consumption, smaller footprint and software investment protection.

A Smarter Approach To PLC Architecture

System-on-Chip (SoC) FPGAs, which combine a processor and FPGA fabric on a single chip, present a unique alternative for overcoming today's PLC design challenges:

- High-performance control – SoC FPGAs can offload the PLC's processor by implementing high-performance algorithms and

HMI in the hardware fabric. Unlike sequential processors, FPGAs are massively parallel and can greatly accelerate algorithm execution. With embedded digital signal processing (DSP) blocks and on-chip memory, FPGAs offer faster hardware acceleration at lower cost and power consumption than conventional processors.

- Connectivity – FPGAs can implement multiple industrial Ethernet protocols simultaneously on a single device with ready-made intellectual property (IP) cores. Designers can enable protocols by downloading the relevant protocol stacks, which execute in the built-in SoC FPGA hard-processor system (HPS). The HPS can also run an OPC server, enabling enterprise communications over OPC-UA. Designers can integrate emerging standards, such as IEEE 802.1 TSN, by reprogramming the FPGA hardware. The high-performance FPGA fabric can easily meet the stringent IEEE 802.1 TSN timing requirements.
- Secure communications – Open SSL encryption implemented in the FPGA fabric provides up to 4X acceleration over processor-based implementations. This encryption enables faster, more-secure enterprise communication channels.
- Cross-platform interoperability – With an integrated processor, SoC FPGAs offer a scalable roadmap using an industry-standard processor.
- Future-proofing – Designers can reprogram the FPGA fabric to incorporate hardware changes, avoiding major re-design of entire systems.

Manufacturers who use SoCs in PLC architectures derive the following benefits:

- High performance;
- 4,600DMIPs for less than 1.8W;
- Up to 1,600GMACs and 300GFLOPS based on >125Gbps processor-to-FPGA interconnect and cache coherent hardware accelerators;
- Lower power – up to 30% less power vs a two-chip discrete solution;
- Reduced BOM and PCB space and layer cost – up to a 55% form-factor reduction;
- Scalability and investment protection – scalable SoC processor roadmap grows with application needs and protects software development investment;
- Flexibility – SoC FPGAs can accommodate software and hardware changes;
- Improved time to market.

In competitive bid situations, manufacturers who have moved to a flexible SoC-based architecture can outbid those who have not. Additionally, a SoC architecture provides advantages in absolute price, total cost of ownership, scalability and protection of customer investment. Cost- and power-efficient SoCs can be rapidly adapted and scaled up according to growing and evolving customer needs driven by Industry 4.0 and Industrial IoT. ●

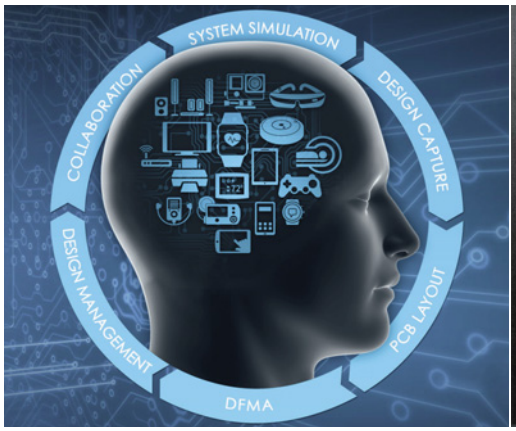

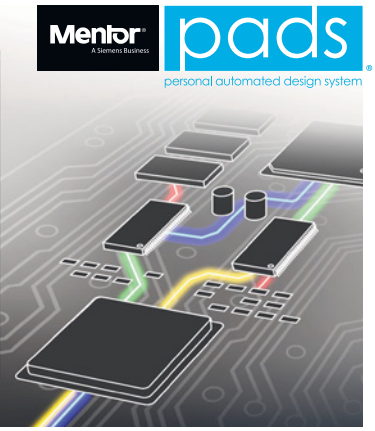

PADS® Strengthens DFT Capability with XJTAG® Boundary Scan Know-How

"Testing with Boundary can help boost test coverage, accelerate design verification and debugging, and increase production-test efficiency for Mentor, a Siemens Business. Mentor PADS users can now leverage XJTAG's experience to maximise the power of boundary scan in their designs without leaving their favourite environment, using the new, free XJTAG DFT Assistant for PADS."

Mentor PADS personal automated design solutions streamline product creation and help designers optimise all aspects of performance and manage their projects from design entry, through simulation and analysis, to sign-off for production. Optional extensions allow users to add capabilities such as advanced board layout, power-delivery analysis, thermal analysis, and support for RF design, high-speed design, and high-density or timing-critical routing.

PADS is now even more powerful with XJTAG's boundary scan test know-how built-in. "Boundary scan can add value from the beginning of the product lifecycle, and is becoming increasingly important to our customers," explains Jim Martens, Product Marketing Manager, PADS Solutions Group. "We saw the opportunity to enhance PADS with class-leading design for boundary scan test capability, by integrating the features of XJTAG's highly regarded DFT Assistant."

Boundary scan can check a high proportion of a board's connections early in the design phase, before any hardware is produced, and only requires the Test Access Port (TAP) pins of JTAG-compliant components to be correctly linked and routed to a connector. The simple four-signal interface allows easy software-based access to I/O pins that are otherwise hard to reach with probes, such as BGA I/O connections. The TAP, and traces comprising the scan chain that links the JTAG pins, occupy minimal real estate on the board.

When designing and prototyping boards, boundary scan tools help check for design errors before any hardware is built. First prototypes can be tested quickly to pinpoint connection errors, potentially saving hours buzzing out boards looking for shorts or opens that may cause errors or prevent the board starting up. In production, boundary scan can quickly check a high percentage of connections to help isolate defective boards and boost overall test efficiency.

Engineers can maximise the test coverage achievable with boundary scan by connecting JTAG compatible components into a JTAG chain. Using the JTAG chain, testing can be further extended to non JTAG compatible devices. PADS users can take advantage of XJTAG's Design-for-Test (DFT) know-how, acquired through years working with clients and refining the XJTAG test development suite, by using the XJTAG DFT Assistant for PADS now included in their favourite design environment.

XJTAG DFT Assistant for PADS features an Access Viewer that gives a graphical view of JTAG chain access across the board, which help users visualise the extent of test coverage and see how their design changes affect testability as the project progresses. In addition, the Chain Checker verifies that all the JTAG and TAP pins are correctly connected and terminated before committing to hardware. The information can be exported directly to the XJTAG test-development environment, where the testing to be carried out can be configured.

"Our customers can now use PADS to produce even better board designs that benefit from higher test coverage, faster debugging and prototyping, and more efficient testing in production. Working with XJTAG enabled us to achieve a high-quality result within a fast turnaround time," concludes Jim Martens.

opinion

Jim Martens
Product Marketing Manager
PADS Solutions Group

"Our customers can now use PADS to produce even better board designs that benefit from higher test coverage, faster debugging and prototyping, and more efficient testing in production. Working with XJTAG enabled us to achieve a high-quality result within a fast turnaround time."

"Boundary scan can add value from the beginning of the product lifecycle, and is becoming increasingly important to our customers. We saw the opportunity to enhance PADS with class leading design for boundary scan test, by integrating the features of XJTAG's highly regarded DFT Assistant."

Data Bank

Mentor® A Siemens Business

Company	Mentor Graphics Corporation A Siemens Business
Nature of business	World leader in electronic hardware and software design solutions providing products, consulting services, and award-winning support for the world's most successful electronic, semiconductor, and systems companies
Location	Corporate headquarters are located at 8005 S.W. Boeckman Road, Wilsonville, Oregon 97070-7777
Web site	www.mentor.com

ACCELERATING WIRELESS CERTIFICATION

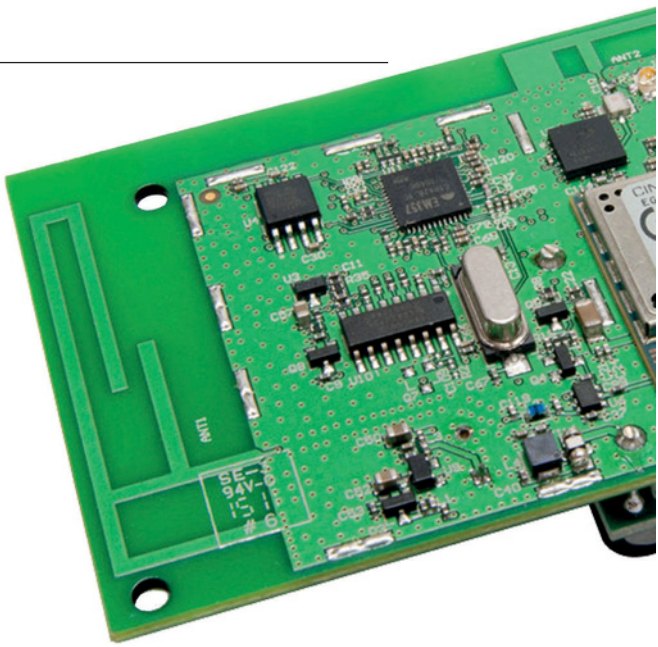

BY **KYLE SPORRE**, RF/HARDWARE ENGINEERING MANAGER AT DIGI INTERNATIONAL WIRELESS DESIGN SERVICES

An important concern in wireless design is how to pass certifications quickly. The complexities involved in designing printed circuit boards (PCBs) and antennas into wireless solutions are challenging enough, but the certification process comes with additional layers of challenges that developers often come to understand too late. That can be costly in multiple ways, from additional board spins to slower time to market.

The following wireless design considerations and best practices can drastically reduce development time and certification headaches:

1. **Be selective in choosing the geographic regions to support:** It can be tempting to choose to design a new product for worldwide release, but it is often best to add geographic regions in phases. Radio frequency (RF) regulations and testing requirements vary by region, and will heavily affect several aspects of the product development, including component selection, product mechanical size, certification cost and end-product unit cost.
2. **Plan for certification:** Before launching the design process, get to know the standards, tests and pass/fail criteria involved in the certification process. While this seems intuitive, even seasoned PCB designers can make the mistake of starting a design before adequately planning for certification. Certification requirements, which we'll cover in this article, can have a tremendous impact on your product design.
3. **Partner with the right experts early in the planning process:** This step can make all the difference in helping understand the regulations and certification requirements and get the design off to a good start. These partnerships can include test labs, cellular carriers – if appropriate – and professional design services that can help create the optimal design to pass manufacturing and certification tests with ease.
4. **Design for first-pass success:** The most obvious way to accelerate certifications is to pass on first attempt. This requires designing the product for first-pass success, instead of waiting until certification testing to fix issues. There are several key design strategies, which we'll discuss later.
5. **Pre-test early and often:** Testing aspects of your design along the way can help ensure you are on track. One should never assume that best practices alone will mitigate all risk, as even a high-quality digital design can fail certain certification tests.

80% of all new cellular designs fail certification

Choosing The Regions

Where do you plan to market your product? You will need to fully understand the regulations in each intended region, to design a successful product that can pass all certifications quickly. As an example, to roll out a product in more than one region around the world with a single SKU, the product will likely need an antenna that supports multiple frequency bands, which is sometimes achieved by a wider frequency bandwidth. This wider bandwidth typically requires a physically larger antenna.

The PCB ground plane is often an important radiation component of the antenna system and will also likely need to be larger. If you are developing a small product, such as a wearable, meeting these multiple requirements can be very difficult with one SKU.

Depending on the regions where the product will be introduced, you will need to work with one or more regulatory bodies. The certification process can vary widely from country to country, and doing some early research to understand the regulatory requirements of each region to launch your product will give you a big leg up in planning the design.

Planning For Certification

While we can't cover all wireless design certification criteria worldwide, FCC certification is an excellent example to examine. All commercial electronic products sold in the US with any oscillation of 9kHz or higher are subject to FCC regulation. Understanding these tests up front will support the highest quality, certification-ready design.

Typical tests required for FCC certification include:

- **FCC Part 15, subpart B – unintentional radiator:** This is a requirement for commercial electronic devices with oscillations greater than 9kHz that don't deliberately generate RF emissions. Unfortunately, the unintentional radiator test is getting much more difficult to pass, because devices today tend to have faster clock speeds and consequent shorter wavelengths. The result is that smaller structures on PCBs essentially act as antennas for the shorter wavelength clock noise, which is then radiated as undesired emissions.

  You can mitigate the issues through effective PCB design using techniques such as shielding and effective grounding. If you plan your design upfront to meet the criteria, you may save board spins, schedule delays and extra certification costs.

  Subpart B is self-declared and unintentional radiators do not necessarily require an FCC ID.
- **Part 15C, 22, 24, 27 – intentional radiator:** These requirements are for devices that intentionally transmit RF signals and evaluate output power, harmonics, out-of-band emissions and other signal characteristics while the radio is transmitting.

  To facilitate testing, it is important to build test modes into your radio. Harmonics are the most common cause of test failure. There are several issues that can cause harmonic failures, such as non-linear power amplifiers in the transmitter chain that generate harmonics which are then radiated by the antenna. Also, non-linear PCB components can pick up the fundamental frequency radiated from the antenna and then generate harmonics of that frequency, radiated by the PCB. In the design phase, you can mitigate these factors with measures such as adding a filter between the radio and the antenna and keeping radiated RF out of the power system.
- **FCC modular approvals:** This is a set of criteria that affect end products typically tested for intentional radiator, but that may potentially bypass that step because they incorporate a pre-certified module. These designs may be exempt from intentional radiator testing if the end product meets all the conditions of the module's FCC grant; the conditions include:
  - The product can contain no other radios within 20cm of each other that can transmit simultaneously;
  - It may not be used within 20cm of the human body;
  - It must use an antenna of equal or lower gain than the certified module.
- **Specific Absorption Rate (SAR):** SAR is a measure of the rate at which energy is absorbed by the human body when exposed to electromagnetic fields. The purpose is to ensure that the device will not cause health hazards due to tissue heating. If the end product is to be used within 20cm of the human body, you must perform a Maximum Permissible Exposure (MPE) calculation to determine whether SAR testing is required.

  In addition to the FCC, cellular products can have a variety of certification requirements that depend on variables such as whether the product uses a certified module, whether there's a specified antenna within 20cm of the device, and the chosen cellular carrier.

  Below is a very common certification example of an LTE data-only product with an internal antenna and integrated certified cellular module, intended for release in North America.

PTCRB

With a few exceptions, most North American cellular carriers require PTCRB certification, in addition to some carrier-specific network testing. PTCRB tests for this example include:

- **SIM card interface testing.**
- **Radiated Spurious Emissions (RSE):** This is an emissions test similar to the FCC's, with more stringent limits. It evaluates emissions from the cellular transmitter when the radio is transmitting (active mode), as well as emissions from the rest of the electronics when the cellular radio is not transmitting (idle mode).

- **CTIA Over-The-Air (OTA) test plan:** The OTA test plan is only required for cellular designs with antennas less than 20cm from the radio. Note that box products with an antenna port where the customer chooses an antenna are not subject to the challenging OTA test plan.

  The OTA tests include:
- **Total Radiated Power (TRP):** measures the total power radiated from the device in a three-dimensional sphere, an indicator of antenna efficiency and impedance match, transmitter output power and PCB RF design. There are no pass/fail limits for TRP for PTCRB certification, but various cellular carriers review the TRP data and apply their own pass/fail limits for carrier certification.
- **Total Isotropic Sensitivity (TIS):** measures the radiated receiver sensitivity of the system when averaged over the entire three-dimensional sphere. It is a function of radio receiver sensitivity, PCB RF design, antenna impedance match and efficiency, but is usually limited by noise radiated from the device electronics. Like TRP, there are no PTCRB pass/fail limits, but various carriers have their own pass/fail limits for TIS.
- **Relative Sensitivity Intermediate Channels (RSIC):** This is a radiated test with PTCRB pass/fail limits that evaluates the relative receiver sensitivity across a subset of supported cellular frequency channels. It ensures against degraded performance on specific frequency channels due to narrow-band noise interference, usually from the host electronics.

> The most obvious way to accelerate certifications is to pass on first attempt

Partnering With Experts

This step can help evaluate all design considerations that impact certification early in the process, when it's still easy to make adjustments. Meet with your test labs to go over the test criteria. If it's a cellular design, work with the cellular carriers to ensure you are well apprised of their requirements.

If you are considering launching FCC and PTCRB testing at once to save time, be aware that FCC certification must be obtained before receiving PTCRB certification. Therefore, FCC failures that lead to design changes could cause wasted PTCRB testing.

Using professional design services can help fast-track your design process and reduce failures, accelerating certification. One test lab manager has estimated that a startling 80% of new cellular designs requiring OTA testing fail certification the first time, resulting in schedule delays and costly additional design cycles. Consulting with knowledgeable experts and/or outsourcing aspects of your design upfront can thwart such issues.

Design For First-Pass Success

Passing certifications the first time is an achievable goal – if right measures are taken at the right time. Antenna design and implementation are vitally important, as is low-EMI PCB design, for the reasons discussed earlier.

Effectively controlling noise, especially for radiated cellular tests such as TIS and RSIC, is the most often overlooked aspect of designing for certification. Don't lock down an industrial/mechanical design before antenna placement and electrical design, as this can make satisfactory performance virtually impossible, driving mechanical changes late in the development cycle.

Finally, at phases throughout product development, ask your test lab to do pre-scans and spot check high-risk areas so you can alter your design as needed. These changes can be very simple or complex, up to and including changing your entire PCB layout and antenna implementation. But, the earlier in the process you can identify issues, the lower the cost and risk you'll ultimately face in the certification process. ●

Electronics WORLD

Your essential electronics engineering magazine and technical how-to-guide

A subscription to Electronics World offers

- 12 monthly issues in digital format
- Regular topical supplements
- The Annual T&M Supplement
- Weekly email bulletin
- Comment and analysis from industry professionals
- Tips and tricks
- News and developments
- Product reviews
- Free invitations to our webinars

Electronics World provides technical features on the most important industry areas, including:

- Nano measurement
- Robotics, and much more...
- DSPs
- Power supplies
- Automotive
- Lighting USB design
- Power supplies
- Test and measurement
- Cables
- Signal processing
- Connectors
- Semi-conductors
- Digi
- Embedded
- Microwave
- RF
- Communica-tions

SUBSCRIBE TODAY FROM JUST £46 BY VISITING THE WEBSITE OR CALLING +44(0)1635 879 361

www.electronicsworld.co.uk/subscribe

Register for our free newsletter, please scan here

OPTIMISED CONDUCTIVE ELASTOMER CONNECTOR STRUCTURE AND APPLICATIONS

BY **JIACHUN ZHOU, KILEY BEARD** AND **KEVIN DEFORD** AT SMITHS INTERCONNECT

Conductive elastomer or rubber is a manufactured elastic material that conducts electricity. Under pressure, this material provides reliable contact between components' uneven surfaces. It's been extensively applied in various industries, including pressure sensors and conductive mats to prevent electrostatic damage to devices. It is also an important electromagnetic interference (EMI) shielding material, providing electrical conductivity across gasket-flange interfaces.

Another valuable application is in testing integrated circuit (IC) chips and packages, where these materials have been used as contactors for over 20 years; they are preferred over other types, because they offer a shorter electrical path.

A compressed conductive elastomer typically has a thickness under 1.5mm, whereas other contactor forms exceed 2mm. This reduced electrical path enables better signal integrity (SI) performance, and the structure provides a reliable contact with low contact resistance and consequent high conductivity.

Elastomer Connector Basic Structures

For EMI shielding or dielectric materials, these semi-conductive elastomers are typically manufactured by distributing conductive particles throughout the elastomer sheet. The particles are typically Ni, Ag or Au, or Au-plated Cu wire. Figure 1 shows a conductive elastomer gasket.

The structure contains conductive columns (Figure 2), which are slightly raised above the elastomer sheet surface for even more reliable contact.

Figure 3 shows a Smiths Interconnect conductive-elastomer connector with another specific structure, called "Contact Sheet", which contains solid Au-plated Cu for greater contact reliability and to avoid contamination.

Elastomer Connector Applications

In the chip test industry (Figure 4), a direct contact is established between package and test board. Here, conductive elastomer contactors perform better than most other contactor technologies, such as spring probe contactors; however, their weakness appears in applications that require many thousands of cycles of compressive contacts.

The direct contact of the device pads on the conductive columns in elastomers may generate contamination on its surface from the migration of solder or metal particles. These contaminations degrade conductive elastomer performance by lowering its conductivity and shortening its service life.

Efforts have been made to make improvements in this industry. Using a separate contact sheet on the elastomer is one solution; see Figure 5. The contact sheet is placed above the elastomer sheet as an interface with the device. Adding another interface above the conductive elastomer will affect performance. To minimise such impact, the following considerations have been made in contact-sheet design:

1. Reduced thickness of the contact sheet, keeping its profile as low as possible, to avoid signal loss and noise.
2. Reduced force while maintaining effective contact. Unlike spring probes, elastomer compliance is low and shows a

Figure 1: Conductive elastomer gasket

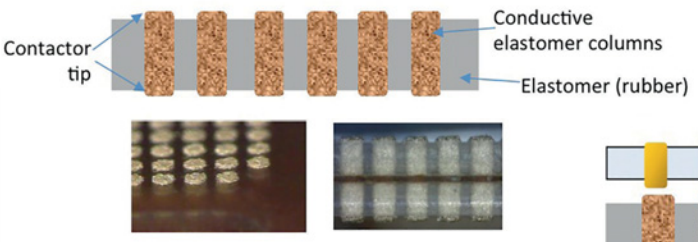

Figure 2: Conductive elastomer connector

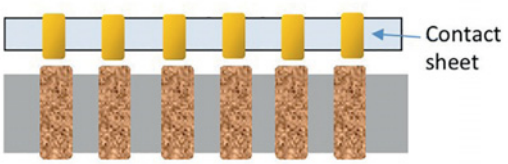

Figure 3: Connector with contact sheet

weakness in large-device applications. To achieve compliance here, high compressive force could be needed, which can prove much higher than the allowable force of the handler in package testing.

The contact sheet and elastomer should have the same compliance. If the contact sheet has limited allowable compliance or needs extra force, the elastomer contactor may not provide enough compliance for large package applications.

Signal Integrity Performance

Generally, the contactor material and mechanical structure determine the SI performance. So, here are some general guidelines:

1. Total length of the contactor (contact sheet head + elastomer thickness). As shown in Figure 6, the bandwidth decreases significantly when the contactor length is >3mm. Most conductive-elastomer thicknesses are ~1mm. After adding a contact sheet, the total contactor length is around 2-3mm. Using a minimum contact-sheet thickness is critical for contact sheet design, since thinner contactors minimise noise in signal transmissions.
2. Contact head diameter should be similar to the elastomer's conductive column diameter to avoid sharp variations in signal path. Generally, larger diameter contactors have less inductance while smaller ones have higher bandwidth at the same pitch structure.
3. Contact sheet head material usually uses Cu plus Au/Ni plating with limited elastic polymer material options, thus they don't have a significant impact on contactor SI.

With all these optimisations, signal integrity performance can be >40GHz at insertion loss IL <-1dB and return loss RL <-10dB, as shown in Figure 7.

Cleaning Methodology

Cleaning and protection are always a concern for conductive elastomer connectors. Direct contact of chip devices on the elastomer affects its life significantly due to surface contamination, as shown in Figure 8.

The contact sheet heads have crown structures and materials similar to other type connectors, including spring probes, to provide reliable contact with balls or flat pads. Figure 9 shows a contact head crown with a ball and Figure 10 contact sheet's bottom tails.

As a general guideline, the cleaning methodologies used on other metal connectors can be applied to conductive elastomer contactors. Some cleaning methods, for both contact sheet and elastomer, are outlined below. The proper cleaning frequency should be selected, based on actual applications and contamination on the contact sheet surfaces.

- **On-line cleaning (contact sheet only):** when the conductive elastomer with contact sheet is kept on the test board with socket.

1. Use Mipox or cleaning sheet with manual compression lid or auto-test handler to clean the contact sheet head crown.
2. Use a soft nylon brush and compressed air (if the environment allows) to clean the contact sheet head crown.

- **Off-line cleaning:** The contactors are removed from test boards and disassembled from sockets.

Contact Sheet:

1. General cleaning with a soft nylon brush and compressed air to blow off debris.
2. The contact sheet can be removed from the socket, allowing both sides to be cleaned with chemicals, such as 99% isopropyl alcohol. If cleaning chemicals are used, the sheet must be completely dry prior to re-installing with the elastomer. The same holds true if the PCB is cleaned with chemicals.

Elastomer:

1. A stiff or wire brush, alcohol or other cleaning chemicals should never be used. If these chemicals are used to clean the contact sheet or the PCB, completely dry surfaces are necessary prior to contact with the elastomer.
2. Compressed air should be used first to blow dust and debris from both sides.
3. If there is debris that can't be removed by air, a soft-hair brush or a light adhesive, such as a Post-it Note, can be used to gently remove any remaining debris. Contaminants can also be removed by carefully using tweezers under a microscope.

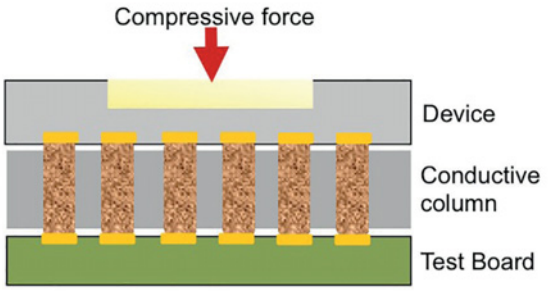

Figure 4: Conductive elastomer connection: device and test board

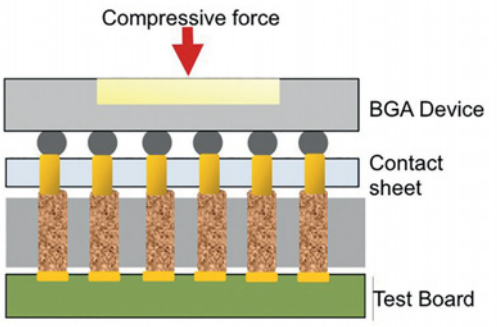

Figure 5: Contact sheet in elastomer connector

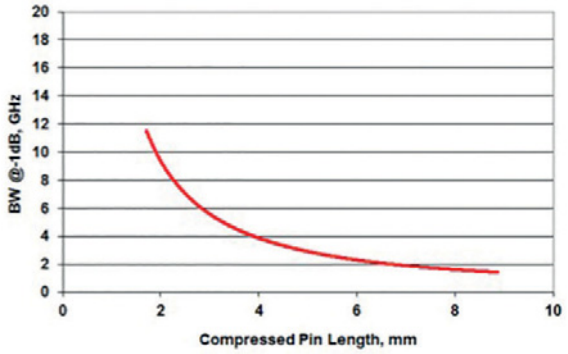

Figure 6: Bandwidth vs connector length

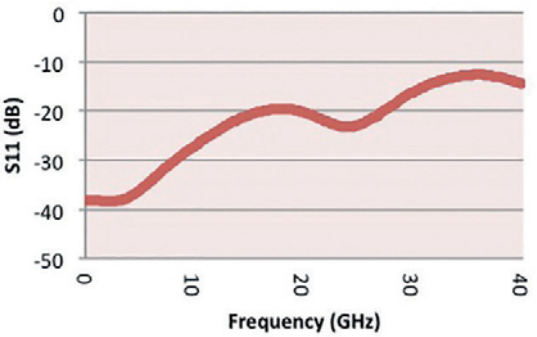

Figure 7: Return loss (dB) vs frequency (GHz)

Summary

For many years, conductive elastomer connectors applied in various industries have demonstrated their strength through excellent signal integrity performance and reliable contacts with proper conductivity. However, rapid performance degradation due to contamination on the elastomer surfaces can affect their wider market acceptance. Using a contact sheet with an elastomer is an effective approach to protecting the elastomer and increasing its service life. To minimise the impact on elastomer performance, the following key factors should be considered when optimising the contact sheet structure:

1. Minimise thickness of the contact sheet for better RF performance.
2. Apply all knowledge of contact geometries from other contactor technologies to reduce contamination and enhance self-cleaning.
3. Maintain enough compliance in the contact sheet to ensure contact while minimising the impact on the compressive behaviour of the elastomer sheet.
4. Avoid extra compression force on the elastomer sheet through proper choice of head structures.
5. Cleaning methods commonly used with other connectors can be applied to conductive elastomer connectors with contact sheets.

Smiths Interconnect has done extensive research to optimise the contact sheet performance with a patented contact sheet structure. The experimental and field application data show that the contact sheet's long service life of over two million touchdown cycles with cleaning every 100K cycles provides a significant advantage to its users. ●

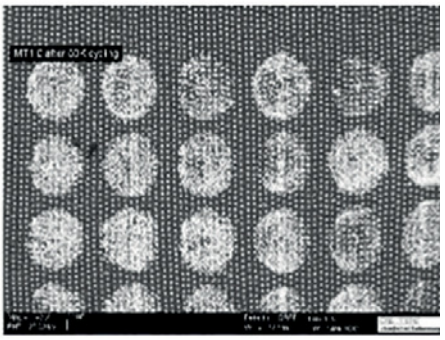

Figure 8: Contaminants on elastomer sheet surface

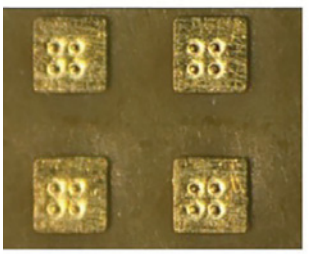

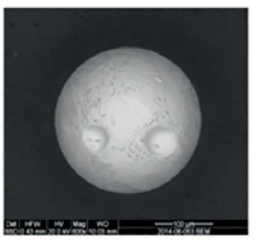

Figure 9: Contact head crown and contact marks on ball

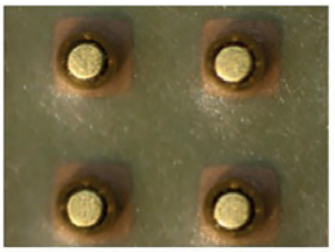

Figure 10: Contact sheet bottom tails (contacting conductive elastomer

Figure 1: Typical battery pack configurations

IMPROVING TAB-TO-TERMINAL CONNECTIONS IN BATTERY PACK MANUFACTURING

THE RIGHT WELDING SYSTEM WILL HELP ACHIEVE AND MAINTAIN PRODUCTION THROUGHPUT TARGETS.
BY **GEOFF SHANNON**, AMADA MIYACHI AMERICA

Battery packs have become an integral part of everyday life, powering a growing range of portable electronic devices, cordless power tool and hybrid and electric vehicles. As such, manufacturers need equipment, systems and automated production lines that meet quality and volume requirements for these products.

Both resistance and laser welding systems are well suited to integration into production lines that may be either standalone or automated. To maintain the required throughput with high quality and yields, users must have a clear understanding of which process is best for a particular battery pack size, tab and terminal material, and type and thickness. In addition, the selected process and integration solution should include process monitoring, process data management and weld quality assessment.

Battery-Pack Basics

Today's battery packs come in a variety of configurations (Figure 1), including cylindrical, prismatic, ultra-capacitor and pouch (Figure 2). Material joining requirements vary depending on the battery's specific type, size and capacity, and include tab-to-terminal connections, internal terminal connections, tab welding, seam welding, fill-port welding, short circuit protection, laser marking and external electrical connections. Figure 2 also shows the typical joining requirements for the different battery types. This article focuses primarily on welding tabs to terminals, since it is one of the key battery-pack manufacturing steps.

In most cases, battery-pack manufacturers receive individual batteries from vendors, so the critical process step for pack manufacturing is joining the individual batteries together using a collector plate, which consists of tabs for the

individual cells to be welded to both terminals. In addition, many packs contain a smaller number of collector-plate-to-busbar connections.

Along with considerations of materials, joint geometry, weld access, cycle time and budget, the welding technology selected will also be affected by manufacturing flow and production. Reviewing all these factors will usually point in the direction of the joining technology most suitable for the application: either laser or resistance welding.

Fibre Laser Welding

The laser welding process is non-contact, has no consumables and offers instantaneous welding once the laser is positioned at the weld point location. Weld size and location on the part can be closely controlled and optimised for each application for strength and conductivity. There are several motion options that can be tailored to each manufacturing environment.

For tab-to-terminal welds, fibre lasers can be used for prismatic, cylindrical and pouch battery types, as well as ultra-capacitors. Tab thickness can vary from 0.006-0.08 inch for both aluminium and copper tab material, depending on battery size.

The fibre laser can weld many material combinations including aluminium to aluminium, aluminium to steel, copper to steel, and copper to aluminium. Whatever the material combination and thickness, the laser must not penetrate the can or overheat the battery. Suitable selection of fibre laser, spot size, weld parameter selection and weld path control enables fine control of both penetration and heating.

Figure 3 shows some examples of common dissimilar material combinations for tab-to-terminal welding. For a lap weld, the tab thickness should ideally be 50% thicker than that of the can, providing a large processing window and high production yields.

The fibre laser source can be sized and matched with an appropriate motion platform according to cycle time and production throughput. For high-volume systems, the welding time can run at <100ms per weld, while low volume production systems run at about 1-2s per weld.

Resistance Welding

Resistance welding has been used in the battery industry for 40 years. A steady stream of advances has offered significantly improved capabilities to control various aspects of the process. For example, the introduction of DC inverter power supplies

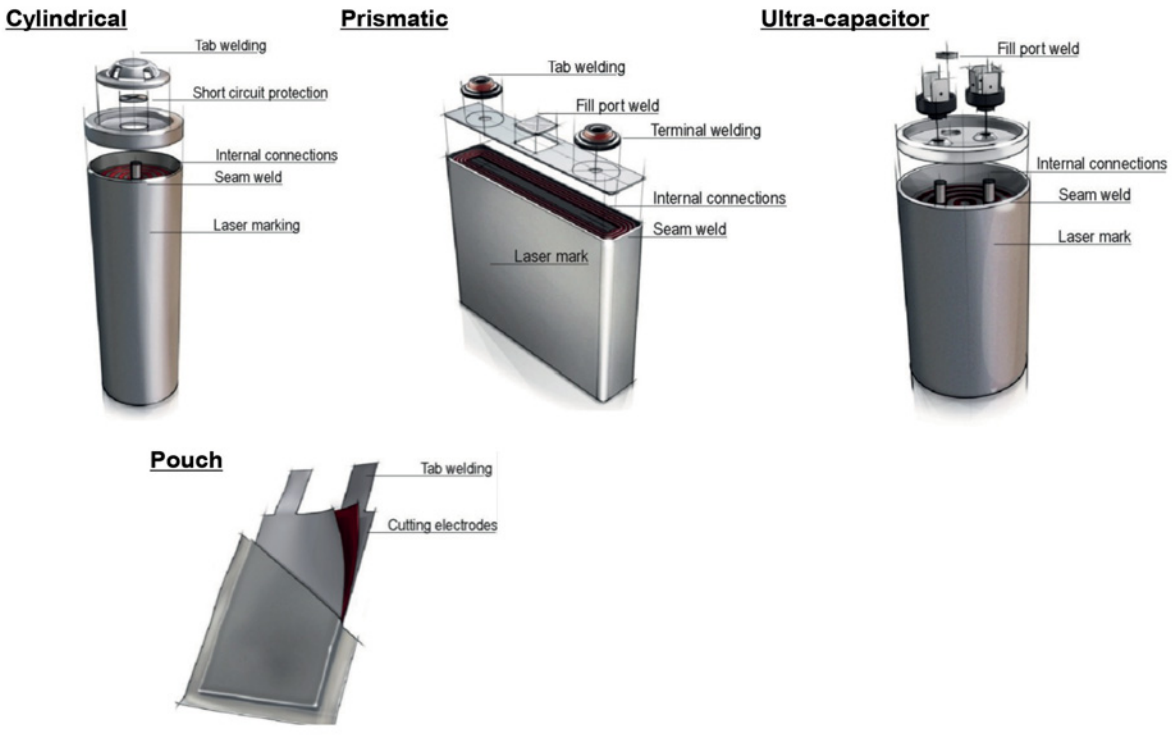

Figure 2: Battery types

Technology	Key Benefits	Battery Pack Applications
Laser	Non-contact High speed welding Tailored weld patterns Weld any joint geometry Weld dissimilar metals	Cylindrical, prismatic, pouch, ultra-capacitor
Resistance welding	Closed loop feedback welding Cost effective Self-tooling	Cylindrical, small prismatic

Table 1: The available methods

Technology	Details	Welds per second
Resistance	< 0.015" thick nickel/steel straps < 0.007" thick copper straps	Up to 1*
Laser	Up to 0.04"+ thick tab material	Up to 20*

**Material and material thickness dependent*

Table 2: Overview of battery-pack joining applications

with basic closed-loop electrical modes helps accommodate changes in the secondary circuit (the electrical loop from the cable connection on the negative side of the power supply or transformer, through the weld head and the parts, returning to the positive side) to specifically address part resistance. Also, polarity switching of capacitance discharge supplies to enable balancing of the weld nuggets and, more recently, the addition of displacement and electrode force measurement, provide manufacturers with more tools to ensure weld quality.

Resistance welding is the most cost-effective method for joining tabs on a wide range of battery types and sizes, using both DC inverter closed-loop and capacitor-discharge power supplies. With fast rise times, closed-loop feedback control, polarity switching and options for displacement and force sensing, the process can be fine-tuned and monitored to ensure both high quality and yield.

For nickel tab thicknesses up to 0.0070 inch, the tab can be welded without modification. Beyond this thickness, and to prevent electrical shunting and excessive electrode wear, a slot and projections are formed in the tab as part of the stamping process. The projections act not only as energy concentrators for the weld, but also greatly increase electrode lifetime.

The important aspects of tab welding for battery packs are the thickness and material of both tab and terminal. Resistance welding is extremely well suited to welding nickel tab material up to 0.015in thickness, and nickel- or steel-clad copper tab material to around 0.012in thickness, to a wide variety of terminal materials. Welding tabs or terminal connections to busbars generally does not require as much heat penetration as the tab-to-terminal welds. The materials, their thickness and their combination determine the best welding technique.

Technology To Suit Manufacture

Battery-pack production volumes are driven by the demands for consumer electronics and electric vehicles. Likewise, manufacturing and joining needs are determined by the pack size, type and thickness of the busbar, and tab and terminal materials. Laser and resistance technologies each have specific features that align well to these joining needs. A clear understanding of the technologies and application is needed to implement an efficient and reliable production battery-pack welding system.

Table 1 gives some guidelines on the available methods and a few parameters, including suitability for a variety of battery pack applications. Table 2 provides an overview of battery pack joining applications and key components of joining solutions required.

A complete production solution for battery pack manufacturing needs to deliver and support the required product flow. The system must take into consideration how the pack is loaded and unloaded, how the unwelded parts are

Figure 3: Weld cross-section of tab-to-terminal material combinations

held prior and during welding, and whether and to what level the system reports to the supervisory control software.

● *Enclosures*

The enclosure requirements for laser and resistance welding are quite different – the key difference is the need for lasers to be housed in a class 1 light-tight enclosure per ANSI 136.1. This requires the complete system to be enclosed, and adds the need for enclosure doors for part-load and unload in addition to access panels for setup and maintenance.

For a resistance welding system, although there may not be any panels, light curtains are used to protect the keep-out zones. If loading into the enclosure is automated, a conveyer is commonly used.

● *Motion*

To support in-line production that uses conveyers and minimise the system's physical footprint, resistance weld heads or laser focus heads should be moved while the pack remains stationary. For resistance welding systems, a defining factor in determining whether the weld head or the pack moves is the length of weld cable needed. Increasing the length requires additional voltage from the power supply to push the current through, which tends to a limit of between three and eight feet, depending on the application and power supply.

Scan heads are being used for more and more laser welding motion solutions. For low-volume applications, they offer cost-effective solutions. For high volumes, scan heads offer a very high-speed motion that can be used to minimise weld time and point-to-point positioning times. Scan heads can also provide a large weld area, with certain configurations providing greater than a 10x10in (250x250mm) weld area. For large pack sizes, this minimises the number of steps needed to weld the entire pack, which can significantly reduce cycle time. For example, point-to-point positioning of the scan heads for a 1in (25mm) travel distance can be completed in 10-20ms.

● *Tooling (for laser welding only)*

The ease of tooling a material is directly related to the tab or busbar thickness and subsequent stiffness. Thick material with high stiffness does not deflect under clamping pressure, so it maintains contact with fairly simple tooling. However, since its thickness for many battery packs is below 0.015in, this means the tab stiffness is not sufficient for simply clamping. Instead, it requires a well-designed tool to apply a localised clamp force that ensures intimate contact without damaging the parts.

Figure 4: Laser welding focus head with integrated tooling

With decreasing tab thickness, tooling becomes even more critical. Welding tab thinner than 0.005in (0.125mm) is not recommended.

Another critical factor for laser welding and part positioning is ensuring the part's weld surface typically in the z or vertical plane matches the laser's focus. This sometimes requires the use of height sensors or vision systems. Figure 4 shows an example of a laser welding focus head with integrated tooling.

● *Communication*

The protocol and control of data to and from such systems depend on the level of automation and the manufacturing environment. Therefore, options and a flexible platform for communication control software and data management are needed to ensure scalability and performance in diverse applications. For example, Industry 4.0 and Smart Factory initiatives demand data transmission speed, networking, storage capacity and processing power far greater than in the past for methods such as statistical process control and machine monitoring. To meet these needs, modern battery assembly systems are equipped with protocols such as EthernetIP, Profinet and Modbus TCP/IP. These Ethernet-based protocols offer ready-made solutions for network configurability, ease of integration, unparalleled data transmission speed and security, and the ubiquitous availability of network hardware and software.

● *Process monitoring*

Even after optimising a process, there are certain production tolerances and setup factors that will cause a weld defect. For laser welding, the key factors are part fit-up and maintaining the laser's focus precisely at the part weld interface. For resistance welding, electrode wear is the key culprit.

There is a comprehensive range of resistance welding process monitoring products that provide information independent of the power supply or weld head, regarding variances in the weld electrical properties, force and displacement. For example, it's possible to monitor the weld's dynamic resistance, voltage or current over the duration of the weld pulse. In addition, the rate of the weld collapse, as well as the overall weld collapse, can also be measured. All this information can be used for process monitoring. At present, data is collected, a set of waveforms defined and then limits or an envelope is placed around the waveform. Future technologies will likely offer better tools to define the good-weld/bad-weld decision, as well as predictive analysis.

Monitoring laser welding is more challenging, because it is a non-contact process without the benefit of electrodes touching. Typically, monitoring techniques capture certain wavelengths of light from the weld area that correspond to ultraviolet from the welding plume and keyhole, infrared from thermal heating of the part, and the back reflection of the laser light itself. With suitable sensing and analysis, these signals can be used as a basis for process monitoring. ●

AMADA MIYACHI SOLUTIONS

Amada Miyachi America offers a complete range of resistance welding power supplies and lasers specifically designed for both low- and high-volume battery pack manufacturing.

Technologies are matched to production needs as either standalone products or fully-integrated system solutions, which can be adapted to accommodate welding of dissimilar metals for battery tabs and tab design optimisation. An optimised software application, developed in Amada Miyachi America laboratories, is delivered with each system.

Figure 5 shows some examples of integrated systems for battery pack manufacturing, including a conveyor-fed automation cell, a laser tab welding system with fire suppression deployment, and a resistance welding system.

Figure 5: Integrated battery pack manufacturing solutions

NEXT-GENERATION GAN HEMT WITH BETTER GAIN AND EFFICIENCY

Wolfspeed launches next-generation 25W, 28V GaN HEMT that delivers better gain and efficiency.
The new 28V GaN HEMT devices are developed using Wolfspeed's proven 0.25µm GaN-on-SiC process, and are designed with the same package footprint as the previous generation of 0.4µm devices, making it possible for RF design engineers to use them as drop-in replacements for the earlier devices in existing designs.

Wolfspeed's CG2H40025 is an unmatched, gallium-nitride (GaN) high-electron-mobility transistor (HEMT). The CG2H40025, operating from a 28V rail, offers a general-purpose, broadband solution to a variety of RF and microwave applications. GaN HEMTs offer high efficiency, high gain and wide bandwidth capabilities, making the CG2H40025 ideal for linear and compressed amplifier circuits. The transistor is available in a screw-down, flange package and solder-down pill packages.

www.wolfspeed.com

VICOR EXTENDS 48V COOL-POWER ZVS BUCK REGULATORS FAMILY

Vicor has extended its Cool-Power 48V ZVS 20A buck regulator portfolio with its newly-released PI3523-00-LGIZ (PI3523). The PI352x family offers 20A solutions complementing the previously released 10A 48VIN PI354x family, enabling scalable power options for 48V direct to point-of-load (PoL) applications. The PI3523 is a 48VIN, 3.3VOUT nominal buck regulator capable of supplying up to 22A. This family of regulators enables 48VIN to 20A PoL voltages spanning 2.2-14VOUT.

Offering all the same industry leading features of Vicor's existing 48V Cool-Power ZVS buck regulators, the PI352x portfolio extends performance by delivering twice the power of the PI354x regulators using only a 40% larger package. The PI3523 requires only an output inductor and minimal passives for a complete cost-effective design that fits into a PCB space of 740mm².

http://www.vicorpower.com/new-products/cool-power-zvs-buck-regulator

READY-WIRED PANEL-MOUNT M16 CONNECTORS

Binder has announced new M16 connectors for panel mount applications, in addition to its other series, including the 581,680, 682, 423,723 and 425. The new connectors come ready-wired for time and cost savings, and are widely applicable across industrial control and measurement, medical and professional audio equipment.

The connectors can be supplied with four to eight contacts and single PVC 0.34mm², AWG 22/7 cUL Style 1007/1569 wires. Rated current is 5-7A and rated voltage 60-250V and, when mated, models meet the requirements of IP40 or IP67.

Binder offers these M16 panel mount connectors in both male and female versions, and as standard they are supplied with 200mm cable lengths, although other sizes are also available.

Binder is a family-owned company with a turnover in excess of 200m employing 1500 people worldwide.

www.binder-connector.co.uk

DIFFERENTIAL HIGH-ACCURACY, HIGH-ISOLATION, CURRENT SENSOR IC

Allegro MicroSystems Europe has introduced a new high-accuracy Hall-effect-based current sensor IC with multiple programmable fault levels, intended for industrial and consumer applications with a focus on motor control and power inverter stage applications.

One of the key benefits of Allegro's ACS720 is to provide high isolation with a reduced bill of materials, made possible by the proprietary IC SOIC-16W package. The ACS720 works from a single 5V supply while maintaining an output voltage swing from 0-3V, with a stable zero current output of 1.5V. This allows the ACS720 to operate from a 5V supply while having an output which is compatible with typical 3.3V ADCs found on many MCUs.
The ACS720's high PSRR rejects supply noise, maintaining high accuracy in noisy environments.

www.allegromicro.com

PICKERING INTERFACES AT EUROPEAN MICROWAVE WEEK 2017

Pickering Interfaces will show its broad range of PXI and Ethernet LXI RF and microwave switching solutions at European Microwave Week (EuMW) at Hall 7A, Stand 141 on 10-12 October 2017 in Nuremberg, Germany. Among the products on show will be its PXI microwave multiplexers (model 40-785B), 4-slot USB/LXI modular chassis (model 60-105), PXI microwave multiplexers (40-784A) and LXI microwave multiplexer (60-800).

Pickering will also highlight its eBIRST Switching System Test Tools and Switch Path Manager Signal Routing Software, designed specifically for its PXI, PCI or LXI (Ethernet) products. The tools simplify switching system fault-finding by quickly testing the system and graphically identifying the faulty relay. The Switch Path Manager speeds up the development of switching system software.

www.pickeringtest.com

NEBULA IOT DEVELOPMENT KIT AT FUTURE ELECTRONICS

Future Electronics has launched the Nebula IoT Development Kit with partners Cypress Semiconductor and Murata, available to buy now. The Nebula board is an IoT cloud-ready board that allows developers to quickly prototype and use IoT ecosystems. Wireless connectivity is supported by the Murata 1DX module, which is powered by the Cypress CYW4343W Wi-Fi and BT/BLE combo SoC.

The SoC includes a 2.4GHz WLAN IEEE 802.11 b/g/n baseband/radio and Bluetooth 4.2 support. In addition, it integrates a high-performance power amplifier (PA), a low-noise amplifier (LNA) for best-in-class receiver sensitivity and an internal transmit/receive (iTR) RF switch.

The board supports applications through the Cypress WICED (Wireless Internet Connectivity for Embedded Devices) Studio development platform.

www.futureelectronics.com

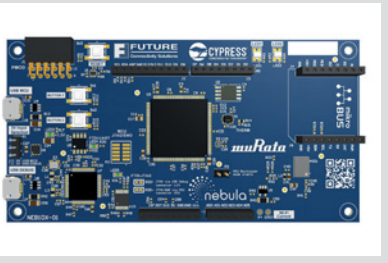

Pickering PXI Switching & Simulation

We Already Have Your Solution.

Designed for Today's Most Demanding Electronic Test & Verification Applications

- **Huge range** of choices ranging from **low level DC to 1000 V, RF/Microwave to 65 GHz**
- **High-quality solutions** for most test applications including Aerospace & Defense, Automotive, Medical, Sensor Emulation, Hardware-in-the-Loop and Semiconductor
- Typical 15 to 20 year full product support
- All modules plug into any PXI chassis and also in our Ethernet controlled LXI Modular Switching chassis with a choice of 18 and 7 slots, plus our new 2-slot USB/LXI chassis
- Full range of supporting cable and connector solutions plus new on-line Cable Design Tool
- **18+ years** in PXI with a strong worldwide customer base
- Learn more at **pickeringtest.com/pxi**

3 Year Warranty — We stand behind our products with a standard three year warranty.

Switching | Programmable Resistors | Instrumentation | Custom Design | Connectivity & Cables

Scan this QR Code or go to
pickeringtest.com/pxi
01255 687900 | uksales@pickeringtest.com

pickering — Pickering Interfaces

FORTE | THE INTELLIGENT BOM TOOL

Master your Bill of Materials with greater intelligence

mouser.com/bomtool

Authorised distributor of semiconductors
and electronic components for design engineers.

MOUSER ELECTRONICS