July/August 2017
Volume 123
Issue 1974
£5.90

www.electronicsworld.co.uk

Electronics WORLD

THE ESSENTIAL ELECTRONICS ENGINEERING MAGAZINE

Supercap Backup Circuit for Uninterrupted Power

By Linear Technology

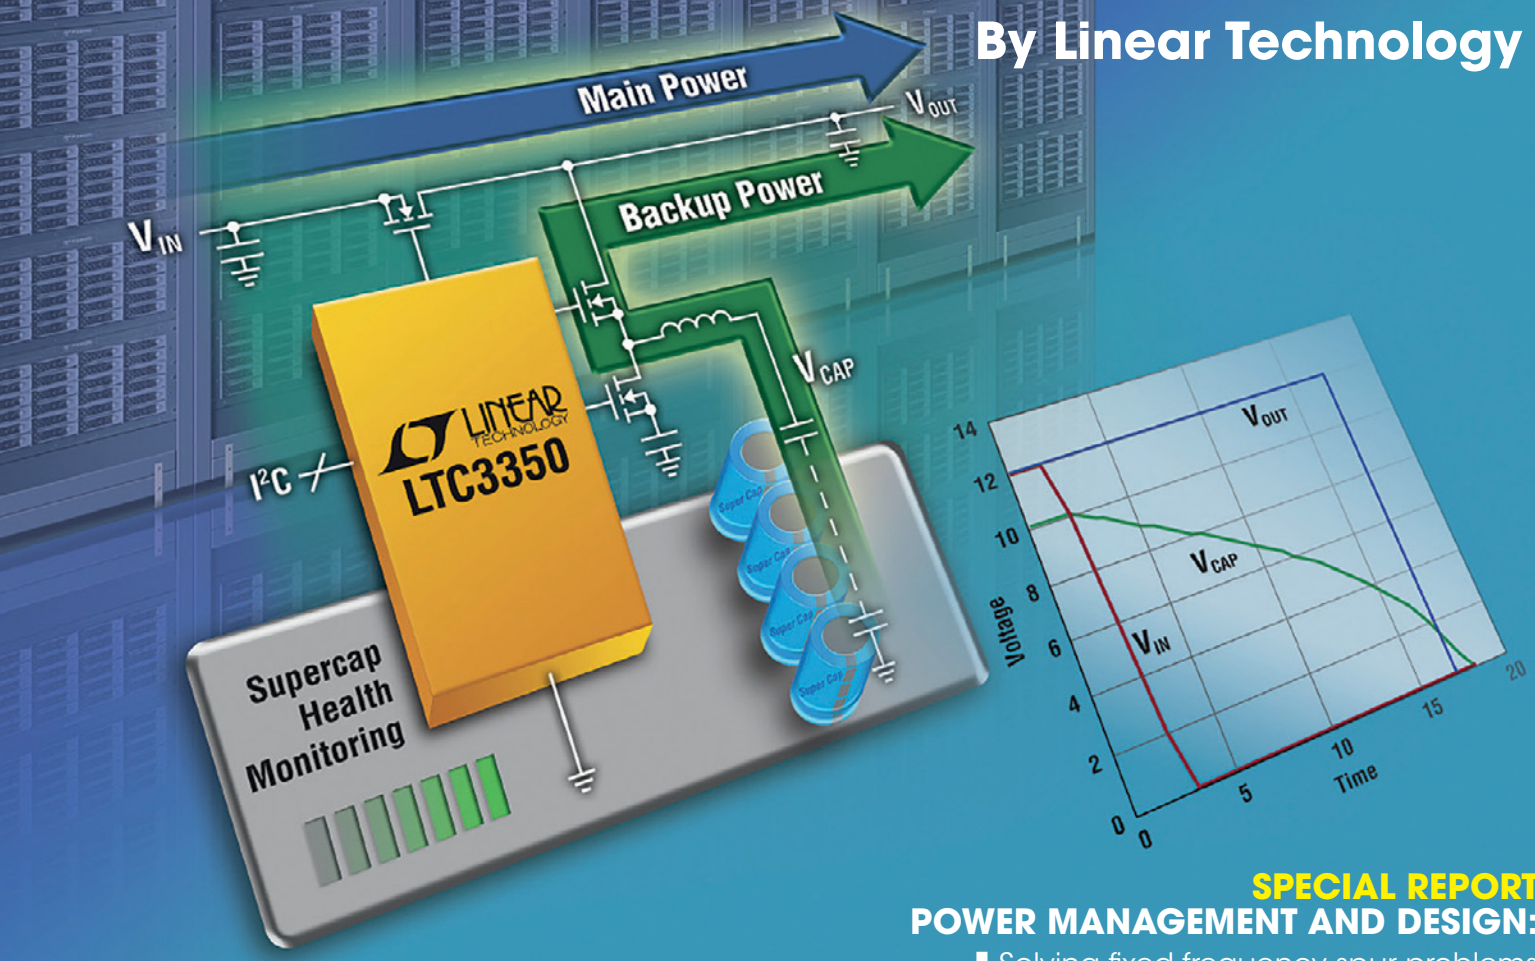

SPECIAL REPORT
POWER MANAGEMENT AND DESIGN:

- Solving fixed frequency spur problems
- Novel multilevel inverter topologies
- Energy-harvesting
- Cooling technologies

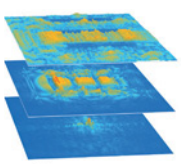

Technology
Leti develops shield that protects ICs from physical attacks

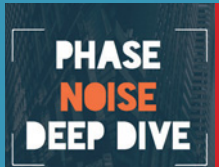

New Column
An in-depth guide to phase noise – from the basics to ways of improving it in designs

Products
High-end analogue RF and microwave signal generator

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

36

Cover supplied by LINEAR TECHNOLOGY NOW PART OF ANALOG DEVICES
More on pages 8-9

REGULARS

42

FEATURES

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

PIC18F "K40" MCUs feature Intelligent ADC with Filtering and Signal Analysis Capabilities

8-bit PIC® MCUs are ideal for Touch and Signal Conditioning

The Core Independent Peripherals (CIPs) in Microchip's PIC18F "K40" family of 8-bit PIC® MCUs support filtering and signal analysis for advanced touch and signal-conditioning applications.

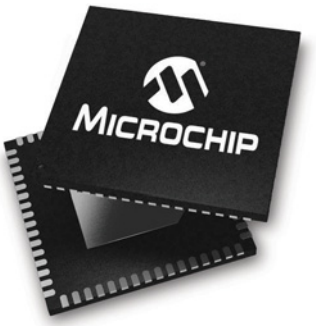

Among the intelligent analog CIPs is an Analog-to-Digital Converter with Computation (ADC2) for averaging, filtering, oversampling and automatic threshold comparison. The MCUs also integrate safety-critical CIPs and hardware PWMs with multiple communication interfaces and generous on-chip Flash and EEPROM. These features, combined with 5V operation, enable the PIC18F "K40" family to increase design flexibility whilst also reducing system cost.

microchip DIRECT
www.microchipdirect.com

MICROCHIP

www.microchip.com/EUPIC18FK40

The Microchip name and logo, the Microchip logo, MPLAB and PIC are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries. All other trademarks mentioned herein are the property of their respective companies. © 2016 Microchip Technology Inc. All rights reserved. DS30010138A. MEC2125Eng11/16

ELECTRONICS NEEDS THE FEMALE CONTINGENT

> The engineering sector needs to stop making things even harder for itself by marginalising its female members

Engineering has always been respected as a noble profession, relying on the endeavours of highly disciplined, conscientious and well-educated individuals. However, the appeal of this vocational path is on the wane.

Many engineers from the 'baby boomer' era are already reaching retirement age, leaving vacancies behind, a prospect that will only get worse over the next few years. Even worse, new vacancies will appear, created by exciting new opportunities created by IoT, home automation, Industry 4.0, renewable energy, autonomous driving and so on – all of which will need skilled workforce. As a result, many believe that Europe is about to face an unprecedented skill shortage. A recent study of the UK's engineering sector revealed that the country need at least 180,000 new engineers each year between now and 2022 to fill such positions, but only around half that figure leave the education system with the relevant training.

Other major economies across Europe are finding themselves in very similar situations. Last summer, when the EC set out its Skills Agenda for Europe, research compiled by Eurofound showed that roughly 40% of European employers are unable to recruit people with the right skills. The commission further estimates that there will be close to 7 million jobs in science technology engineering and mathematics (STEM) to be filled between now and 2025, yet in many European countries the number of students of these subjects is in decline.

With escalating demand but diminishing supply, can the electronics industry really keep on ignoring the potential of female recruits, or half the population? Currently even the most progressive of European countries cannot claim above 20% of its STEM jobs held by female employees – in the UK it is still below 10% – yet they have much to contribute.

On top of strong academic abilities, starting with girls normally outperforming boys in both secondary and tertiary education, it's no secret that girls are also better communicators. As a result, studies have shown that their inclusion in engineering teams has a positive impact on productivity – bringing closer cooperation between team members, better defined team objectives and quicker project completion.

Based on this, the industry needs to have a plan of action. It must start at the grass-roots level, by encouraging female pupils to consider careers in STEM. At the same time, companies should be taking a much more diverse approach when it comes to hiring, not simply sticking to their same old strategies. This will help improve team dynamics within their engineering workforce, allow new concepts to be experimented with and project ventures carried out much faster.

There also needs to be a change in how females in existing engineering roles are treated. I know from personal experience, having been in the technology business for over thirty years, that women have to work a lot harder to prove themselves, since their opinions are not always valued as much as that of their male counterparts, and their ideas not always given the serious consideration they deserve. The time has come for this to change!

The engineering sector already faces many difficult challenges, so it must stop making things even harder by marginalising its female contingent.

Moving forward, the efforts of male and female engineers need to be recognised accordingly. This will lead to greater diversity, strengthened numbers and more effective collaboration.

By Françoise Chombar, CEO of Melexis (www.melexis.com)

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

LETI INVENTS A SHIELD THAT PROTECTS MICROCHIPS FROM PHYSICAL ATTACKS

France-based technology research institute Leti has developed a shield that can help protect electronic devices against physical attacks from the back of the chips.

Physical attacks can occur when hackers have access to the device and exploit weaknesses of its embedded ICs to steal sensitive information or to corrupt their functioning. The shield proposed by Leti protects chips from invasive and semi-invasive attacks by infrared lasers, focused ion beams (FIB), chemicals and other means.

The shield consists of a thin, serpentine-shaped metal, sandwiched between two polymers, one opaque to infrared light and serving as a physical barrier against FIB attacks. It also hides the design of the chip's serpentine and combines with the polymer underneath to detect chemical attacks. Altering the serpentine typically triggers the IC to delete sensitive data.

The shield is fabricated using standard packaging processes, which demonstrates that hardware cybersecurity can be implemented at low additional cost.

"Implementation of multiple hardware and software countermeasures is making integrated circuits more secure, but the back of the chip is still considered vulnerable to physical attacks," said Alain Merle, Leti's Security Strategic Marketing Manager. "Our team designed, fabricated and tested a novel protection structure combining several elements that will trigger an alert if hackers use the back of the chip to access the active parts of the IC."

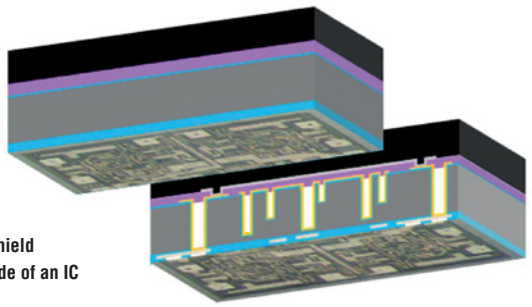

View of the security shield placed on the underside of an IC

NEW RESEARCH REVEALS POWER ENGINEERS STRUGGLE WITH SPEC CHANGES

The biggest device-power design challenges identified by engineers are downstream changes to the specifications after the design process has started. According to research conducted by Vicor Corporation, 87% of power system designers around the world identified specification changes as the most frequent impediment in designing power systems.

"When creating power system designs, engineers need to consider more flexible approaches to contend with specification changes," said Rob Russell, VP of Product Marketing at Vicor.

The research also highlights the challenges of meeting project timelines. Some 80% of power engineers are already struggling to meet time constraints when designing power systems, and 79% feel that project schedules are tighter due to time-to-market pressures. Spec changes negatively impact schedules, say 65% of the surveyed engineers.

Vicor's in-depth research identifies the biggest problems of designing power systems in an increasingly demanding business environment.

Rittal – The System.

Faster – better – everywhere.

Box smarter with Rittal

- Superior quality
- Extensive size range
- Comprehensive product options
- Instant availability
- Competitive pricing

ENCLOSURES | POWER DISTRIBUTION | CLIMATE CONTROL

FRIEDHELM LOH GROUP

WIFI-ROUTER-ENABLED HOLOGRAPHY MAY HELP TRACK DEVICES AND PEOPLE

Scientists at the Technical University of Munich (TUM) have developed a holographic imaging process that depicts the radiation of a Wi-Fi transmitter to generate three-dimensional images of the surrounding environment. Among many other applications, industrial facility operators could use this to track objects as they move through the production hall.

While optical holograms require elaborate laser technology, generating holograms with the radiation of a Wi-Fi transmitter requires merely one fixed and one movable antenna. This will be particularly useful in Industry 4.0, where automated industrial facilities will need a greater ability to localise parts and devices.

There are already processes that allow the localisation of a microwave radiation source, even through walls, or in which changes in a signal pattern show the presence of a person; what's new is that an entire space can be imaged via holographic processing of Wi-Fi or mobile phone signals.

Generating images from microwave radiation requires special-purpose, large-bandwidth transmitters. Using holographic data processing, the very narrow bandwidths of typical household Wi-Fi transmitters operating in the 2.4GHz and 5GHz bands were sufficient for the researchers.

"One can use a larger number of antennas to obtain a video-like image," says Philipp Holl, who executed the experiments. "Future Wi-Fi frequencies, like the proposed 60GHz IEEE 802.11 standard will allow resolutions down to the millimetre range."

Going forward, microwave images can be combined with that of a camera, to help trace a radio tag attached to a lost item, for example.

Further advances in this technology may aid in the recovery of victims buried under snow or a collapsed building. While conventional methods only allow point localisation of victims, holographic signal processing could provide a spatial representation of destroyed structures, allowing first responders to navigate around heavy objects and use cavities in the rubble to find the easiest approach to quickly reach victims.

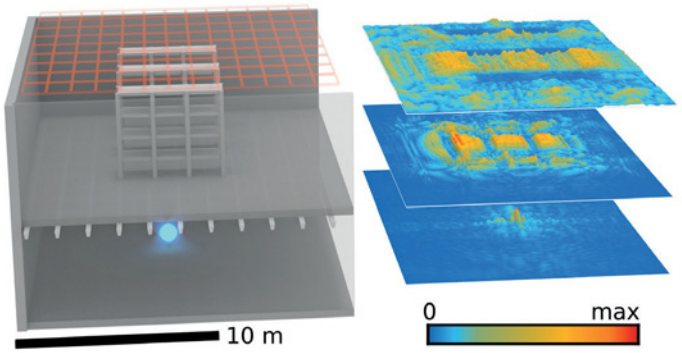

An entire space can be imaged via holographic processing of Wi-Fi or mobile phone signals

SUPERCAP BACKUP CIRCUIT PROVIDES RELIABLE UNINTERRUPTED POWER

By Samuel Nork, Director, Boston Design Center, Linear Technology, now part of Analog Devices

Temporary backup power is a common requirement for a wide range of applications whenever the main power source is suddenly unavailable. Examples include data backup applications ranging from servers to solid-state drives, power fail alarms in industrial or medical applications, and a host of other "dying gasp" functions where orderly power-down must be assured and system status communicated to a powered host. In the past, these types of high reliability systems used batteries to provide an uninterrupted power source whenever the main supply of power was inadequate or unavailable. However, many trade-offs accompany battery backup, including long charge times, limited battery lifetime and cycle life, safety and reliability concerns, and large physical size. With the advent of high value electric double layer capacitors, better known as supercapacitors, alternate backup architectures may be employed which eliminate many of these tradeoffs.

Batteries vs. Capacitors

Systems relying on batteries for backup power require that a fully charged battery is available at all times with suitable capacity to keep volatile memory alive or alarms sounding until power is restored. Typically, systems employing battery backup enter a low-power standby state whenever the main power fails, and only the critical volatile memory or alarm sections of the systems remain powered. Since power failure duration is impossible to predict, such systems require oversized batteries to avoid the possibility of data loss during a lengthy outage.

Capacitor-based backup systems use a different methodology. Unlike battery-based systems which provide continuous power during the entire backup time, capacitor-based systems require only short-term backup power in order to transfer volatile data into flash memory or provide "dying gasp" alarm operation for a minimum necessary amount of time. Once the required data has been saved and the power fail alarms have been properly issued, the power restoration time is unimportant.

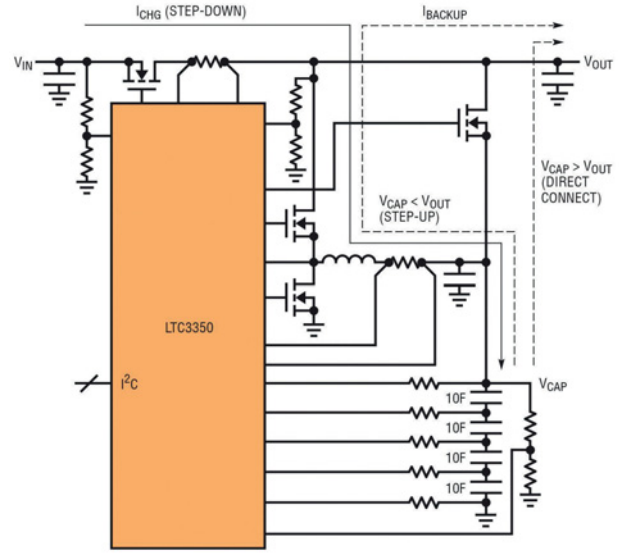

Figure 1: High Current Supercapacitor Charger & Backup Controller

With the emergence of small, relatively inexpensive supercapacitors capable of storing numerous Joules of energy, the number of backup applications that can be satisfied with capacitors instead of batteries has grown considerably.

All capacitor-based backup systems share many common elements. PowerPath™ control and power fail detection are required to supply power to the load from the proper source and to alert the system when transitioning from normal operation into backup mode. The storage capacitor needs to be charged, and ideally this is done in a fast, efficient manner. Since proper backup is not possible unless an adequate number of Joules are stored on the backup capacitor, many applications require that charging is completed by the time the system boots up and is ready for operation. Hence, high charge currents are generally desirable, and since supercapacitors typically have a max operating voltage of 2.7V, it is common and often necessary for several to be stacked in series. In such cases, provision must be made for balancing and protecting the capacitors as they charge to prevent damage and lifetime degradation due to overvoltage.

Figure 1 shows a simplified schematic for the LTC3350, a capacitor charger and backup controller IC designed specifically to address capacitor backup applications. The LTC3350 includes all the features necessary to provide a complete, standalone backup controller for applications needing capacitor-based backup. The device can charge, balance and protect up to four capacitors in series. Input power fail threshold, capacitor charge voltage and regulated minimum backup voltage can all be programmed with external resistors. In addition, the device contains a very accurate 14-bit internal measurement ADC which monitors input, output and capacitor voltage and current. The internal measurement system also monitors parameters associated with the backup capacitors themselves including capacitor stack voltage, capacitance and stack ESR (Equivalent Series Resistance). All system parameters and fault status can be read back over a two-wire I²C bus, and alarm levels can be set to alert the system to a sudden change in any of these measured parameters.

The LTC3350 Supercharger

The charging circuitry in the LTC3350 consists of a high current, synchronous buck controller with a resistor programmable max charge current and max stack voltage (Figure 2). Since the charger is powered from the same supply that is powering the

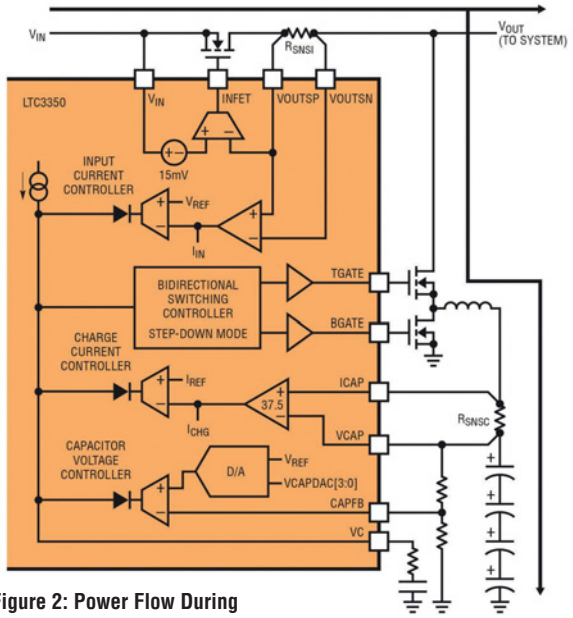

Figure 2: Power Flow During Normal Operation

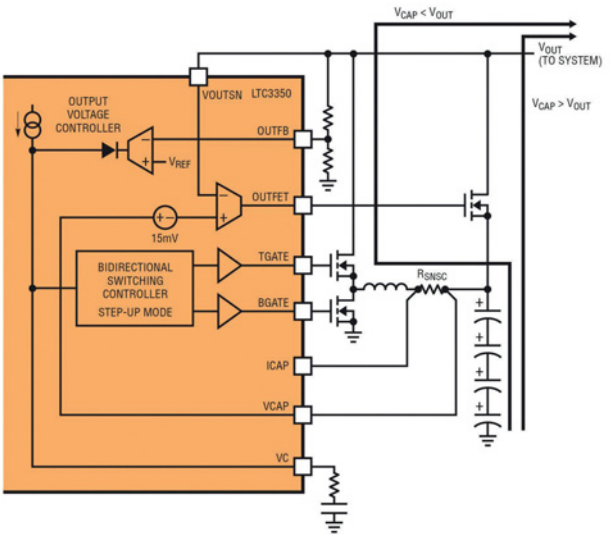

Figure 3: PowerPath Operation in Backup Mode

load, the LTC3350 also contains a separate programmable input current limit which automatically reduces charge current to the capacitors under heavy V_{OUT} load conditions. Internal, low current balancers force all cells to within 10mV of each other up to a max voltage of 5V per cell. Internal protection shunts will automatically reduce charging current and shunt the remaining charge current around any capacitor that has reached the 2.7V default or a user-programmed max cell voltage. In addition, the stack charge voltage may be reduced under software control in order to optimize capacitor lifetime for a given backup energy requirement.

Backup Mode

Once the backup capacitor stack is charged, the system is able to provide backup power. Charge mode and backup mode are determined by the voltage on the PFI (Power Fail Input) pin. If the V_{IN} voltage drops such that the PFI comparator trips low, the part immediately enters backup mode (see Figure 3). V_{OUT} will drop as V_{IN} drops, and once the V_{OUT} voltage falls below the capacitor stack voltage, the OUTFET ideal diode conducts to prevent V_{OUT} from falling further. Once V_{OUT} falls to a voltage programmed by a resistor divider on the OUTFB pin, the capacitor charger operates in the opposite direction as a synchronous boost backup DC/DC converter using the V_{CAP} stack as its input source and V_{OUT} as its regulated output. The boost backup converter will continue to run until it can no longer support the V_{OUT} load conditions and the voltage on V_{OUT} falls below the 4.5V UVLO point. This allows virtually all of the usable energy in the supercap stack to be transferred to the load during backup since the boost will continue to run when the stack voltage is well below 4.5V. A typical backup scenario is also shown in Figure 3. In this example, a stack of four series capacitors is charged to 10V, and during backup mode V_{OUT} is regulated to a minimum of 8V until all energy is depleted from the backup capacitors.

"Health" Monitoring Assures Reliability & Optimizes Performance

In high reliability systems requiring short-term backup power, adequate energy must be stored and available in order to perform critical functions immediately following a main power failure. It is essential that the backup energy source is able to deliver the necessary backup power. Supercapacitors are an excellent choice for such applications due to their extremely high capacitance per unit volume and very low ESR. However, like batteries, their performance will degrade over time. Capacitor lifetime is commonly (and somewhat arbitrarily) defined as the time required for capacitance to drop by 30% and/or ESR to increase by 100. Capacitor degradation is accelerated by either high operating voltages or elevated temperatures. Since both capacitance and capacitor ESR are critical for ensuring that the system can perform a reliable backup, it is important that the system is able to monitor and report the "health" of the backup capacitors as they age.

The LTC3350 automatically monitors both stack capacitance and stack ESR at a time frequency chosen by the user once the capacitor stack is fully charged. The part employs a precision current source, precision timing circuit and its internal 14-bit ADC to accurately monitor the stack capacitance. A precise, programmed current is pulled from the top of the capacitor stack while the charger is forced off. The time required for the capacitor stack to drop by 200mV is precisely measured, and the stack capacitance is calculated from these parameters. Once the capacitance test is completed, the ESR test is done by measuring the stack voltage with and without the high current charger running to recharge the stack.

Once the stack capacitance and ESR values are known, it is straightforward to compute the minimum stack voltage necessary to assure a reliable backup for a given application.

Linear Technology (UK) Ltd • Tel: 01628 477066
Email: uksales@linear.com • www.linear.com

eXtremely (normal) low power

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

In the history books of embedded control, right after the Flash Revolution, came the Low Power Age. The advent of the MSP430, in the early years of the new millennium, focused the attention of the entire microcontroller industry on power consumption. A myriad of new techniques were introduced – and with them new acronyms, such as Ultra Low Power (ULP), nanoWatt, picoWatt and, eventually, eXtreme Low Power (XLP). Soon the marketing departments of all the major vendors were busy counting the new low-power modes, painting all shades of grey between the full-on 'RUN' mode and the 'off 'state or 'SLEEP' mode.

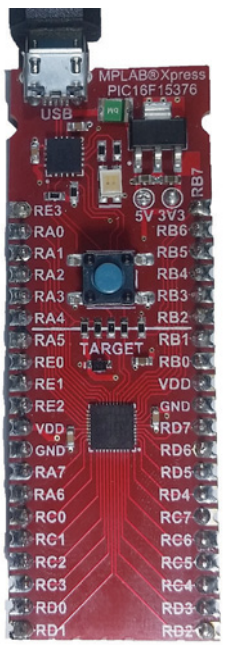

Figure 1: PIC16F15376 XPRESS evaluation board

However, the rules of physics are the same for all CMOS devices. Microcontroller "power modes" are just changing the contribution of static and dynamic power consumption components that can be represented by this simplified formula:

$$I_{tot} = I_s(V) + I_d(V) * F_{clock}$$

The first term (I_s) is the static current consumption. It represents the total contribution of leakage currents (resistive) characteristic of the fabrication process used. It is (inversely) proportional to the geometry (thickness of the isolation layers). It also grows linearly with voltage and exponentially with temperature.

The dynamic current consumption term (I_d) grows similarly with temperature and voltage, but depends on capacitive (gate) parasitic effects, and its contribution is multiplied by the operating clock frequency.

Naturally, to compute the power consumption we need to multiply the total current by the supply voltage, which makes the voltage contribution to both power components (dynamic and static) quadratic! This is the reason why much of the focus in low-power designs is on reducing the operating voltage whenever possible.

It must be noted that the limiting factor here is rarely the microcontroller itself (which would happily operate at 1.8V, for example), but rather what goes around it. For noise reasons, a 5V supply is often preferred in automotive and industrial applications, or a 3V (regulated) supply might be required for wireless modules and/or displays.

When lithium batteries are used as power sources, the low-voltage operating limit of the microcontroller is rarely the true limiting factor. Most such batteries' output will fall off a cliff well before 2.5V and, by the time 2V is reached, primary protection mechanisms would have kicked in.

Too Little To Measure

There are many myths around ultra-low power; some are so extreme they are hilarious. For example, when bragging

about the extreme (low) static power consumption of a given microcontroller family, vendors will often refer to the number of years of battery life an application could supposedly achieve. You might have heard of claims of battery life of 10 or 20 years – from a coin cell! Unfortunately, this is solely based on the arithmetic of battery capacity divided by static current consumption. This might have been fine to compute battery life when devices had static current consumptions of the order of tens of microamperes; but today, most 8-bit PIC microcontrollers have static current consumption of less than 50nA, making the battery's "self-discharge effect" approximately equivalent to 1uA for common alkaline batteries – a dominant limiting factor.

An excellent article written by Jack Ganssle in his muse #249 does an excellent job of presenting this and many similar misconceptions!

Too Many Modes

If in the early days of the Low Power Era counting low-power modes was a thing, and since only two or three control knobs where available to the designer to tune an application, it produced four to eight low-power modes. Today the situation is a lot more complex. As an example, recent members of the PIC microcontroller family have the following standard set of low-power options: SLEEP, IDLE, DOZE, LDO, DCO(x3), PLL, BOR, FSCM. As we compute the number of modes (10 bits giving 1024 unique combinations), even excluding the odd and improbable cases, it can easily run in the hundreds.

Fact is, there is no such a thing as a typical low-power application. All applications require special considerations, and more options is always better than fewer. So rather than reducing the conversation to a few modes, I will try to review some of the fundamental ideas behind the latest and greatest features in a modern microcontroller.

More precisely, in the following examples I will refer to latest-generation PIC microcontrollers, easily identified by the part number PIC16F1 followed by four more digits. I will use data from the latest PIC16F153xx family datasheet and/or occasionally measured by me on a breadboard, using a low-priced MPLAB XPRESS series board.

DCO Vs. Dividers

If old microcontrollers had one or two oscillator options, modern versions offer at least a dozen internal and external ones. Further, if you are looking for low power and low cost, it won't be enough to use an internal (precision) oscillator, you will want to use a digitally-controlled oscillator (DCO) instead.

DCOs produce many different output frequencies (with tight tolerances) using the minimum possible current without using "dividers". Table 1 clarifies the difference.

Even though the values in Table 1 were obtained empirically on a single sample (PIC16F15376), and therefore cannot be claimed to represent a statistically valid typical value, I believe the vast difference in power consumption achieved with a dedicated 1MHz DCO over the 16MHz oscillator (divided by 16) is so large (2.5x) as to leave no doubts. The PIC16F153xx series DCO, for example, can produce six frequencies: 1MHz, 2MHz, 4MHz, 8MHz, 16MHz and 32MHz, efficiently covering a wide and common range of clock frequencies.

Low-Power Oscillators

To reach much lower active power consumption figures though, the clock frequency must be reduced further and two special oscillators can be used. If accuracy is not of primary importance, the internal low-power oscillators (~31kHz) are the best option, giving a typical value of just 500nA at 3V. But if accuracy is required, to support correct timekeeping or some form of serial communication, then an external secondary oscillator (SOSC), with an inexpensive 32kHz crystal, will deliver the best performance at the cost of just 800nA.

PLLs

A phase locked loop (PLL) circuit is a frequency multiplier, often used to extend the frequency range of an oscillator. PLLs do, however, add a bit of power consumption on their own, so it's always worth checking if a DCO/oscillator with the desired natural frequency is available first.

Table 2 shows bench measurements on a PIC16F15376 device, illustrating the effect. It can be seen that the PLL applied to the DCO operating at 8MHz gives a 16MHz output with a slightly higher current consumption. Just 14uA would have been a small price to pay if the 16MHz native option was not available.

ADCs And Power-Hungry Peripherals

The data collected in the examples here was measured while the device was running (executing code in an empty 'while' loop) without any peripherals active. While each active microcontroller peripheral will add some current consumption, only a few are significant contributors. Analogue peripherals are among

CLOCK	CURRENT CONSUMPTION	CONDITIONS
16MHz	901uA	3V
1MHz	537uA	3V, 16MHz/16 divider
1MHz	215uA	3V, 1MHz DCO

Table 1: DCO outputs

CLOCK	CURRENT CONSUMPTION	CONDITIONS
16MHz	901uA	3V RUN
16MHz	915uA	3V, 2 x 8MHz, RUN

Table 2: Bench measurements on a PIC16F15376

PERIPHERAL	CURRENT ADDER	CONDITIONS
ADCC	250uA	3V
Analogue comparator	25uA	3V
FVR	10uA	3V

Table 3: Type of peripherals

the worst offenders, and in particular the ADCs, operational amplifiers (if present) and analogue comparators; see Table 3.

The obvious choice is to keep such peripherals disabled most of the time, enabling them only when needed. This requires some serious considerations for the startup or settling time, the time required for the peripheral to become available and to produce a reliable/stable output.

Peripheral Module Disable

All peripherals – analogue and digital – will use some power, even when not enabled! This is because the microcontroller clock signal must reach each peripheral's control registers (and most of its internal logic), even when the peripheral is not in use. In the new PIC microcontrollers, there is a fine control of the clock "branches" via a mechanism called Peripheral Module Disable, or PMD. When a peripheral PMD bit is set, the peripheral effectively disappears!

As shown in Table 4, the impact of the PMD (used here to disable all device peripherals) changes with the clock frequency and voltage. Comparing row 1 and 2, we notice that at 3V and a relatively high frequency (2MHz) the PMD can reduce power consumption by as much as 25%. At the lower frequency (500kHz) and voltage (2V) the difference (between rows 5 and 6) is just 7%. So, as expected, the PMD produces the greatest effect when the clock frequency is the highest, which is perhaps when it's needed the most.

Robustness Vs. Low Power

Robustness and safety features such as the Fail-Safe Clock Monitor (FSCM), Windowed Watchdog (WWDT) and Brown-Out Reset (BOR) are all contributing to the power consumption of the microcontroller. But a BOR circuit can be automatically disabled when the device goes into a sleep mode. As long as the lower Power On Reset (POR) threshold is not passed, a safe minimum voltage needs to be re-checked only upon device wakeup, saving precious current (0.8uA typ) during the sleep phase.

Similarly, the watchdog circuit can be used to wake the device at regular intervals during sleep, but this will cost a bit of current (0.6uA typ) that can be avoided, once more automatically, by disabling the WWDT feature upon entry into sleep mode.

TEST	CLOCK	CURRENT CONSUMPTION	CONDITIONS
1	2MHz	262uA	3V RUN
2	2MHz	216uA	3V, RUN + PMD
3	2MHz	174uA	2V RUN
4	2MHz	147uA	2V RUN + PMD
5	500kHz	115uA	2V RUN
6	500kHz	108uA	2V RUN + PMD

Table 4: PMD impact

Of Sleeping And Awakenings

Most microcontrollers have a mode where both CPU and peripherals are completely stopped, commonly called SLEEP. This mode has been refined in recent years to account for new 'tricks' required by the smallest geometries.

Deep-sleep modes have been introduced to further reduce power consumption by removing power entirely from special islands of power, often starting with the RAM memory array. For large (16/32-bit) devices this can be the only way to keep static current below the 1uA threshold. The price is that by losing most or all the RAM contents, the application will have to recover its "state" based on saved information, in EEPROM or in a special tiny section of RAM that is reserved for the purpose. This eventually translates into more time needed during "awakening", which decreases the benefits of the deeper slumber. Luckily, this is totally unnecessary on most small 8-bit microcontrollers, which can easily achieve <200nA power consumption without losing any RAM contents.

A special case must also be considered for those devices that can/must operate in 5V environements. In the latest generations, most such devices employ an internal voltage regulator (LDO) to keep the core running at <2.5V (for lower power consumption) while the I/O structures are kept at full Vdd. Such LDOs, however, can become a liability once in sleep. Their quiescent currents ~10uA (required to achieve strong regulation during RUN time) can seem huge when compared with nA sleep modes. So, modern microcontroller families provide special modes for LDOs themselves, reducing the LDO quiescent current in sleep by several orders of magnitude. In the PIC16F153xx family, for example, look for the VREGPM configuration bit – and make sure it is set before going to sleep!

Another important consideration is required when we want to re-start a precision (external crystal) oscillator upon awakening from sleep. Such oscillators often require thousands of cycles before they reach their nominal frequency. This can translate into a long wait (up to several milliseconds), during which the processor is drawing power but unable to perform any activity. This can wreak havoc on an application's power budget.

All PIC16F1 and PIC18 microcontrollers can perform what is called a "Two-Speed Startup". It means, they can wake up on an internal oscillator (non-accurate, but low-power and fast), and immediately start being productive. Only when the external (accurate crystal) oscillator is finally stabilised, they switch automatically to it and are ready to perform the activities (often communication) that require high-precision timing.

CPU Speed Vs. Peripheral Speed

Then there are the applications where processor performance and speed are not critical but the performance of peripherals matters. Using IDLE mode is an effective but drastic way of keeping the peripherals running while the CPU is completely stopped.

Modern microcontrollers (both PIC18 and PIC16F1) also offer a DOZE mode, where the main clock is set to the required (higher) frequency by the peripherals to achieve a desired resolution (PWMs, timers, etc.) while the CPU is not stopped but given a fraction of the clock, say from 1/2 to 1/512th. This feature can be used to optimise the power budget to a very fine degree. For example, the CPU clock can be divided during normal operation (in the main loop) but, when needed, interrupts can be used to awaken it (full-speed operation). To achieve this, simply set the Resume On Interrupt (ROI) bit in the CPUDOZE register.

Additionally, CPU full-speed operation can be restored during the interrupt, but can immediately and automatically return to the lower speed as soon as the interrupt service routine (ISR) is exited. To achieve this, set the Doze On Exit (DOE) bit in the CPUDOZE register as well.

The main benefit of using such DOZE features is that all application timings can be computed for a single clock – the peripheral's full-speed clock. All CPU speed changes are completely automatic and transparent to the application. As I often confess to being quite the lazy developer, the simplification introduced by the DOZE modes is a great comfort, since it guarantees useful power-saving for very little design effort.

Hands-On Experience Matters

For all the tools and features listed above, though, nothing beats some real hands-on low-power design experience. It is way too easy to spend hours poring over tables and interpolating values from the datasheet (AC characteristics tables), only to forget the proper I/O configuration of some unused pins (to make them into analogue inputs by setting the proper ANSEL registers) or, even more easily, ignoring the leakage of capacitors (electrolytics are the worst offenders) and/or the impact of a pull-up resistor here and there.

Low-power design is an iterative process, often involving a good deal of trial and error. Modern PIC microcontrollers make eXtreme Low Power not just possible, but quite normal, I'd say! ●

FANTASTIC MODERN POWER SUPPLY ONLY IU HIGH PROGRAMMABLE

LAMBDA GENESYS	PSU GEN100-15 100V 15A Boxed As New	£325
LAMBDA GENESYS	PSU GEN50-30 50V 30A	£325

IFR 2025	Signal Generator 9kHz - 2.51GHZ Opt 04/11	£900
Marconi 2955B	Radio Communications Test Set	£800
R&S APN62	Syn Function Generator 1HZ-260KHZ	£195
HP3325A	Synthesised Function Generator	£195
HP3561A	Dynamic Signal Analyser	£650
HP6032A	PSU 0-60V 0-50A 1000W	£750
HP6622A	PSU 0-20V 4A Twice or 0-50V 2A Twice	£350
HP6624A	PSU 4 Outputs	£350
HP6632B	PSU 0-20V 0-5A	£195
HP6644A	PSU 0-60V 3.5A	£400
HP6654A	PSU 0-60V 0-9A	£500
HP8341A	Synthesised Sweep Generator 10MHZ-20GHZ	£2,000
HP83731A	Synthesised Signal Generator 1-20GHZ	£1,800
HP8484A	Power Sensor 0.01-18GHZ 3nW-10uW	£75
HP8560A	Spectrum Analyser Synthesised 50HZ - 2.9GHZ	£1,250
HP8560E	Spectrum Analyser Synthesised 30HZ - 2.9GHZ	£1,750
HP8563A	Spectrum Analyser Synthesised 9KHZ-22GHZ	£2,250
HP8566B	Spectrum Analsyer 100HZ-22GHZ	£1,200
HP8662A	RF Generator 10KHZ - 1280MHZ	£750
Marconi 2022E	Synthesised AM/FM Signal Generator 10KHZ-1.01GHZ	£325
Marconi 2024	Synthesised Signal Generator 9KHZ-2.4GHZ	£800
Marconi 2030	Synthesised Signal Generator 10KHZ-1.35GHZ	£750
Marconi 2305	Modulation Meter	£250
Marconi 2440	Counter 20GHZ	£295
Marconi 2945/A/B	Communications Test Set Various Options	£2,000 – £3,750
Marconi 2955	Radio Communications Test Set	£595
Marconi 2955A	Radio Communications Test Set	£725
Marconi 6200	Microwave Test Set	£1,500
Marconi 6200A	Microwave Test Set 10MHZ-20GHZ	£1,950
Marconi 6200B	Microwave Test Set	£2,300
Marconi 6960B with	6910 Power Meter	£295

MARCONI 2955B Radio Communications Test Set - £800

CAN BE SUPPLIED WITH OPTIONAL TRANSIT CASE

PROPER 200MHZ ANALOGUE SCOPE - £250

FLUKE/PHILIPS PM3092 Oscilloscope 2+2 Channel 200MHZ Delay TB, Autoset etc

Tektronix TDS3052B/C	Oscilloscope 500MHZ 2.5GS/S	£1,500
Tektronix TDS3032	Oscilloscope 300MHZ 2.5GS/S	£995
Tektronix TDS3012	Oscilloscope 2 Channel 100MHZ 1.25GS/S	£450
Tektronix 2430A	Oscilloscope Dual Trace 150MHZ 100MS/S	£350
Tektronix 2465B	Oscilloscope 4 Channel 400MHZ	£600
Farnell AP60/50	PSU 0-60V 0-50A 1KW Switch Mode	£195
Farnell H60/50	PSU 0-60V 0-50A	£500
Farnell XA35/2T	PSU 0-35V 0-2A Twice Digital	£75
Farnell LF1	Sine/sq Oscillator 10HZ-1MHZ	£45
Racal 1991	Counter/Timer 160MHZ 9 Digit	£150
Racal 2101	Counter 20GHZ LED	£295
Racal 9300	True RMS Millivoltmeter 5HZ-20MHZ etc	£45
Racal 9300B	As 9300	£75
Fluke 97	Scopemeter 2 Channel 50MHZ 25MS/S	£75
Fluke 99B	Scopemeter 2 Channel 100MHZ 5GS/S	£125
Gigatronics 7100	Synthesised Signal Generator 10MHZ-20GHZ	£1,950
Seaward Nova	PAT Tester	£95
Solartron 7150/PLUS	6 1/2 Digit DMM True RMS IEEE	£65/£75
Solatron 1253	Gain Phase Analyser 1mHZ-20KHZ	£600
Tasakago TM035-2	PSU 0-35V 0-2A 2 Meters	£30
Thurlby PL320QMD	PSU 0-30V 0-2A Twice	£160-£200
Thurlby TG210	Function Generator 0.002-2MHZ TTL etc Kenwood Badged	£65
HP33120A	Function Generator 100 microHZ-15MHZ	£260-£300
HP53131A	Universal Counter 3GHZ Boxed unused	£500
HP53131A	Universal Counter 225MHZ	£350

INDUSTRY STANDARD DMM ONLY £325 OR £275 WITHOUT HANDLE AND BUMPERS

HP 34401A Digital Multimeter 6 ½ Digit

YES! AN HP 100MHZ SCOPE FOR ONLY £75 OR COMPLETE WITH ALL ACCESSORIES £125

HP 54600B Oscilloscope Analoque/Digital Dual Trace 100MHZ

STEWART OF READING

17A King Street, Mortimer, near Reading, RG7 3RS

Telephone: **0118 933 1111** Fax: **0118 9331275**

USED ELECTRONIC TEST EQUIPMENT

Check website **www.stewart-of-reading.co.uk**

(ALL PRICES PLUS CARRIAGE & VAT)

RFID systems in embedded applications

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

Radio frequency identification (RFID) is based on the transfer of data using electromagnetic fields to identify tags attached to objects. Information about an object is stored on a carrier device known as a tag or transponder, which consists of a built-in antenna and a chip containing the identifier data; see Figure 1. The data can either be a simple unique serial number, or more information about the object. As with a barcode, identification requires an electronic reader, but unlike barcode systems, the RFID reader can be in any orientation and not in a line of sight.

There are two types of tags: passive and active. Passive tags have no power, obtaining their energy from the electromagnetic field of the reading device, making them small, economical and maintenance-free. On the other hand, because there is no power source, their detection ranges are rather limited.

Active tags typically use batteries, and can operate hundreds of metres from the reader. However, this also makes them larger and more expensive than passive tags. Active tags are usually used with large objects where tracking is required at long distances.

RFID Versatility

RFID tags can be used to identify an array of objects or even animals, but are more commonly used in access control. Other applications include identifying railway tracks a train is travelling on. Here, the RFID tag is mounted on the tracks with the reader inside the train; reading is possible at very high speeds.

Passive RFID tags are available in various formats and shapes (Figure 2), and can be embedded in various materials including glass, plastic and paper.

They offer several advantages compared to barcodes:

- Operatation in harsh environments;
- They cannot be contaminated by dust or oil;
- They can store additional data;
- Read and write functionality;
- Better security.

RFID Operating Frequencies

RFID systems come in three frequency ranges:

- Low frequency (LF): 30-300kHz. Typically, these systems operate at 125kHz, though some operate at 134kHz. They have low data rates, but work well in water and through most materials. The detection range is around 10cm. They are used in animal identification, access control and factory data collection.
- High frequency (HF): 13.56MHz. These systems have greater data rates and longer detection ranges (10cm to 1m). They are commonly used in smart cards.
- Ultra-high frequency (UHF): 850-960MHz. These systems

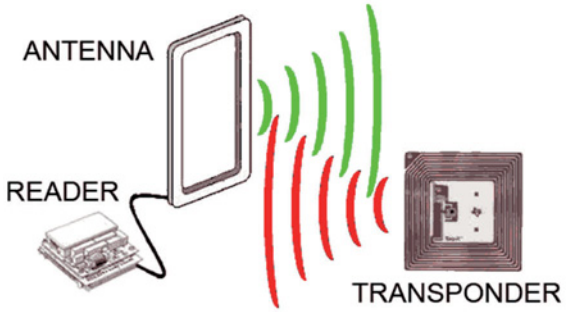

Figure 1: Typical RFID system

Figure 2: Passive RFID tags

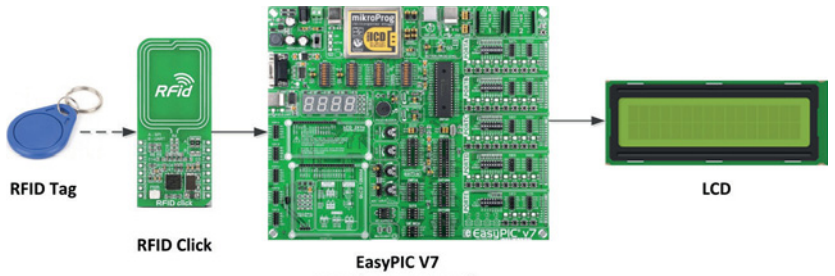

Figure 3: Block diagram

have much higher data rates and longer detection ranges (up to 10m); however, their signals do not pass through materials. The passive UHF variety are perhaps the most commonly used RFID systems with low-cost tags. Active RFID systems, with their much longer range, typically operate in the UHF range.

Selection Criteria

Selecting the correct RFID system for an application depends on many factors:

- Tag reading distance;
- Memory size needed;
- Memory type (read/write);
- Material used;
- Required level of security;
- Size of the tag, antenna and reader;
- Robustness of the tag and application;
- Overall system cost.

In general, for low-cost applications UHF systems are the most popular.

RFID Standards

Active and passive RFIDs each have unique standards. Passive UHF RFID is the only type regulated by a single standard. Currently, Generation-2 ("Gen-2") is the most widely accepted and used standard, covering UHF RFID systems operating in the frequency range 860-960MHz. This standard is maintained by EPCglobal, and defines the logical and physical requirements of the tags and readers.

Depending on their functionality, RFID tags are grouped in five classes, as shown below:

Class	Description
0	UHF passive read-only tag
1	UHF or HF passive write-once, read-many-times tag
2	Passive read-write tag
3	Read-write tags with on-board sensors, can be semi-passive
4	Read-write active tags, can communicate with other tags
5	As 4, but can power other tags and can be readers

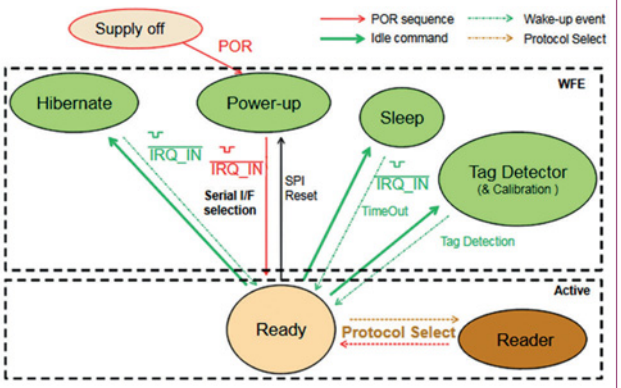

Figure 4: CR95HF operating states

RFID System Design

Figure 3 shows the block diagram of an example. The EasyPIC V7 microcontroller development board is used, with a microBUS-compatible RFID Click board, and several tags used for testing the system. The identity of the tags is read and displayed on an LCD connected to the development board. EasyPIC V7 is equipped with two microBUS-compatible sockets; the RFID Click is plugged into microBUS socket 1.

The RFID Click board is based on the CR95HF transceiver chip, and has the following features:

- 13.56MHz contactless transceiver;
- Internal on-board antenna;
- Serial UART and SPI interface;
- +3.3V operation;
- NFC forum type 1-4 tag compatible;
- ISO_IEC_14443 protocol compatible.

The tags used in the design are small, plastic, 13.56MHz NFC forum type 1-4-compatible with read and write capabilities, and on-board memory sizes from 96 bytes to 4Kbytes.

As shown in Figure 4, the CR95HF has two operating modes: Wait for Event (WFE) and Active. In the Active mode, the CR95HF communicates with a tag, or an external host.

The WFE mode includes four low-power states: power-up, hibernate, sleep and tag detector. The power-up mode is entered after power is applied to the chip. In hibernate and sleep, the chip

waits to be awakened by its interrupt pin. The tag detector state wakes up the chip whenever a tag is detected.

In this example, the SPI interface is used to communicate with the RFID Click board. The software is developed using the mikroC for PIC programming language and IDE.

Listing 1 shows operation of the software as a program description language. Listing 2 shows the complete program listing.

At the beginning of the program, the interface between the LCD and the microcontroller are defined, and so are the CR95HF commands; the RFID Click interface is also configured. The main program then initialises the SPI bus library and the LCD library. The program waits until an echo is received from the CR95HF chip, and then sets the protocol to ISO_IEC_1443. The remainder of the program is executed in an endless loop, where the program calls function GetNFCTag to check if a tag has been recognised and, if so, its ID received and displayed on the LCD. Notice that calibration of the tag detection process may be required for optimum response.

```
        BEGIN
Define LCD interface signals
Define CR95HF chip commands
Define RFID Click interface signals
Configure I/O ports as digital outputs
Initialize the SPI bus library
Initialize the LCD library
Wait until CR95HF is detected
Set the protocol to ISO_IEC_14443
DO FOREVER
        IF tag detected
        Get tag ID
        Display tag ID
        Wait 1 second
                        ENDIF
                ENDDO
        END
```

Listing 1: Operation of the program

```
// LCD module connections
sbit LCD_RS at RB4_bit;
sbit LCD_EN at RB5_bit;
sbit LCD_D4 at RB0_bit;
sbit LCD_D5 at RB1_bit;
sbit LCD_D6 at RB2_bit;
sbit LCD_D7 at RB3_bit;

sbit LCD_RS_Direction at TRISB4_bit;
sbit LCD_EN_Direction at TRISB5_bit;
sbit LCD_D4_Direction at TRISB0_
bit;
sbit LCD_D5_Direction at TRISB1_bit;
sbit LCD_D6_Direction at TRISB2_bit;
sbit LCD_D7_Direction at TRISB3_bit;

// CR95HF Commands Definition
#define ProtocolSelect 0x02
#define SendRecv 0x04
#define Idle        0x07
#define RdReg       0x08
#define ECHO        0x55

// RFID Click Connections
sbit SSI_1 at LATA.B2;
sbit SSI_0 at LATE.B1;
sbit IRQ_IN at LATC.B0;
sbit CS at LATE.B0;

sbit SSI_1_Direction at TRISA.B2;
sbit SSI_0_Direction at TRISE.B1;
sbit IRQ_IN_Direction at TRISC.B0;
sbit CS_Direction at TRISE.B0;

unsigned short sdata[18];
unsigned short rdata[18];
unsigned short res = 0, dataNum = 0;
unsigned short j = 0, tmp = 0;
char ID[10] = {0};
char txt_hex[3];

// Write command to the CR95HF
void writeCmd(unsigned short cmd,
unsigned short dataLen)
{
 unsigned short i = 0;
 CS = 0;
 SPI1_Write(0x00);
 SPI1_Write(cmd);
 SPI1_Write(dataLen);
 while (dataLen == 0)
 {
 CS = 1;
 break;
 }
 for(i=0; i<dataLen; i++)
 {
 SPI1_Write(sdata[i]);
 }
 CS = 1;
}

// Poll the CR95HF
void readCmd()
{
 unsigned short i = 0;
 while(1)
 {
 CS = 0;
 SPI1_Write(0x03);
 res = SPI1_Read(0);
 CS = 1;
 if((res & 0x08) >> 3)
 {
 CS = 0;
 SPI1_Write(0x02);
 res = SPI1_Read(0);
 dataNum = SPI1_Read(0);
 for(i=0; i<dataNum; i++)
 rdata[i] = SPI1_Read(0);
 CS = 1;
 break;
 }
 CS = 1;
 Delay_ms(10);
 }
}

// Select the RF communication
protocol (ISO/IEC 14443-A)
void Select_ISO_IEC_14443_
Protocol()
{
 sdata[0] = 0x02;
 sdata[1] = 0x00;
 writeCmd(ProtocolSelect, 2);
 readCmd();
 for(j=0; j<18; j++ )          // Clear
read/write buffers
 {
 rdata[j] = 0;
 sdata[j] = 0;
 }
}

 // Read the tag ID
char GetNFCTag()
{
 sdata[0] = 0x26;
```

```
sdata[1] = 0x07;
writeCmd(SendRecv , 2);
readCmd();
sdata[0] = 0x93;
sdata[1] = 0x20;
sdata[2] = 0x08;
writeCmd(SendRecv , 3);
readCmd();
if(res == 0x80)
{
for( j = 1; j < dataNum - 3; j++)
{
ByteToHex(rdata[j], txt_hex);
strcat(ID, txt_hex);
LATD.B0 = 1;
}
ID[10] = 0;
return 1;
}
else
{
LATD.B0 = 0;
return 0;
}
}

initialize()
{
ANSELA = 0;
ANSELB = 0;
ANSELC = 0;
ANSELE = 0;

CS_Direction = 0;
IRQ_IN_Direction = 0;
SSI_0_Direction = 0;
SSI_1_Direction = 0;

IRQ_IN = 1;
CS = 1;
// Set SPI mode
SSI_1 = 0;
SSI_0 = 1;
SPI1_Init_Advanced(_SPI_MASTER_
OSC_DIV16, _SPI_DATA_SAMPLE_
MIDDLE,
_SPI_CLK_IDLE_LOW, _SPI_
LOW_2_HIGH);
Lcd_Init();
}

// Get Echo reponse from CR95HF
char EchoResponse()
{
CS = 0;
SPI1_Write(0x00); // Send cmd to
CR95HF
SPI1_Write(ECHO);
CS = 1;
while(1)
{
CS = 0;
SPI1_Write(0x03);
tmp = SPI1_Read(1);
CS = 1;

if((tmp & 0x08) >> 3)
{
CS = 0;
SPI1_Write(0x02);
tmp = SPI1_Read(1);
CS = 1;
if(tmp == ECHO)
{
return 1;
}
return 0;
}
}
}

void main()
{
Initialize();
while (!EchoResponse()) // Until
CR95HF is detected
{
IRQ_IN = 0;
Delay_ms(10);
IRQ_IN = 1;
Delay_ms(10);
}
Select_ISO_IEC_14443_Protocol();
Lcd_Cmd(_LCD_CLEAR);              //
Ready for Tag scanning
Lcd_Out(1, 1, "Tag ID:");
while(1){
if(GetNFCTag())
{                                 //
Get tag ID
Lcd_Out(2,1, &ID);
for(j=0; j<10; j++)
// Clear ID buffer
{
ID[j] = 0;
}
Delay_Ms(1000);
}
}
}
```

Listing 2: Program of the example design

In this example, the tag returned its 32-bit ID and the LCD display showed:

Tag ID:
A1B05D7A ●

Electrical Review

The invaluable resource for electrical professionals informing the industry for over 140 years

- ✔ Lighting
- ✔ Cable management
- ✔ Data centre management
- ✔ Voltage optimisation
- ✔ Smart grids
- ✔ Renewable energy
- ✔ Transformers
- ✔ Safety
- ✔ Training
- ✔ Test and measurement
- ✔ Drives and controls
- ✔ Lightning protection
- ✔ UPS/standby power and batteries
- ✔ Plus product news, a supplier directory and much more

Register now for your free subscription to the print and digital magazines, and our weekly enewsletter

Subscribe for free today

www.electricalreview.co.uk/register

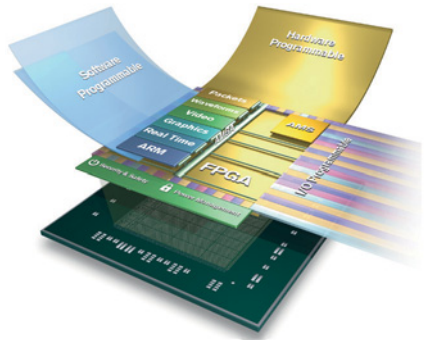

Efficient sensor fusion in embedded vision systems

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

This series of articles has considered several aspects of developing embedded vision systems, including sensor selection, interfacing and development of the signal chain comprising vision-signal processing algorithms. In a sophisticated embedded vision application, such as an automotive Advanced Driver Assistance System (ADAS), some functionality may be dependent on combining the results from two or more sensors. This is sensor fusion, and enables the system to acquire information that can't be provided by one sensor alone.

In the context of vision-based systems, sensor fusion is usually done in real time, to enable immediate decision making. The alternative is offline sensor fusion, where sensor data is extracted and fused, and decisions made at a later point in time. Moreover, in an application such as ADAS, sensor fusion may involve combining several channels of data from sensors of the same type, or could demand fusion of data from different type sensors. An object-detection and distance-monitoring application (Figure 1) provides a good example for comparison of these homogeneous and heterogeneous approaches to sensor fusion.

A system relying on a single forward-looking vision sensor could detect and identify objects, but at least one more vision sensor is needed to calculate distances to detected objects using a parallax algorithm.

Alternatively, combined object-detection, recognition and range-finding can be enabled by fusing vision-sensor data with radar or lidar. Other application examples involving fusion of differing images include x-ray, MRI and CT for medical applications, or visible and infrared images in security systems.

Processing Demands

Crunching data from multiple vision sensors requires considerable processing power. If using colour image sensors, pre-fusion processing such as colour-filter interpolation, colour-space conversion, resampling and image correction are required. The sensor-fusion algorithm itself must be performed, and an ADAS system requires subsequent background subtraction, thresholding and contour detection to locate objects using the simplest approach, whereas some systems may use an even more processing-intensive HoG/SVM classifier. Moreover, demands for higher frame-rate or larger image size further increase the computation required to pre-process the image and extract the information.

Of course, this is literally only half of the story: in a homogeneous system, the same image-processing pipeline must be implemented for the second sensor. Similarly, a heterogeneous system must configure, drive, receive and extract the information from the accompanying channel/s.

Benefit Of All Programmable Architectures

Within embedded vision systems it is common to use All Programmable FPGAs or All Programmable SoCs to implement the image-processing pipeline. If these make sense for traditional embedded vision applications, then they really stand out for embedded vision fusion applications.

An embedded vision application typically uses a processor for supervision, control and communication. In an All Programmable SoC, this is a hard core with many supporting peripherals and interface standards. If an All Programmable FPGA is used, the processor can be a soft core with customised peripheral and interface support. Taking advantage of other features of these embedded processors, such as SPI or I2C interfaces, allows additional sensors such as accelerometers, magnetometers, gyroscopes and even GPS to be connected. This enables the software to quickly and easily obtain required information from a host of different sensor types, and provides for a scalable architecture.

While the image-processing pipeline required to extract information from the image sensor can be implemented easily in programmable logic fabric, this fabric can also implement pipeline for other heterogeneous sensors such as radar and lidar, or multiple pipeline instances in the case of a homogeneous system.

The tight coupling between the processor memory and programmable logic in All Programmable Zynq-7000 or All Programmable Zynq UltraScale+ MPSoCs ensures the application

software can easily access the resultant datasets for further processing and decision making. Because the separate sensor chains are implemented in logic they operate in parallel, which is beneficial when synchronisation is required, such as with stereoscopic vision. Moreover, implementation can be accelerated by using high-level synthesis (HLS) to develop the algorithms directly for implementation within the programmable logic fabric.

Example Architecture

Both homogenous and heterogeneous approaches can be demonstrated in an All Programmable SoC. While the sensor types will will vary between both applications, the end objective of both architectures is to place two data sets within the processing system DDR memory while maximising the performance provided by the programmable logic fabric.

Considering the homogeneous approach first, the resulting implementation is a stereoscopic vision system in which each channel uses a CMOS imaging sensor. A major advantage is that only one image-processing chain needs to be developed; the same design can be instantiated twice within the programmable logic fabric for both image sensors. This enables a significant saving in development costs, even though the algorithms for calculating parallax require intensive processing.

One of the most important requirements in such a system is to synchronise the two image-processing chains. When implementing the chains in parallel within the programmable logic fabric, this requirement can be met by applying the same clock to each chain, subject to appropriate constraints.

The architecture of the homogenous approach (Figure 2) shows the two image-processing chains, which are based predominantly upon available IP blocks. Image data is captured using a bespoke sensor-interface IP module and converted from parallel format into streaming AXI. This allows for an easily extensible image-processing chain, the results from which can be transferred into the PS DDR using the high-performance AXI interconnect combined with video DMA.

Figure 1: Vision-based ADAS requiring sensor fusion

If a heterogeneous implementation using differing sensor types is considered, this could combine the image-sensor object-detection architecture described earlier with radar to perform distance detection. There are two options for implementing the radar: pulsed approach (Doppler) or continuous wave. The best option will depend upon the requirements for the final application; however, both follow a similar approach.

The radar implementation (Figure 3) can be considered in two parts: signal generation including a high-speed digital-to-analogue converter to produce a continuous-wave or pulsed signal, and signal reception using a high-speed analogue-to-digital converter to capture the received continuous-wave or pulsed signal. When it comes to signal processing, both approaches will use FFT-based

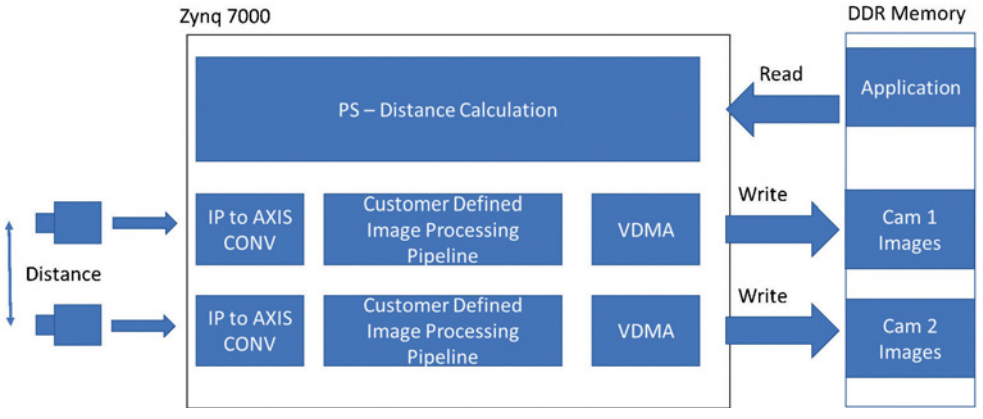

Figure 2: Homogeneous sensor fusion in All Programmable FPGA

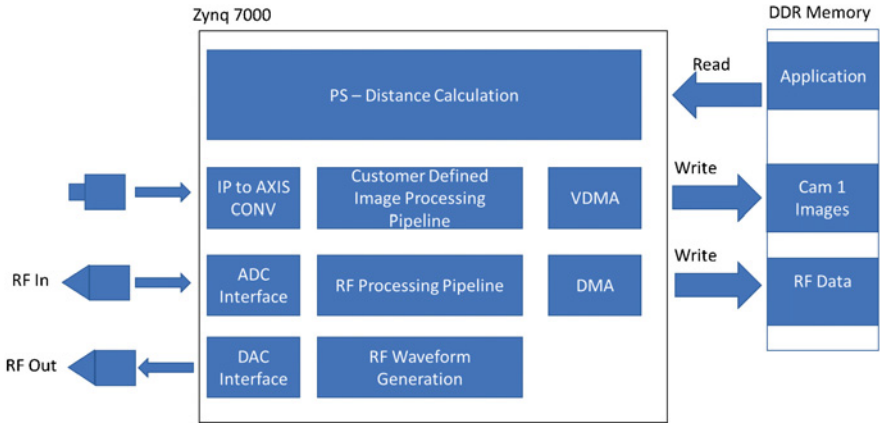

Figure 3: Heterogeneous sensor fusion

analysis implemented with the programmable logic fabric; the resultant data sets can be transferred to the PS DDR using DMA.

For either implementation, the fusion algorithm for both datasets is performed with the PS, using software. It is worth noting that designers often find that fusion algorithms can impose intensive demand for processing bandwidth. One option available to create higher performance is to use the SDSoC design environment. This enables software functions to be transferred seamlessly between the processor and programmable logic of a SoC, using Vivado HLS and a connectivity framework. Both are transparent to the software developer. The use of HLS to develop the processing chains of both homogenous and heterogeneous implementations can be extended further to create a custom SDSoC platform for the chosen implementation and then use SDSoC to harness uncommitted logic resources to further accelerate the performance of the overall embedded vision system. ●

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

A five-step guide to phase noise

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

In this series, I will present an in-depth guide to phase noise, from its basic concepts to how to improve it in applications.

Understanding Phase Noise

Phase noise is the frequency-domain representation of rapid, short-term, random fluctuations in the phase of a waveform, caused by time-domain instabilities, or jitter. Jitter is a measure of the stability of an oscillator in the time domain; it combines all the noise sources and shows their effects in time.

For simplicity, the concept of phase noise can be broken down into five steps; see Figure 1.

Step 1: Spectral Density

This is the measure of the signal's power intensity in the frequency domain – a useful way to characterise a random signal's amplitude versus frequency. For example, in Figure 2 each red dot is the spectral density at a selected frequency point.

Step 2: Signal Power Density

When plotting each spectral density point at varying frequency intervals of your choosing – in case of Figure 2 every 1Hz – it's known as signal power density of the noise.

Step 3: Noise Power Density

The upper sideband of the graph from "f_{start}" to "f_{stop}" is known as the "single sideband"; see Figure 3.

The plotted part of the single sideband (SSB) is noise; anything above the nominal oscillator frequency "F_{osc}" and not harmonically related is considered phase noise. The technical term for this part is noise power density, measured in dBW (log(Watts)).

Step 4: SSB Noise Density

When we combine the single sideband and noise power density, we are measuring what's called SSB (single sideband) noise density.

Step 5: Phase Noise

Figure 4 shows a "jittery" waveform – that's jitter.

Because the jitter is much smaller than one complete period (see Figure 4), we can say it's caused by phase fluctuations, instead of frequency fluctuations. Since these fluctuations are noise, that's phase noise.

So, "SSB noise density" is, in fact, "phase noise".

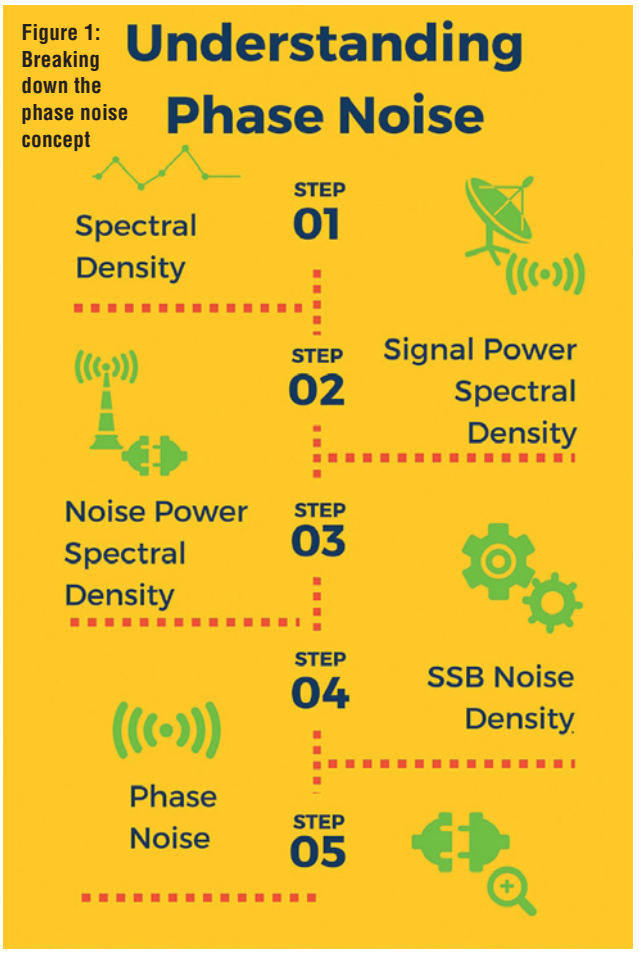

Figure 1: Breaking down the phase noise concept

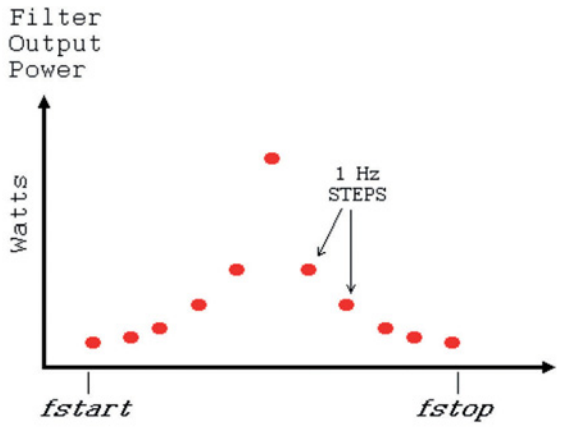

Figure 2: Noise signal power density

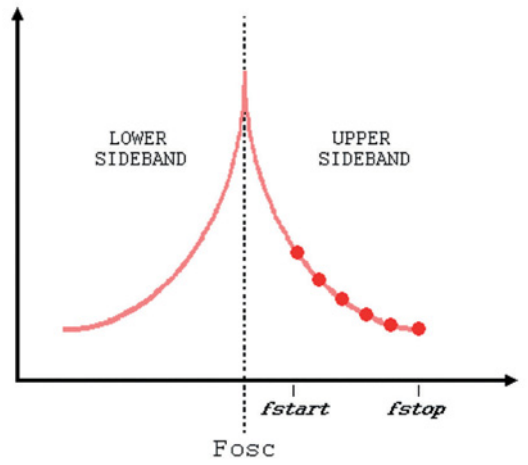

Figure 3: Single sideband

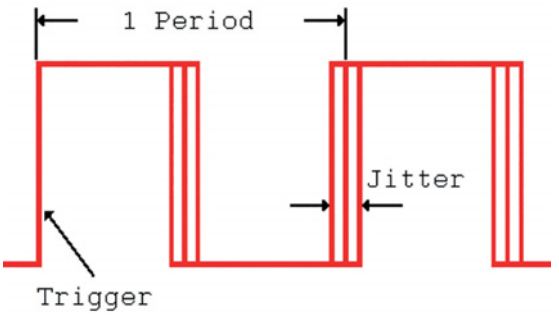

Figure 4: Jitter

Figure 5: Spectrum analyser

Common Phase Noise Types

Some common types of phase noise in crystal oscillators are caused by vibration (jitter); here are some commonly-encountered types:

- Thermal (Johnson) noise;
- Shot noise;
- Flicker noise (pink noise);
- Crystal defects (ageing).

Tools For Measuring Phase Noise

To find and ultimately prevent phase noise from negatively-effecting crystal oscillators or other components, it first needs to be measured. The most common technique is with a spectrum analyser; see Figure 5.

Spectrum analysers measure the power of radio frequency (RF) signals in a specific bandwidth, so they can also be used to measure phase noise levels. Some newer spectrum analysers even have specific features like peak find, peak tracking and normalisation to 1Hz bandwidth, making it much easier to track and measure phase noise.

Other common methods and tools include:

- Delay line discriminators, which demodulate the signal to provide an output that represents the FM noise of a given source;
- Quadrature technique, which uses two oscillators at equivalent frequencies to measure phase noise;
- FM discriminators, which use mixer and a reference source to measure phase noise;
- Direct digital measurement, which uses high-speed digitisers to allow digital data to quickly be exported to computing hardware.

What's Next?

The next article in this series will delve deeper into measuring phase noise with a crystal oscillator. ●

SOLVING FIXED-FREQUENCY SPUR PROBLEMS IN HIGH-PRECISION ADC SIGNAL CHAINS

BY **STEVEN XIE**, PRODUCT APPLICATIONS ENGINEER WITH THE CHINA DESIGN CENTRE OF ANALOG DEVICES (ADI)

Current high-resolution successive-approximation-register (SAR) and Σ-Δ ADCs provide high resolution and low noise, but system designers can't achieve the rated datasheet signal-to-noise ratio (SNR) performance. It's even harder to achieve the optimum spurious-free dynamic range (SFDR) that is a clear noise floor without spurs in the system signal chain. Spurs are typically introduced by the circuit around the ADC or external interference.

This article will introduce approaches for determining the root causes of spur issues in high-resolution, precision ADC applications and present solutions to solve them. These techniques and methods will help improve end-system EMC capability and reliability.

We will also discuss five different application cases of specific design solutions for reducing spurs caused by:

1. DC-to-DC power supply radiation from the controller board;
2. AC-to-DC adaptor noise through external reference;
3. The analogue input cable;
4. Interference coupled on the analogue input cable;
5. Room lighting.

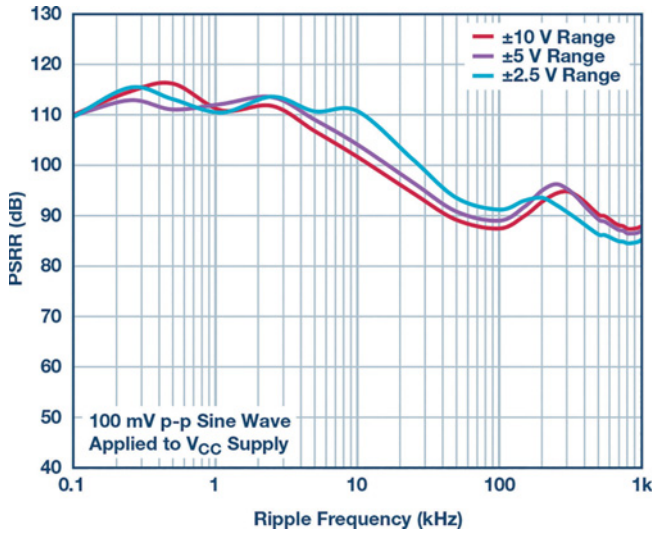

Figure 1: AD7616 PSRR vs. ripple frequency

Spurs And SFDR

SFDR represents the smallest power signal that can be distinguished from a large interfering signal. For today's high-resolution, precision ADCs, the SFDR is typically dominated by the dynamic range between a fundamental frequency and the second or third harmonic of the fundamental frequency of interest. However, there are spurs that can occur and limit the performance due to other aspects of the system.

The spurs can be categorised as input-frequency-dependent and fixed-frequency, where input-frequency-dependent spurs are related to harmonics or non-linear performance. Here we will focus on fixed-frequency spurs caused by power supplies, external references, digital interfacing, external interference, and more. Based on the application, these types of spurs can either be reduced or eliminated to help achieve maximum signal chain performance.

DC-To-DC Power Supply Noise Spurs

Typically, low dropout regulators (LDOs) are the suggested solution for generating low-noise power supply rails for precision ADCs in precision measurement systems, because of the DC-to-DC switching regulator's higher ripple noise. Fixed frequency or pulse-width-modulated switching regulators provide switching ripple usually at a fixed frequency, ranging from tens of kHz to a couple MHz. Noise at the fixed frequency can feed into the ADC conversion codes via the ADC's power supply rejection ratio (PSRR) mechanism.

Some designers may use DC-to-DC switching regulators for precision ADC applications due to limited budget or board space, but to achieve signal chain performance they'll need to limit the ripple noise or use ADCs with high PSRR to make sure that the ripple noise is below the ADC noise floor. Otherwise, there could be spurs at the switching frequency in the ADC's output spectrum, which may degrade signal chain dynamic range.

The AD7616 is a 16-bit digital analogue simulator (DAS) that supports dual, simultaneous sampling of 16 channels for power line monitoring. It has a very high PSRR and does a good job of rejecting/attenuating switching ripple. For example, when a DC-to-DC switching power supply with 100mV p-p ripple noise at 100kHz powers the AD7616 (VCC=5V) running at its ±10V input range, the digital code noise is:

$$\frac{100\ \text{mV p–p}}{10^{\frac{88\ \text{dB}}{20}}} = 3.98\ \mu\text{V}\ \textit{or}\ 0.013\ \text{LSB}$$

This level of ripple in the ADC output is extremely low for a 16-bit converter. So high PSRR performance in ADCs makes it possible to use switching regulators in precision measurement systems.

DC-to-DC Power Supply Radiation Spurs

Using a high PSRR ADC doesn't ensure that switching regulators won't cause problems in precision measurement systems, since the ripple noise from switching regulators can feed into the ADC's digital codes through other ways.

The Analog Devices AD4003 is a low-noise, low-power, high-speed 18-bit, 2 MSPS precision SAR ADC. During EVAL-AD4003FMCZ evaluation board AC performance testing, a spur at a level of about -115dBFS was found at around 277.5kHz; the spur and its second harmonic are shown in Figure 2.

First, it was confirmed that the AD4003's power supplies were not causing the spurs. Then, tests were performed to determine if the spurs were coming from the analogue input:

1. The spurs decreased when the differential analogue input conditioning circuitry was removed.
2. The spurs also decreased when a narrow-bandwidth RC filter (such as 1kΩ/10nF) was inserted at the front end of the AD4003's buffer amplifier, ADA4807-1.

These results show that noise causing the spurs may pass through the conditioning circuitry and into the AD4003's analogue inputs.

Next, the sensor output was disconnected and the conditioning circuitry removed, leaving only the $V_{REF}/2$ CM voltage at the non-inverting input of the ADA4807-1. However, the spurs remained almost unchanged.

It was then suspected that the interference source was somewhere in the EVAL-AD4003FMCZ signal chain. To prove this, a copper-foil shield was placed at various locations on the EVAL-AD4003FMCZ and controller SDP-H1 boards. When the copper-foil shield was placed over the DC-to-DC power supplies on the SDP-H1 board, as shown in Figure 3, the spurs disappeared. The spur frequency of 277.5kHz matches the programmed switching frequency of the ADP2323 regulator. Figure 4 shows the 3.3V VADJ_FMC switching frequency power as captured by the EVAL-AD7616SDZ GUI FFT.

We concluded that the DC-to-DC switching frequency interference was being emitted by the 8.2µH inductor, L5. The

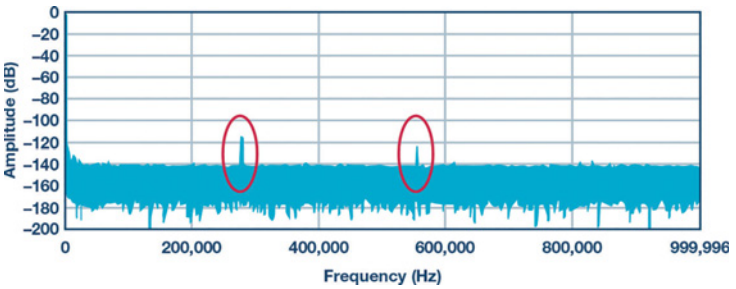

Figure 2: Spurs as seen on an EVAL-AD4003FMCZ eval board

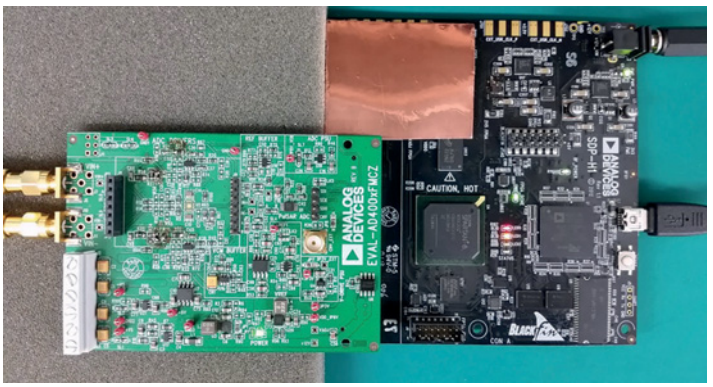

Figure 3: VADJ_FMC inductor L5 covered by copper foil shield

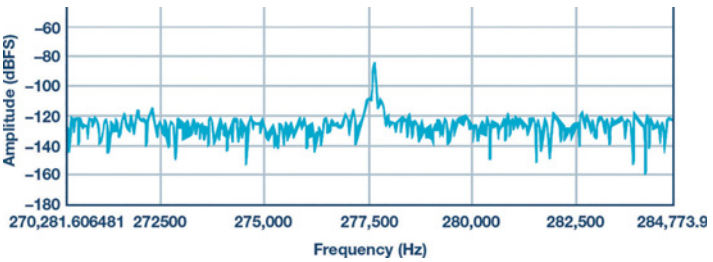

Figure 4: VADJ_FMC 3.3 V switching ripple captured by EVAL-AD7616SDZ GUI FFT

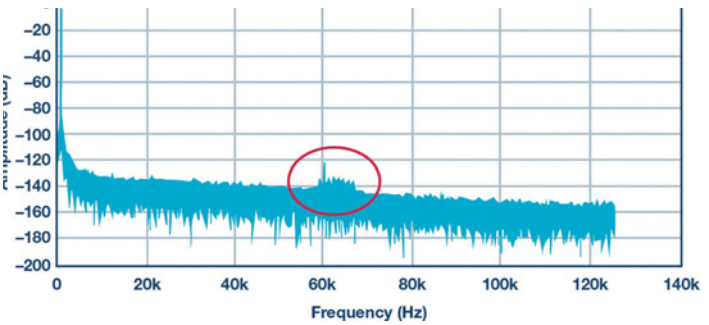

Figure 5: Spurs seen on the EVAL-AD7175-2SDZ board

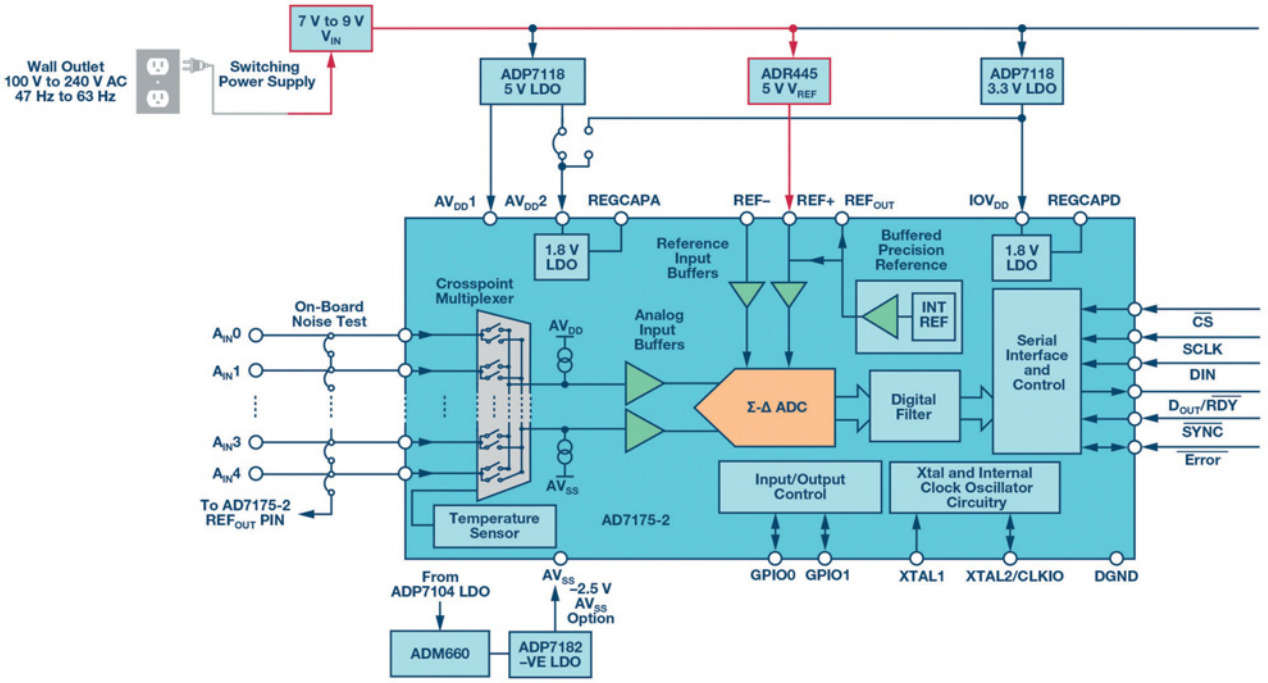

Figure 6: The EVAL-AD7175-2SDZ board connections

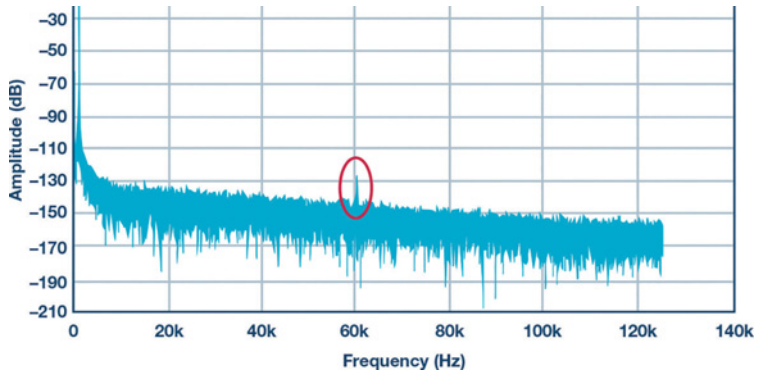

Figure 7: The cluster of spurs removed on EVAL-AD7175-2SDZ board

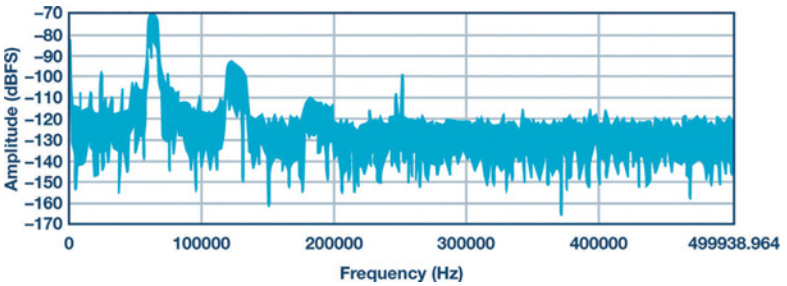

Figure 8: 3.3 V VADJ_FMC switching ripple captured by EVAL-AD7616SDZ GUI FFT

interference was injected into the signal chain at the input of buffer amplifier ADA4807-1, and then went into the AD4003 ADC's analogue input.

Potential solutions to this spur issue caused by the DC-to-DC power converter include:

1. Using a low-pass filter at the front end of the AD4003 ADC to attenuate the coupled DC-to-DC switching frequency interference to a level that meets the design target (that is, the spur buried in noise), if the application bandwidth allows it.
2. Using the new SDP-H1 board (BOM Rev 1.4), which adds shielding to L5. Radiated interference power is reduced, so the spurs seen in AD4003 ADC's spectrum are much lower.
3. The VADJ_FMC voltage level can be programmed by the EEPROM on the EVAL-AD4003FMCZ board. It was found that using a lower voltage level, such as 2.5V for VADJ_FMC, also caused the spurs to disappear.

Noise Coupling Through External Reference Spurs

ADCs quantise an analogue signal into digital code referred to the ADCs' DC reference voltage level. Therefore, the noise on the DC reference input will directly feed into the ADC's output digital codes.

The AD7175-2 is a low-noise, fast-settling, multiplexed, 2-/4-channel (fully/pseudo differential) Σ-Δ ADC for low bandwidth inputs. During the EVAL-AD7175SDZ signal chain test, a cluster of spurs around 60kHz was captured, as shown in Figure 5.

The AD7175-2 ADC's power supplies and analogue

conditioning circuits were evaluated and found to be good. However, as shown in Figure 6, the AD7175-2's 5V reference input is generated by the ADR445 reference, supplied by a 9V DC from an AC-to-DC adaptor external to the evaluation board. Next, a bench 9V DC power module was substituted for the adaptor, and the cluster of spurs disappeared, leaving only a narrow spur at 60kHz; see Figure 7.

The 9V output AC-to-DC adaptor was tested with the EVAL-AD7616SDZ GUI FFT while supplying the EVAL-AD7175-2SDZ board with 320mA output. The switching frequency power at the ADR445 reference's power pin is about -70dBFS with an AD7616 ±10V input range, which means 6.325mV p-p or -64dBFS at ±5V input.

$$20\ \text{V} \times 10^{\frac{-70\ \text{dB}}{20}} = 6.325\ \text{mV p-p}$$

$$20 \times \log\left(\frac{6.325\ \text{mV p-p}}{10\ \text{V}}\right) = -64\ \text{dBFS}$$

This power-switching ripple noise feeds into the AD7175-2 ADC and shows up in the digital codes with some attenuation as explained below:

1. The ADR445 reference's datasheet specifies a PSRR of 49dB at 60kHz.
2. The ADR445 output impedance is about 4.2Ω at 60kHz. This combines with the 4.8µF reservoir caps, giving a further 18dB of attenuation.
3. In addition, the AD7175-2 ADC's digital filter sinc5 + sinc1 adds about 3dB of attenuation at 60kHz, when ODR is 256kSPS.

$$-64\ \text{dBFS} - 49\ \text{dB} - 18\ \text{dB} - 3\ \text{dB} = -134\ \text{dBFS}$$

This calculated -134dBFS level is very close to the level of the captured -130dBFS cluster of spurs (not including the highest narrow spur) shown in Figure 5. This confirms that the cluster of spurs is caused by the AC-to-DC adaptor's switching ripple feeding through the external reference ADR445. The remaining narrow spur will be resolved in the next section.

Signal Chain Interference Spurs

In the hardware system, there is generally a long signal chain from the input sensor to the input of the precision converters. This signal chain includes connecting cables, connectors, routing wires, scaling and conditioning circuits, ADC drivers, and more. There is high likelihood for external interference to enter the analogue signal chain and cause ADC spurs.

Power Cable Interference Related Spurs

During the investigation of the remaining narrow spur on the EVAL-AD7175-2SDZ board's spectral output, we noticed that there was a digital oscilloscope elsewhere on the test bench. As shown in Figure 9, the scope's 220V AC power supply cable

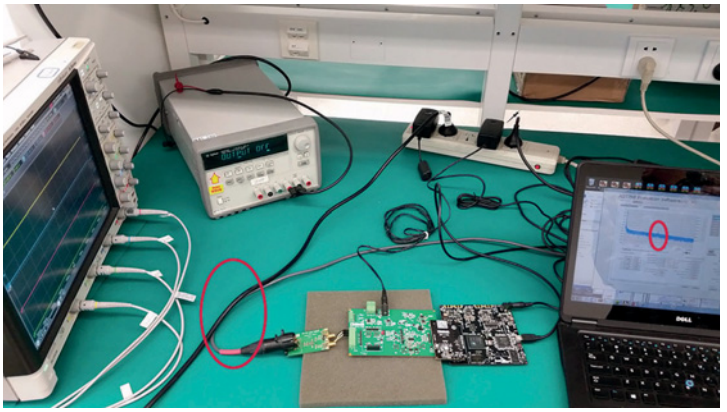

Figure 9: Spur caused by oscilloscope power supply cable

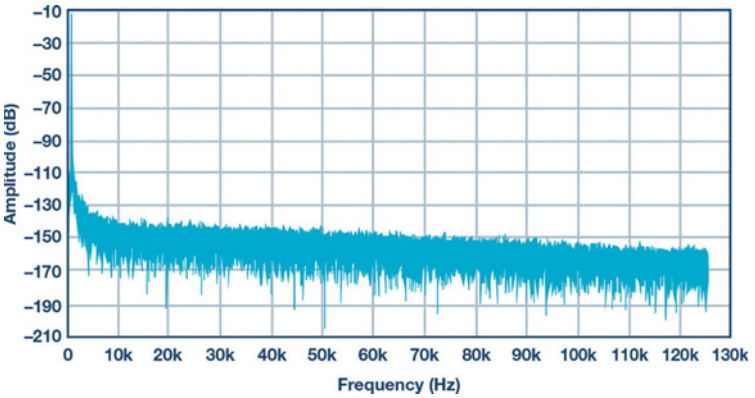

Figure 10: All spurs removed on EVAL-AD7175-2SDZ board

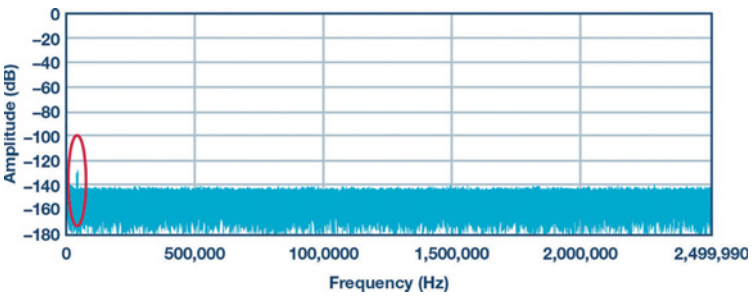

Figure 11: Spurs on EVAL-AD7960FMCZ by fluorescent lighting radiation

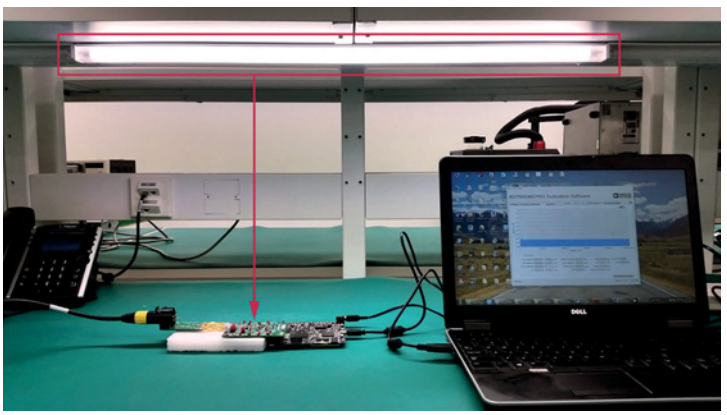

Figure 12: Fluorescent lighting in proximity to the EVAL-AD7960FMCZ board

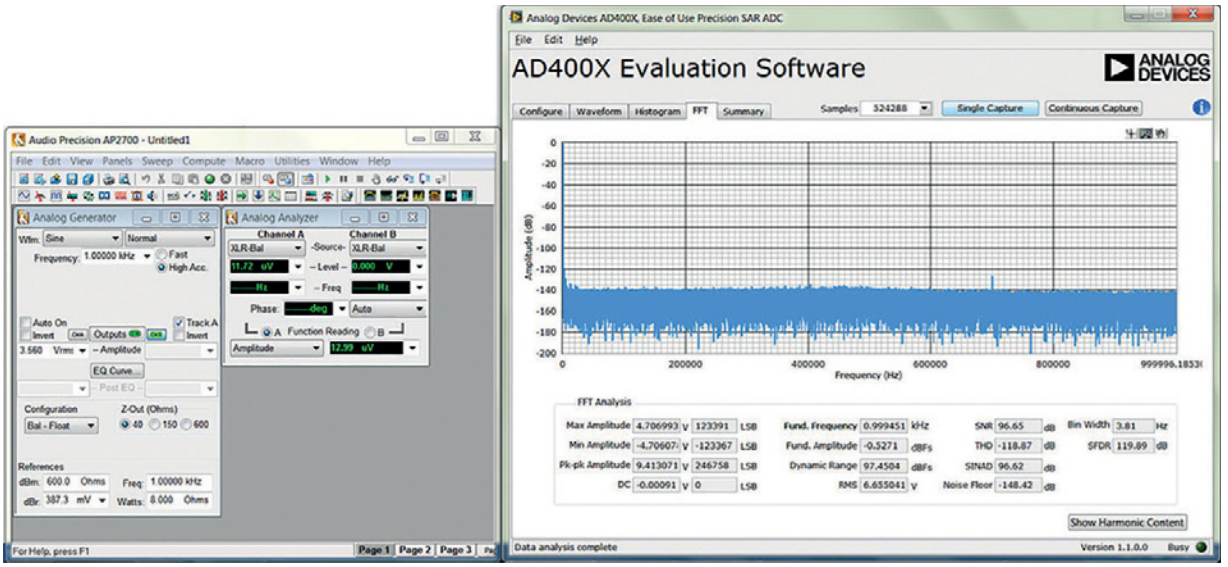

Figure 13: Spur on EVAL-AD4003FMCZ board caused by the XLR cable

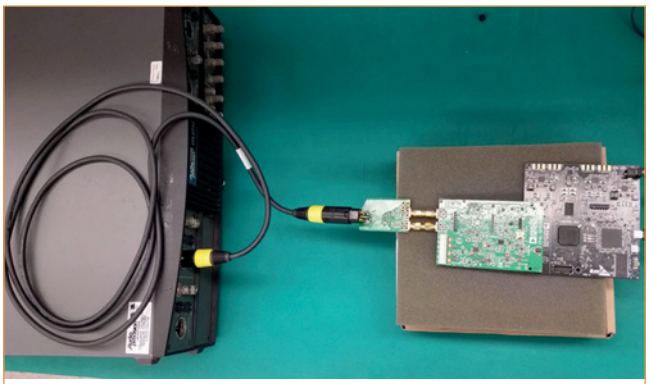

Figure 14: AP driving the EVAL-AD4003FMCZ board through a long XLR cable

(black) overlapped the EVAL-AD7175-2SDZ EVB's analogue input cable (grey). When the oscilloscope was turned off or its power cable moved away from the analogue input cable, the narrow spur at 60kHz disappeared; see Figure 10.

Care should be taken in the system cabinet, when routing the cables from the sensor to the DAQ board. It is good practice to keep the low-level sensitive analogue signals separated from the high-current power lines.

Lamp Radiation Spurs

Another spur appeared on the FFT spectrum while testing the EVAL-AD7960FMCZ board. As shown in Figure 11, its level was about -130dB at 40kHz. Its frequency seemed unrelated to any of the signal frequencies that appear on the EVAL-AD7960FMCZ or its controller board, SDP-H1.

The step to finding the source of the spur was to clear the test bench, in case something there was generating external interference. When the fluorescent light on the bench rack was turned off, the spur disappeared. Furthermore, it was found that as the EVAL-AD7960FMCZ board was moved closer to the light, the 40kHz spur grew. When an additional RC filter (such as 1kΩ/10nF) was added to the input of the buffer amplifier ADA4899-1, the spur decreased about 10dB. That meant the fluorescent light was radiating into the signal chain path at the non-inverting input of the buffer amplifier.

For systems that operate in a lighted environment, installing a shielding case over the front-end circuitry can help protect it from radiating interference, and hence optimise signal chain performance.

Long Analogue Input Cable-Related Spurs

During evaluation of the EVAL-AD4003FMCZ board, an AP SY2712 signal generator was used to send a low-noise and low-THD sine wave to the analogue inputs through an XLR microphone cable about 2m long. In this setup, a spur was apparent at about -125dB at 700kHz, as shown in Figure 13.

Through investigation, three ways to solve it were found:

- Bypass the long XLR microphone cable and short the AP balanced output XLR male connector to the interposer XLR female connector.
- Set the signal source SY2712's output impedance from Z_{Out} = 40Ω to Z_{Out} = 600Ω.
- The spur becomes smaller when a narrow-bandwidth RC filter (such as 1kΩ/10nF) is inserted in the signal chain at the front end of the AD4003's buffer amplifier, the ADA4807-1.

It was concluded that the mismatch in the signal source's output impedance and the too-long XLR cable caused the high frequency spur at 700kHz. ●

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

RF ENERGY HARVESTING FOR MOBILE DEVICES

BY **MAURIZIO DI PAOLO EMILIO**, TELECOMMUNICATIONS ENGINEER AND ELECTRONICS DESIGNER

Energy-harvesting systems capture energy from the environment and convert it into electricity. Energy sources are everywhere, from light, sound, vibration, movement, and even heat dissipation.

One interesting energy source is the transmitted radio frequencies (RF). However, the challenge here is recovering it from power transmitted in telecommunications, broadcasting or mobile communications. Collecting this type of energy would charge and power small, portable devices without using batteries. Self-powered devices can be wearable or part of the Internet of Things (IoT).

RF Energy Devices

A typical RF energy-harvesting system consists of a receiving antenna, rectifier circuit and a power management system that includes a DC-DC converter, with a system for storing that energy, such as a battery or capacitor.

There's a distinction between non-radiative and radiative RF, with the first type being based on inductive couplings, whereas the latter uses the transmission and reception of radio waves (electromagnetic waves).

All RF energy-harvesting systems need a suitable transducer or antenna, one of the most important parts of the RF energy harvesting system, which of course must be designed for maximum energy capture. To help with the design, certain equations are used, such as Maxwell's, and the Friis transmission equation, which helps determine the transmission and reception of RF signals in free space.

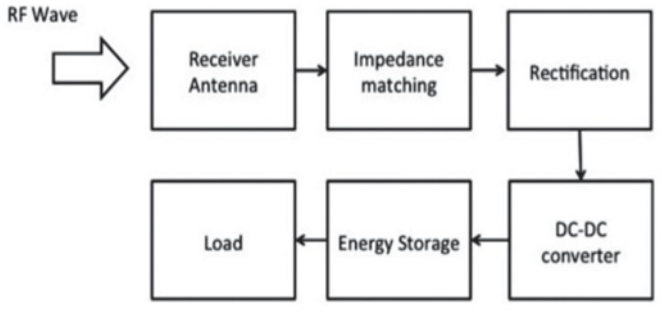

Figure 1: Generic layout of an RF energy-harvesting system

In general, a matching circuit is applied between antenna and rectifier, and a DC-DC conversion circuit is connected to the energy storage system, which is then connected to the load; see Figure 1.

Antenna Design

When an electric field hits the antenna with a certain intensity, the induced voltage it generates must then be suitably rectified. Figure 2 shows a ferrite core with many turns, a suitable solution for low frequencies and, in particular, for medium-wave transmissions.

Figure 3 shows the antenna's equivalent circuit, where R_R is the radiation resistance of the loop proportional to the loop section and the ferrite's magnetic permeability, Q is the quality factor and R_f is the loss resistance of the rod.

The circular spiral antenna in Figure 4 with a microstrip line ensures a high power-density and is one of the most used for recovering energy from radio waves.

There are various configurations for microstrip antennas (L, E and U shape), but most are narrow-band. Figure 5 shows the equivalent circuit of a microstrip antenna, composed of two inductors and a shunt capacitance.

Power Management

The rectifier circuits provide a DC output voltage to the corresponding load, which in the case of an RF energy harvesting system is the input impedance of the DC-DC converter.

There are three main configurations for the rectifier: single-diode, diode-bridge and voltage-multiplier. The first two lead

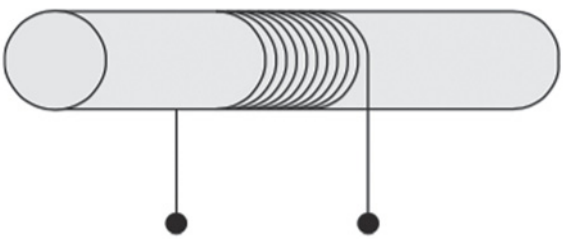

Figure 2: Ferrite rod antenna

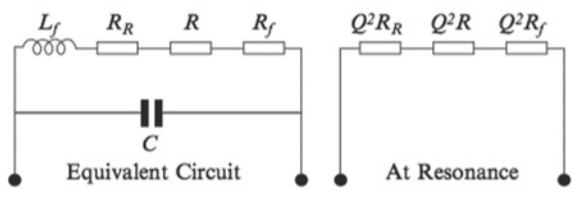

Figure 3: Equivalent circuit of a ferrite rod antenna

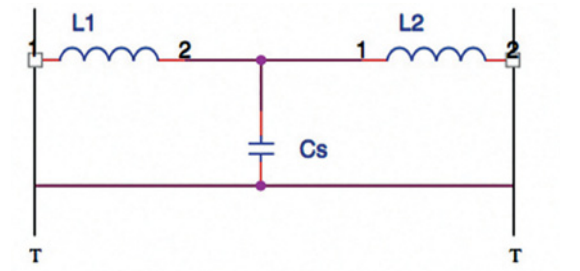

Figure 5: Equivalent circuit of a microstrip antenna

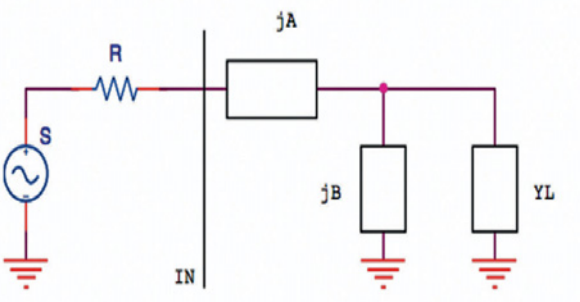

Figure 7: Antenna equivalent circuit and L-matching network

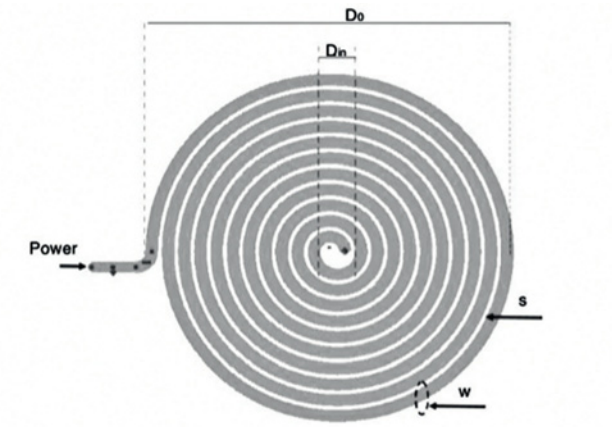

Figure 4: Circular spiral inductor antenna with main parameters that describe its geometry

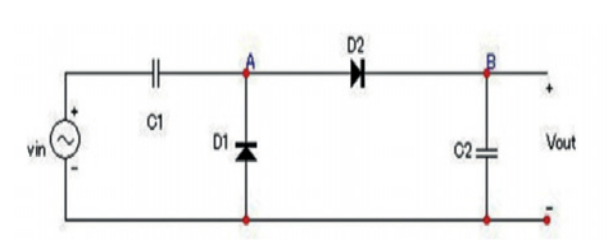

Figure 6: Voltage rectifier

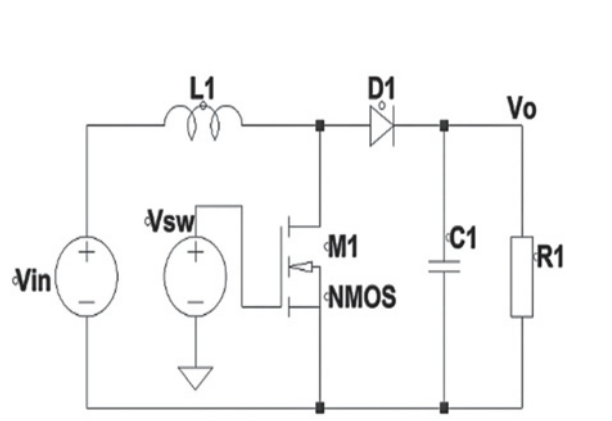

Figure 8: Boost converter

to a decrease in output voltage, whereas the multiplier-voltage rectifier allows increases in the signal's peak amplitude; see Figure 6.

In general, an impedance-matching circuit is applied between antenna and rectifier. The input voltage of the rectifier/multiplier is then dictated by the available power from the antenna and by the matching circuit's Q factor (Figure 7).

The voltage values supplied by the antenna are low, so a DC-DC converter needs to have a useful voltage level (e.g., 2-5V) to power a common electronic circuit or node sensor in the WSN. In this case, the circuit is a boost converter; see Figure 8.

Super-Capacitor

Super-capacitors (SC) accumulate energy in two capacitors in a dual-layer configuration, or ELL (electrochemical double layer). The simplest super-capacitor consists of two polarisable electrodes, a separator and an electrolyte. The electric field is stored in the interfaces between the electrolyte and electrodes.

To maximise the energy efficiency of the load, the residual energy is reduced by changing to a super-capacitor series configuration, increasing the output voltage above the minimum DC-DC converter input voltage.

ENABLING AT-HOME MEDICAL CARE WITH NEXT-GENERATION POWERING DEVICES

BY **TONY ARMSTRONG**, DIRECTOR OF PRODUCT MARKETING FOR POWER PRODUCTS AT LINEAR TECHNOLOGY

The global medical electronics market was valued around $3bn in 2015, growing at an annual compounded growth rate of 5.4%. It is expected to reach $4.4bn in 2022, according to market research site Marketsandmarkets.com. Some key factors driving this growth are ageing population, the demand for personalised, easy-to-use, advanced healthcare devices, and the adoption of wearable medical electronic devices. At the same time, the costs associated with keeping a patient in hospital are becoming unsustainable, spurring hospitals to find ways of reducing these costs without compromising patient care.

One way of reaching this goal is to release the patient, equipped with remote monitoring and diagnostic devices, setting a trend for portable and wireless medical instrumentation in outpatient care. This typically includes monitors for heart rate, blood pressure, breathing rate, sleep apnea, blood glucose levels and body temperature. Most portable electronic monitoring systems must incorporate RF transmitters so data gathered from the patient can be readily sent to a hospital or clinic for analysis.

Power is another important parameter in medical products, which must operate properly and switch seamlessly between different power sources, such as AC mains outlets, battery backups and even harvested ambient-energy sources. Equally, great lengths must be taken to protect against – and tolerate – various fault conditions, maximise operating time when devices are battery-powered and to ensure reliable system operation, with 99.999% integrity of wireless data transmission.

Patient-Monitoring Systems

So, for medical equipment to reside in a patient's home, in addition to being compact and efficient, it also needs to be robust and flexible.

Since many applications in medical electronic systems require continuous power, even when the mains supply is interrupted, a key requirement is low quiescent current to extend battery life. This usually calls for switching regulators with standby quiescent current of below 9µA, or even nano-amps, as may be the case with systems that use a combination of battery and energy-harvested power. Switching regulators tend to generate more noise than linear regulators, but they offer far greater efficiency.

Noise and EMI levels have proven manageable in many sensitive applications, if the switcher behaves predictably. EMI is minimised if a switching regulator switches at a constant frequency in normal mode, and the switching edges are clean and predictable, with no overshoot or high-frequency ringing. EMI radiation is further minimised with small package sizes and high operating frequencies, allowing for a tight layout. Furthermore, if the regulator is used with low ESR ceramic capacitors, input and output voltage ripple, which are additional sources of noise in the system, can be minimised.

The number of power rails in today's feature-rich patient monitoring devices has increased whilst operating voltages have continued to reduce. Nevertheless, many of these systems still require 3V, 3.3V or 3.6V rails for powering sensors, memory, microcontroller cores, I/O and logic circuitry. Furthermore, since their operation is sometimes critical, many of them have a battery backup system, should the main power supply fail.

Traditionally, their voltage rails have been supplied by step-down switching regulators or low-dropout regulators. However, these types of ICs do not capitalise on the battery cell's full operating range, thereby shortening potential battery runtime. Therefore, when a buck-boost converter is used (which can step voltages up or down), the battery's full operating range can be harnessed. This increases the operating margin and extends device runtime as more of the battery's life is usable, especially as it nears the lower end of its discharge profile.

Energy Harvesting As A Power Source

Recently, there has been a great deal of innovation in energy harvesting, especially using body heat as an energy source to power electronic monitoring systems or recharge a battery that powers them. Such advances enable modification of the size and shape of medical electronics components to accommodate a milliwatt or microwatt power range. This means that many complex electronic systems and devices, such as wearable medical and autonomous systems, can consume power less than 250µW.

Furthermore, wireless sensor networks with power levels of µWs-100mWs routinely operate from battery. However, due to the intrinsic limitations of battery power, such as charge longevity and periodic recharging where applicable, the use of ambient energy

sources such as heat or vibration has risen to the task.

Linear Technology launched its first energy-harvesting IC, the LTC3108, in December 2009. This is an ultra-low-voltage DC/DC converter and power management IC, designed to collect and dispense surplus energy, creating extremely low voltages from heat sources, in temperature gradients of 1°C or more.

A more recent introduction is the LTC3107, a highly-integrated DC/DC converter designed to extend the life of a primary battery in low-power wireless systems, by harvesting and managing surplus energy from extremely-low-input voltage sources, such as thermoelectric generators (TEGs) and thermopiles.

With the LTC3107, a point-of-load energy harvester requires little space, just enough room for the LTC3107's 3mm × 3mm DFN package and a few external components. By generating an output voltage that tracks the primary battery, the LTC3107 can be seamlessly adopted to bring the cost savings of free thermal energy harvesting to new and existing battery-powered designs. Figure 1 shows the LTC3107 harvesting thermal energy to power wireless senor nodes (WSNs).

The LTC3331 is a multifunctional ambient energy harvester that forms a complete regulating energy-harvesting solution that delivers up to 50mA of continuous output current to extend battery life when harvestable energy is available; see Figure 2. It requires no supply current from the battery when providing regulated power to the load from harvested energy, and only 950nA when powered from the battery under no-load conditions.

LTC3331 Internals

The LTC3331's energy-harvesting power supply consists of a full-wave bridge rectifier accommodating AC or DC inputs and a high-efficiency synchronous buck converter. The IC harvests energy from piezoelectric (AC), solar (DC) or magnetic (AC) sources, and a 10mA shunt enables simple charging of a battery with harvested energy while a low-battery disconnect function protects the battery from deep discharge.

The rechargeable battery powers a synchronous buck-boost converter that operates from 1.8-5.5V and is used when harvested energy is not available to regulate the output, whether the input is larger, smaller or equal to the output.

The LTC3331 battery charger has a very important power management feature that can't be overlooked when dealing with micropower sources. It incorporates logical control of the battery charger such that it will only charge the battery when the energy-harvested supply has excess energy. Without this logical function the energy harvested source would get stuck at startup at some non-optimal operating point and not be able to power the intended application. The LTC3331 automatically transitions to the battery when the harvesting source is no longer available. This has the added benefit of allowing the battery-operated WSN to extend its operating life from 10 years to over 20 years if a suitable ambient-energy power source is available at least half of the time, and even longer if the ambient energy source is more prevalent.

Benefits

So, the smart medical wearables market can benefit greatly from such devices. The core architecture for a smart wearable device depends on the product type, but essentially consists of a microcontroller, MEMS sensor(s), wireless connectivity, battery and support electronics.

Thus, having the current wave of feature-rich energy-harvesting and/or IoT solutions that can utilise many forms of ambient energy to power health-monitoring devices can facilitate getting patients to their homes faster without compromising their recovery.

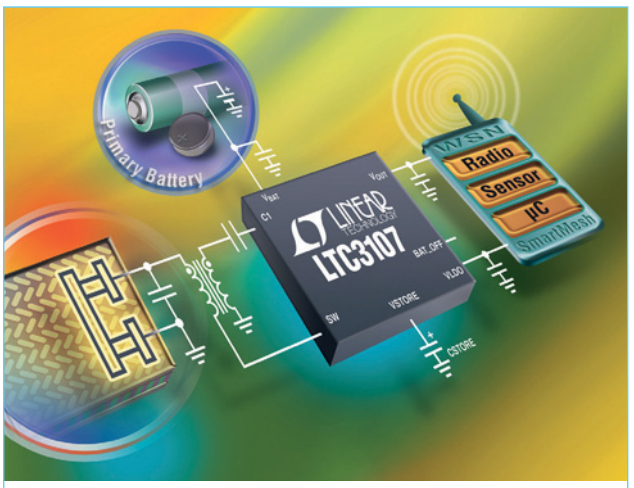

Figure 1: LTC3107 harvesting thermal energy to power a WSN and/or charge a battery

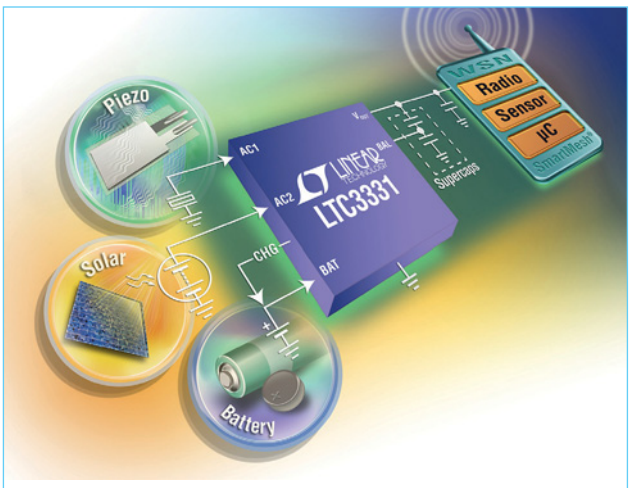

Figure 2: The LTC3331 converts multiple energy sources and can use a primary rechargeable battery

NOVEL MULTILEVEL INVERTER TOPOLOGY

LU ZHENG, **GUO ZHEN** AND **LU YU** FROM THE NANYANG INSTITUTE OF TECHNOLOGY IN CHINA
PRESENT A NEW MULTILEVEL INVERTER TOPOLOGY WITH LOW SWITCHING LOSSES

The cascaded multilevel inverter is an important topology in medium-to-high power applications. These inverters use many switches, making the drive circuits complex and switching losses large.

Multilevel inverters' advantages include high output voltages, output waveforms close to sinusoidal waves, low harmonic content and low du/dt. These characteristics make them widely useful in industries such as steel rolling, papermaking, concrete, coal, railway and shipping, also for unified power flow controllers (UPFC) and active power filters (APF), among other things.

Multilevel Inverter Types

Typically, there are three commercial topologies of multilevel inverters: diode-clamped, flying-capacitor and cascaded H-bridge. There's a capacitance voltage balancing problem in diode-clamped and flying-capacitor inverters, which doesn't exist in the cascaded H-bridge type. The latter is also used for easily-built modular systems, which are convenient to maintain, without any limitations imposed by a clamped diode or capacitor. The cascaded H-bridge inverter also outputs several voltage levels and reduces harmonic content using a relatively simple control method.

> There's a capacitance voltage balancing problem in diode-clamped and flying-capacitor inverters, which doesn't exist in the cascaded H-bridge type

Beside the conventional multilevel inverter topology structure, there are dozens of other, improved, multilevel inverter topologies such as the five-level H-bridge middle clamped structure, the modular multilevel converters, and others.

New Structure

In this study we propose a novel multilevel inverter topology, based on a cascaded structure with a half-bridge circuit and H-bridge type inverter. This approach enables cascaded-topology modular construction of systems that can easily be extended. It also boasts low harmonic content and low total

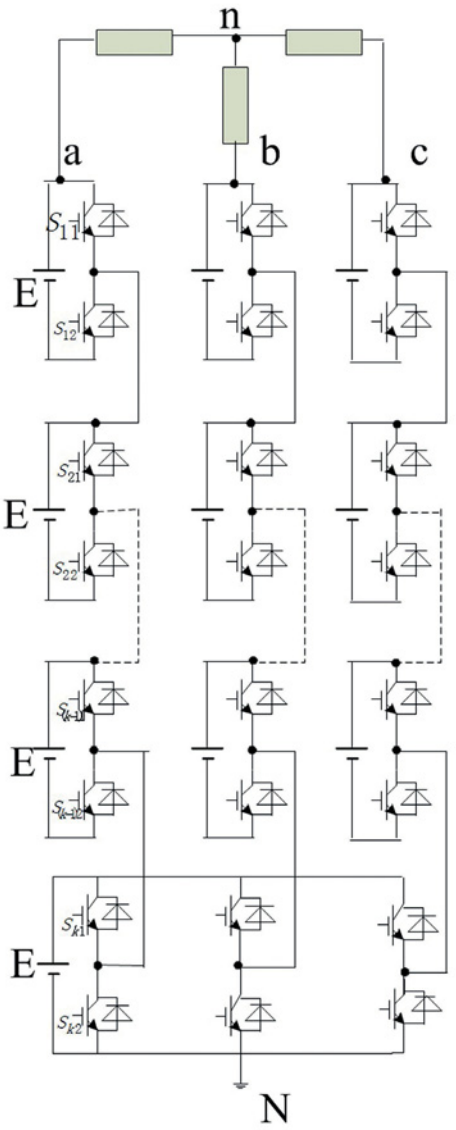

Figure 1: Proposed topology

(a) Line voltage waveform, X-axis: 1div = 10ms, Y-axis: 1div = 1000V

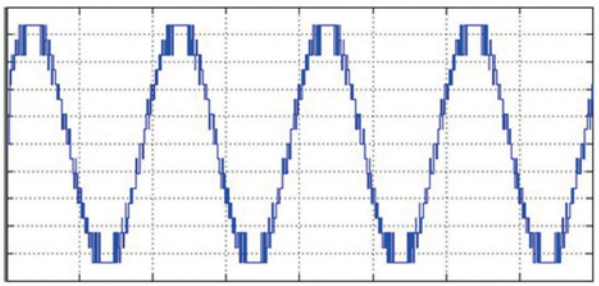

(b) Harmonic analysis

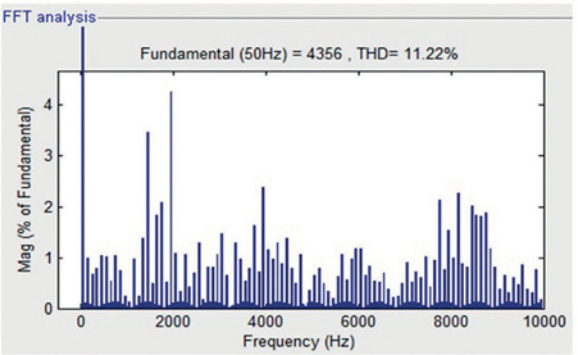

Figure 2: Line voltage waveform (U_{ac}) and its harmonic analysis

harmonic distortion (THD), as well as a lower number of switching devices, which in turn keeps switching losses low.

The traditional H-bridge cascaded topology structure offers advantages such as:

1. Independent DC sources, so no need for voltage sharing;
2. Convenient modular construction for easy design, manufacture, installation and control, which is also highly reliable;
3. Compared to diode-clamped and flying-capacitor inverters, the number of components needed by the cascaded structure is lowest;
4. It results in higher voltages and lower harmonic content.

The proposed structure is shown in Figure 1. It includes a three-phase bridge circuit and 3(k-1) half-bridge circuits. The number of switching devices is 6k, and the number of independent DC sources needed is 3(k-1)+1. There are k-1 half-bridge circuits in every phase, and a three-phase bridge inverter is cascaded in the corresponding phases. When the up-bridge switch in every phase conducts, it can output high voltage E, otherwise when the down-bridge switch conducts its output can be as low as zero.

To avoid straight-through conduction from the source, the switching devices of the up- and down-bridge can't conduct simultaneously. In this case, the outputs are k+1 level phase voltage and 2k+1 level line voltage.

Supposing the output line voltage level is 21, it needs 13 independent DC sources and 30 switching devices for the topology proposed in this study; whereas the traditional H-bridge cascaded topology structure needs 15 independent DC sources and 60 switching devices.

Simulation Results

We verified the validity of the topology through simulation, with a MATLAB model consisting of a 21-level inverter with cascaded three-phase inverter cells and half-bridge inverter cells, with each cell powered from a 540V DC voltage. The frequency of the reference voltage is 50Hz and the sampling frequency is 850Hz.

Figure 2 shows the waveform of the output line voltage U_{ac} and its harmonic analysis. The waveform is quite similar to a sinusoidal wave, and the THD is 11.22%.

Figure 3 shows the current waveform (i_a) and its harmonic analysis; the THD is only 8.28%. ●

(a) current waveform, X-axis: 1div = 10ms, Y-axis: 1 div = 10A

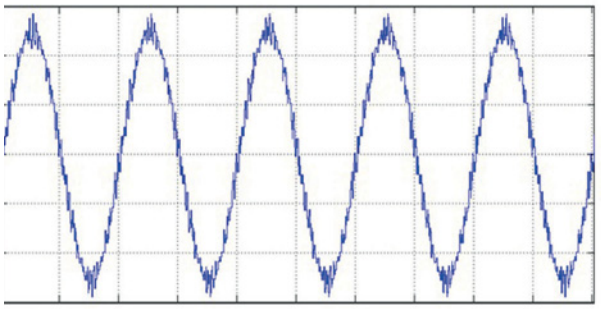

(b) Harmonic analysis

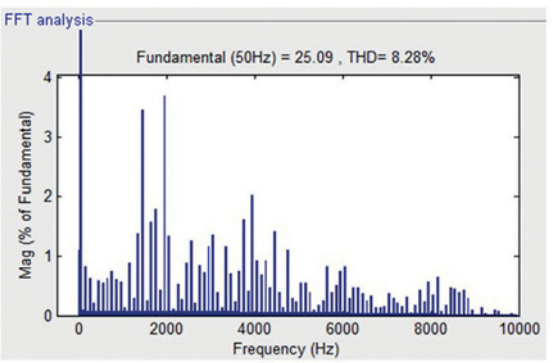

Figure 3: Current waveform (i_a) and its harmonic analysis

A GUIDE TO CHOOSING COOLING SYSTEMS FOR INDUSTRIAL APPLICATIONS

STEVE HUGHES, MANAGING DIRECTOR OF POWER QUALITY SPECIALIST REO UK, EXPLORES THE BEST WAYS OF COOLING INDUSTRIAL EQUIPMENT

Each summer, something very unexpected occurs in the UK; we glimpse a fleeting moment of sunshine. During this period, there is often a debate about the best way to cool down, from using fans to a dip in the pool. A similar debate often occurs when businesses look to keep their equipment cool, too.

Industrial equipment gets hot during use, damaging components and hindering functionality. This is particularly true of electrical equipment where, in addition to the typical operating temperatures, there is also heat generated by electrical currents. Hence, plant managers and original equipment manufacturers (OEMs) build cooling systems into their equipment.

However, cooling electrical infrastructure and industrial equipment isn't as simple and straightforward as just turning on a fan. Even something seemingly unrelated like the location of the device is vital in ensuring the system is cooled effectively and safely.

There are two main ways of regulating device temperature: air- and water-cooling. Air cooling is the most commonly used and, as the name suggests, uses fans to circulate air around the device. Water cooling, on the other hand, uses liquid coolants to transfer heat out of equipment and is often met with apprehension by plant managers. Used correctly, both have their place within industrial environments.

Air-Cooling Considerations

The humble fan, what many people initially think of when asked about cooling, is not without its limitations. For example, including a fan often means the enclosure of a device needs to be bulkier to accommodate it, which can cause problems with integrating devices into certain industrial setups.

This is especially true with electrical power controllers, since many are integrated into transformer cabinets. In these instances, the bulkier design of products with fans may cause problems and requiring the sleeker design of a water-cooled power controller.

Yet, air cooling does offer benefits; these systems are so popular because they are relatively inexpensive and typically require less maintenance beyond the replacement of filters. As long as maintenance staff are vigilant, fan-based systems will not encounter too many operational difficulties.

Fans also provide effective heat transfer in remote applications. Liquid cooling requires a series of tubes for the coolant to flow through, perhaps taking the heated fluid to another part of the plant for use in another process. This may sound efficient, but some applications cannot handle such a system, making the uncomplicated design of fans desirable.

In addition, there is an element of peace of mind with air cooling. Plant managers have safety and stability concerns with alternative systems such as liquid coolants, and many electrical engineers have reservations about liquid-based systems close to electrical equipment.

Liquid Coolants And Water Cooling

Initial safety concerns aside, liquid-cooling boasts a range of benefits for industrial applications beyond those offered by traditional air-based systems. In fact, these systems deliver where fans often fall short.

Water-based systems operate on a very simple principle of heat exchange: pipes containing a liquid coolant circulate around an enclosure and out of the device to keep components

cool. The excess heat generated during use is transferred to the water in the pipes. This heated water is then available for other industrial processes, such as plant heating.

The development of cooling technology in recent years means it now poses almost no risk of causing an electrical fire, due to the self-contained housing these systems include as standard, ensuring that no water can reach critical components in the unlikely event of a coolant leak.

Despite this added safety measure, the system remains compact so it can easily be fitted into equipment without increasing the product footprint. This solves the problem faced by air cooling and means that water cooling can now be used where space is limited.

The biggest selling point for water cooling is its energy efficiency. Not only fans are limited by their reliance on ambient temperatures to cool components – after all, air conducts less heat than water – to provide sufficient cooling, multiple fans may be necessary. This increases the cooling system's energy consumption, not to mention its space requirements.

Conversely, water has a high heat conductivity so systems require only a single pump for effective cooling. Coolant can also be stored in the pipes at a temperature below a plant's ambient level for even faster cooling.

> The development of cooling technology in recent years means it now poses almost no risk of causing an electrical fire

Engineers looking to specify liquid cooling need to consider that, although it will have an eventual higher return on investment, the initial setup of liquid-cooled systems is more complex and costly than air-cooled ones.

Making The Best Choice

The first step in selecting the right cooling system is to identify the device and its typical operating challenges. For example, businesses in the electroplating sector will find that a lot of the copper sulphate particulates produced during production can get into enclosure vents of air-based systems and corrode core components and wiring. In these situations, liquid-cooled alternatives are obviously preferable.

Likewise, high-voltage electrical infrastructure generates more heat than a standard electronic device. To this end, OEMs such as REO UK can work with electrical engineers to determine the best solution for specific requirements.

The Reotron MDW 700 power controller, for example, is designed for use with industrial transformers and is available in both air- and water-cooled versions. For transformers expected to operate at more than 50A, it is advisable that businesses chose either additional fans for the air-cooled model or opt for the water-cooled systems, such as Reotron MDW 700 WK, to eliminate the risk of overheating.

The last step in choosing the right cooling system is to realistically determine how regularly it will be maintained. Fan filters need frequent replacement to avoid a build-up of contaminants that lead to overheating, so using fans in difficult-to-reach or infrequently-maintained applications is not recommended.

Just as it can be difficult to choose between standing in front of a fan and standing in a pool of water on a warm day, so too is selecting the right cooling system for industrial applications. But considering the factors mentioned here that affect equipment performance further down the line leads to the right choice every time. ●

Figure 1: Liquid-cooling occupies a smaller footprint than air-cooling systems

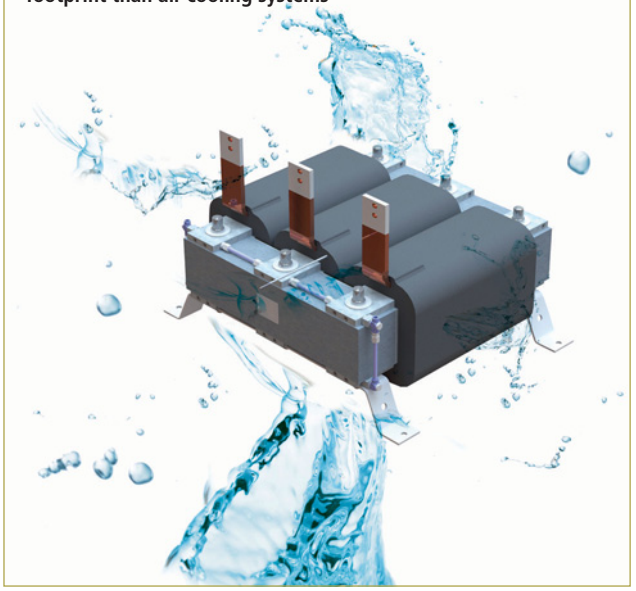

Figure 2: Water-cooled components have become increasingly popular in electronic design in recent years

AN EFFECTIVE CRITERION FOR LOW POWER ENCODING IN NETWORKS ON CHIP

BY **MEHDI TAASSORI** AND **SENER UYSAL** FROM EASTERN MEDITERRANEAN UNIVERSITY IN TURKEY

The Network on Chip (NoC) enables on-chip communication between several IP cores. An important goal of NoC design is to keep power consumption low, and low-power encoding algorithms help achieve that.

Here, we introduce an effective scheme for low-power encoding algorithms in the transport and data link layers of NoCs.

Power Consumption

Power consumption in NoCs is mostly from links, routers and network interfaces. Our low-power encoding method uses transition signalling (meaning a logical 1 is sent by a transition, which is a rising or falling edge, while 0 is represented by a lack of transition on the bus); it aims at reducing the number of 1s by data encoding to decrease the number of transitions on the bus, which will lead to lower power consumption.

To assess the efficiency of these methods, we consider the power consumption of a coder and a decoder, which are treated as design overheads. With our adaptive low-power encoding approach, a chip designer can easily determine if the applied low-power encoding scheme is cost-effective or not, and under which conditions.

Encoding techniques can be broadly grouped into two types: those that focus on power reduction through self-capacitance – where the impact of coupling capacitance is largely ignored, and those that work by minimising coupling transitions. The first group of coding algorithms is not appropriate for use in NoCs due to the coupling capacitance's significant impact on the total on-chip capacitance.

Transport Layer

In this article, we discuss the way our proposed scheme works in two cases: First, when the coder and decoder are located between the process element and the router, situated in the transport layer; this is known as end-to-end. Second, when the data is coded and decoded in each node, i.e. the coder and the decoder are in the data link layer.

In the transport layer, the data packets are coded, leaving the packet header unchanged. The header's information is used for routing data between sender and receiver.

Data encoding in the transport layer is shown in Figure 1. A transport layer has two advantages over a data-link layer. First, the reduced transition signalling coded data passes through intermediate routers, thus the power dissipation in various parts of the router (such as buffers) decreases. In the data link layer, uncoded data is delivered to the router.

Second, in the transport layer the distance between encoder and decoder increases, therefore the data encoding approach helps reduce power even more.

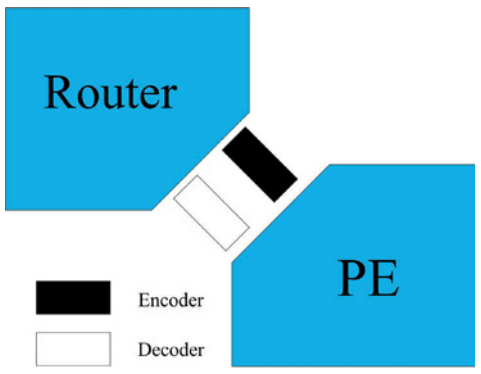

Figure 1: Data encoding algorithm in the transport layer

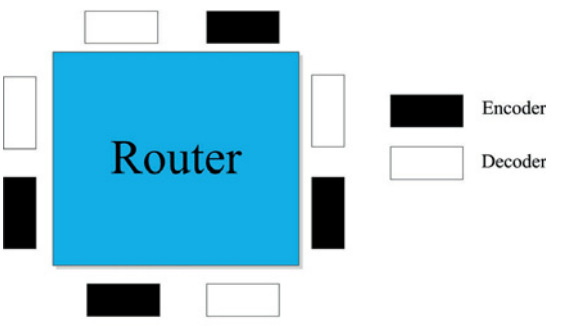

Figure 2: Data encoding algorithm in the data link layer

In the data link layer, the encoder-decoder distance is between adjacent routers.

The aim of the proposed scheme is to determine if the applied data encoding algorithm helps reduce power dissipation in NoCs; i.e., whether power dissipation after applying the encoding technique is lower than before, when not applying it.

Sending a message between nodes i and j in the NoC before using the encoding algorithm consumes power according to the following equation:

$$P = \beta P_R + \beta P_{NI} + (\beta - 1)P_L \quad (1)$$

where β is the number of nodes, P_R is the router's power dissipation, P_{NI} is the power used by the network interface, and P_L is the link's power consumption.

The power required to send a message between nodes i and j in the NoC after applying encoding algorithm is:

$$P_C = P_{cod} + P_{dec} + \beta P_{RC} + \beta P_{NIC} + (\beta - 1)P_{LC} \quad (2)$$

where P_C is power consumption when the data encoding algorithm is applied, P_{cod} and P_{dec} are the power dissipation of the encoder and decoder respectively, P_{RC} is the power consumption of the router after coding, P_{NIC} is the power dissipation of the network interface and P_{LC} is the power of the link. P_{cod} and P_{dec} are overheads of the low-power encoding approach. In fact, data-encoding algorithms compensate for this overhead in addition to reducing the power consumption.

The following condition should be satisfied to decrease power consumption using the low-power encoding algorithm:

$$P_C < P \quad (3)$$

Using Equations 1-3, we can determine β, which is the number of nodes:

$$\beta = \frac{P_{cod} + P_{dec} + P_L - P_{LC}}{P_R - P_{RC} + P_{NI} - P_{NIC} + P_L - P_{LC}} \quad (4)$$

$$\min(\beta) = \frac{P_{cod} + P_{dec} + P_L - P_{LC}}{P_R - P_{RC} + P_{NI} - P_{NIC} + P_L - P_{LC}} \quad (5)$$

$$(\min(\beta))' = \frac{P_R - P_{RC} + P_{NI} - P_{NIC} - P_{cod} - P_{dec}}{(P_R - P_{RC} + P_{NI} - P_{NIC} + P_L - P_{LC})^2} \quad (6)$$

The value of β in Equation 5 is the minimum number of nodes needed to apply the low-power encoding algorithm; otherwise the data encoding approach won't be efficient. In other words, if $P_R - P_{RC} + P_{NI} - P_{NIC} + P_L - P_{LC}$ is greater than zero then Equation 5 shows the minimum distance between sender and receiver that will work to reduce power dissipation based on the encoding technique. It means that the number of nodes between sender and receiver should be greater than β.

If $P_R - P_{RC} + P_{NI} - P_{NIC} + P_L - P_{LC}$ is less than zero, then using these encoding approaches is useless because the power reduction in the link is less than the encoding overhead.

If the value of $P_R - P_{RC} + P_{NI} - P_{NIC}$ is equal to $P_{cod} + P_{dec}$, then the value of β is one and the proposed scheme is useless. Hence, for a useful criterion those values should not be equal. It is impossible that $P_R - P_{RC} + P_{NI} - P_{NIC} = P_{cod} + P_{dec}$ because β can't be less than one.

Equation 5 is a descending function, so whenever $P_L - P_{LC}$ increases, the minimum value of β decreases. In other words, with increasing the link's power dissipation after using the data encoding algorithm, the minimum number of nodes is increased and a greater distance between them is needed for an effective encoding algorithm.

It is obvious that as the number of nodes increases, the efficiency of the low-power encoding approach decreases. According to Equation 5, with a decreasing $P_R - P_{RC} + P_{NI} - P_{NIC}$, the number of the nodes increases. Based on this analysis, it can be concluded that if the distance between sender and receiver is smaller than β then the data encoding algorithm is useless, and encoding increases, not decreases power dissipation.

In advanced VLSI technology, the power dissipation of encoder and decoder decreases. On the other hand, the wire length between routers is almost constant, or even longer. Based on Equation 5, in the transport layer, with decreasing power consumption of encoder and decoder, the minimum distance between the routers also decreases. In other words, using the data encoding algorithm is useful even for adjacent nodes.

Data Link Layer

In the data link layer of the NoCs, coder and decoder are placed among the routers. Data is coded and decoded in the nodes and, consequently, uncoded data is transmitted between them. Figure 2 shows the data-encoding algorithm in the data link layer.

The data link layer has two advantages compared to the transport layer. First, when using a data encoding algorithm with spatial redundancy, bus width becomes greater in the transport layer because data from buffers and switches grows, which in turn increases power consumption. Whereas in the data link layer, uncoded data is delivered to the routers.

The second advantage of the data link layer relates to data consecutiveness. The packets from various sources (depending on the size of the NoC, there could be many nodes, or routers, each with coder and decoder) could suffer interference in the transport layer, reducing the algorithm's efficiency.

The power required to send a message between adjacent nodes i and j in the NoC, before applying low power encoding, is:

$$P = 2P_R + P_L \quad (7)$$

Power dissipation after using the low-power encoding algorithm is:

$$P_C = P_{cod} + P_{dec} + 2P_R + P_{LC} \quad (8)$$

According to Equation 8, after applying the low-power encoding technique, the router's power requirement is the same as before using this approach. As mentioned earlier, in the data link layer, uncoded data enters the routers, so power consumption of the router is the same as without the algorithm. The link's power dissipation is given by:

$$P_L = (\alpha_L C_L + \alpha_C C_C) V_{dd}^2 f \quad (9)$$

where α_L is the link's transition, C_L is its capacitance, α_C is the coupling's transition, C_C is its capacitance, V_{dd} is the input voltage of the system and f is the clock frequency.

The link's power consumption after applying data encoding algorithm P_{LC} is:

$$P_{LC} = (\alpha_L^C C_L + \alpha_C^C C_C) V_{dd}^2 f \quad (10)$$

where α_L^C and α_C^C are link- and coupling transitions respectively within the encoding scheme. According to the International Technology Roadmap for Semiconductors, the value of the coupling capacitance is n times the link's capacitance, where n is a constant value; see below:

$$C_C = nC_L \quad (11)$$

The criterion for effective encoding in the data link layer is given by:

$$(\alpha_L - \alpha_L^C) + n(\alpha_C - \alpha_C^C) = \frac{P_{cod} + P_{dec}}{C_L V_{dd}^2 f} \quad (12)$$

$$\min\left[(\alpha_L - \alpha_L^C) + n(\alpha_C - \alpha_C^C)\right] = \frac{P_{cod} + P_{dec}}{C_L V_{dd}^2 f} \quad (13)$$

Equation 13 gives the minimum differences between numbers of link and coupling transitions before and after the low-power encoding algorithm has been used. Based on this equation, the number of transitions after using the encoding technique should reduce further to compensate for the power overhead of the scheme. Otherwise, the coding algorithm will not work as it will be unable to compensate the overhead's power consumption. In other words, to have a successful encoding approach, the number of transitions should reduce following coding to compensate for the power of the coder and the decoder. ●

RELEC ELECTRONICS LTD

RIA12 compliant train-borne dc dc converter

The URB series from Mornsun are a range of rugged ultra-wide input dc dc converters, when used in combination with the specially designed FC series input filters they conform to the challenging requirements of EN50155 and RIA12 for train-borne applications.

Available in 3 input ranges:
24Vdc input (range 9 to 36Vdc)
48Vdc input (range 18 to 75Vdc)
72V, 96V, 110V & 120Vdc input (range 40 to 160Vdc)
Output voltages: 3.3V, 5V, 9V, 12V, 15V & 24Vdc
Power rating: 6W, 10W, 15W & 20W
Mounting: PCB; chassis mount or DIN rail
Efficiency: 90%
Isolation: 1.5kVdc
Cooling: convection
Protection: reverse polarity; output short circuit; over voltage
Lead time: 4 weeks

By virtue of their design for the harsh environment of the railway, they are also suitable for many other applications requiring a compact rugged dc dc solution. Applications include: passenger reading lights; on-board Wi-Fi; passenger USB hubs; sensor control modems.
The URB series and the filters are very competitively priced, for further information or to discuss your application please contact our technical sales team.

POWER DISPLAYS EMC

www.relec.co.uk

Tel: 01929 555800
e-mail: sales@relec.co.uk

Professional 8, 10 & 20 ports USB 2.0 charge only or charge & sync hubs

The Powersolve 8, 10 and 20 port USB charge only, or charge & sync hubs offer the perfect solution to charge and sync several tablet PC's or other USB devices simultaneously. Also ideal for automated duplication tasks such as USB memory sticks etc, where data can be transferred to multiple devices at the same time. Ideal for use in schools, business or any application where multiple charging and or data transfer is required.

Model PLV60-USB 5V USB Hub
This charge only model has 10 output ports which are switchable for 10 x 1A or 5 x 2.4A 5V outputs. Integral power supply will operate from a 90-264VAC universal AC input.

Model PLV120-USB 5V USB Hub
This 10 port, charge only, model is available with 10 independent 5V 2.4A outputs. Features smart charging IC and can charge any device using USB charging technology. Integral power supply operates from 90-264VAC input.

New!

Model PSUSB-0824 5V charge & sync hub
Provides 8 output ports with up to 5V 1.5A (CDP mode) or 5V 2.4A (DCP mode) Supports iOS and Android devices. Supports high speed 480 Mbps, full speed 12 Mbps and low speed 1.5 Mbps operation. Up to 8 units can be cascaded to increase output ports to 64. Housed in compact aluminium case measuring 170 x 80 x 30mm.
Powered by external 120W power supply with 90-264VAC input.

Model PSUSB-20CH charge & sync hub
Charges and syncs up to 20 devices. Charge current 1.1A per port in charge mode. Supports high speed 480 Mbps, full speed 12 Mbps and low speed 1.5 Mbps operation. Compatible with all USB compliant devices. 2 x 20 port hubs can be connected in cascade mode increasing number of ports to 40.
Housed in metal enclosure which measures 268 x 102 x 40mm
Powered by external 150W power supply with 90-264VAC input.

Models PSUSB-1024 & PSUSB-2024 10 & 20 port charge & sync hubs
Provides up to 1.5A (CDP mode) or 2.4A (DCP mode) Supports iOS and Android devices. Features Green Energy mode switch. The 10 & 20 Port hubs can be cascaded with another unit to double the number of ports. Supports high speed 480 Mbps, full speed 12 Mbps and low speed 1.5 Mbps operation. These models feature integral fan cooling via temperature controlled fans and will shutdown with over temperature. Compact metal housing measures 268 x 102 x 40mm
Powered by external 150W (10 ports) or 288W (20 ports) power supply with 90-264VAC input.

New!

www.powersolve.co.uk

Tel: 44-1635-521858 Email: sales@powersolve.co.uk

F-P FIBRE OPTIC ULTRASONIC SENSORS IN PARTIAL-DISCHARGE MONITORING OF POWER TRANSFORMERS

BY **WANG WEI**, **GAO CHAOFEI**, **SHANG CHAO** AND **LIU HAN** FROM THE STATE KEY LABORATORY OF ELECTRICAL POWER SYSTEMS FROM RENEWABLE ENERGY SOURCES IN BEIJING, AND **LI FUPING** FROM THE HANGZHOU ELECTRIC POWER SUPPLY BUREAU IN CHINA

Transformers are an important part of power systems, increasing and decreasing voltage in the transmission, distribution and use of electricity. Because they handle a lot of electricity, they need cooling.

There are three types of cooling mechanisms, oil-immersed, dry and inflatable (SF6). High-voltage transformers are mostly of the oil-immersed type.

The typical power transformer has several coil layers to handle different voltages; each layer is wrapped in insulating paper and immersed in oil.

During the operation of a high-power transformer, a partial discharge (PD) can be generated in the insulation system, causing malfunctions and even failure.

PD also causes significant losses, and unstable and unsafe operation of the power system. Therefore, monitoring it is an important part of the system's condition monitoring and fault diagnosis.

Partial Discharge Detection

The insulation paper ages, deteriorates and warps, and with that the electric field distorts – this is when partial discharge occurs the most. The energy released during partial discharge can produce local perturbations on the transformer's oil-paper insulation in the form of ultrasonic waves. Unencumbered by electromagnetic interference, these waves can be located and measured.

The traditional detection method based on the piezoelectric effect is to convert acoustic into electrical signals. By placing a piezoelectric ultrasonic sensor on the transformer's casing, PD can be pinpointed and measured. Since the transformer's enclosure causes multipath and attenuation effects on the ultrasonic signal, this measurement method is less sensitive and hence not so accurate.

Rapid developments in optical fibre technology have given rise to another PD measurement – the fibre optic sensor. This type is less susceptible to electromagnetic interference, highly sensitive, smaller in size and can easily be built into the oil tank.

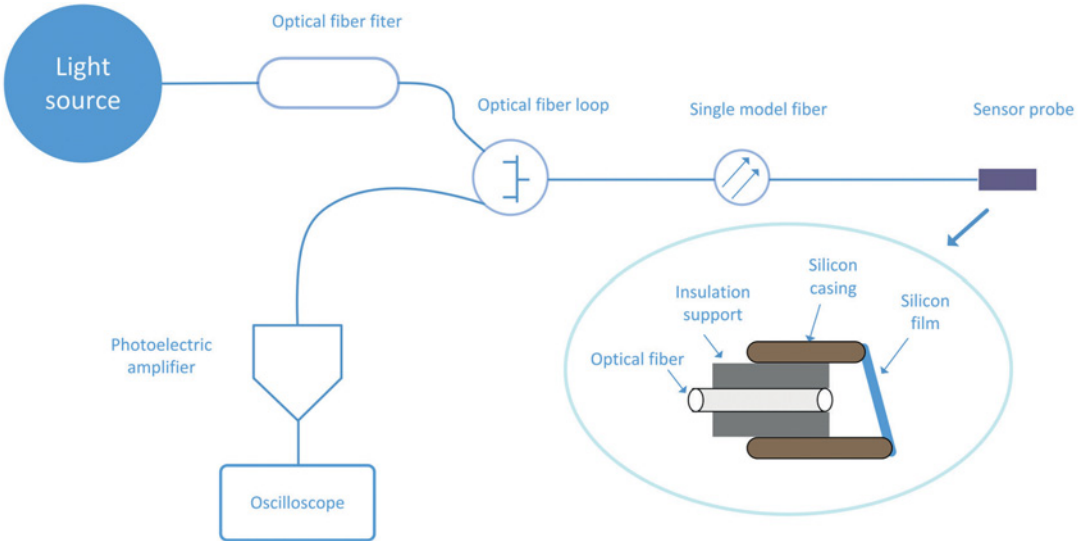

Figure 1: Structure of the F-P optical fibre ultrasonic detection system

Figure 2: Solid transformer field-test setup

At present, optical sensors for PD detection mainly include the Michelson, Mach-Zehnder and Fabry-Perot (F-P) types. Figure 1 shows the structure and size of an F-P sensor.

Here we describe the principle and design of an F-P fibre optic ultrasonic sensor and analyse its frequency response characteristics, test results and applications.

F-P Sensor Design

As shown in Figure 1, the F-P fibre optic ultrasonic detection system consists of a sensor probe, light source, coupler, photoelectric signal processor and single-mode optical fibre, while the sensor probe consists of an optical fibre, a casing pipe and a diaphragm.

The F-P interference cavity is between the optical fibre and the inner surface of the diaphragm. The diaphragm detects the ultrasonic signal to be measured and vibrates. This vibration is converted into optical phase changes through F-P cavity beam interference, with varying light intensity. Detecting this light enables the signal measurement.

The sensitivity and frequency response of the diaphragm is crucial, and is largely influenced by the diaphragm's material. The diaphragm's sensitivity can be improved by increasing its radius, however this comes at the cost of its inherent frequency response. Therefore, it's important to choose the size of the diaphragm to achieve the best tradeoff between sensitivity and frequency response for the design.

For our quartz diaphragm, we selected a radius of 0.9mm and thickness h of 30um, which yields 60nm/kpa sensitivity and the highest frequency response of 101kHz.

Optimal Sensor Design

To improve the sensitivity of the sensor, in addition to diaphragm sensitivity, light intensity is just as important, since according to the Fabry-Perot interference principle, more light improves the output's intensity.

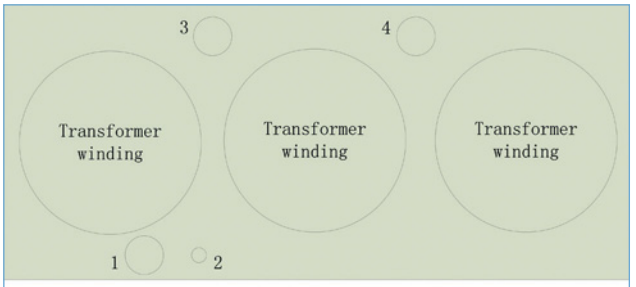

Figure 3: Test plan

The sensitivity of the sensor is related to the interference phase of the two beams, which are the light source and the reflection from the silicon film. When the power of the light source is below 2mW, loss in the optical path is larger, but it improves significantly at 2mW and above.

The needle-plate electrode model is a commonly-used discharge model in engineering; it is a steel needle perpendicular to a plate,

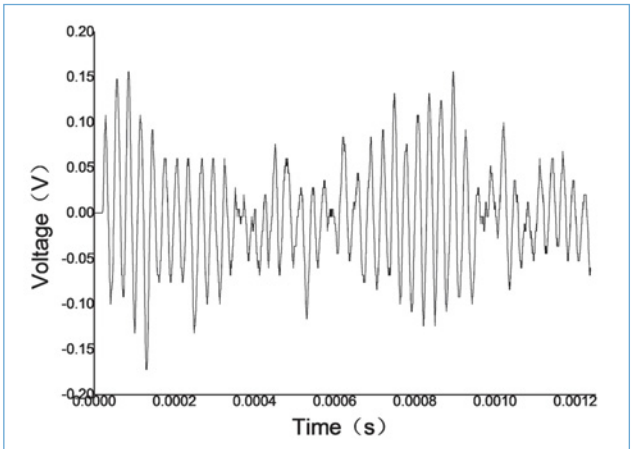

Figure 4: Partial discharge waveform of the sensor in position 2 at a 2m distance

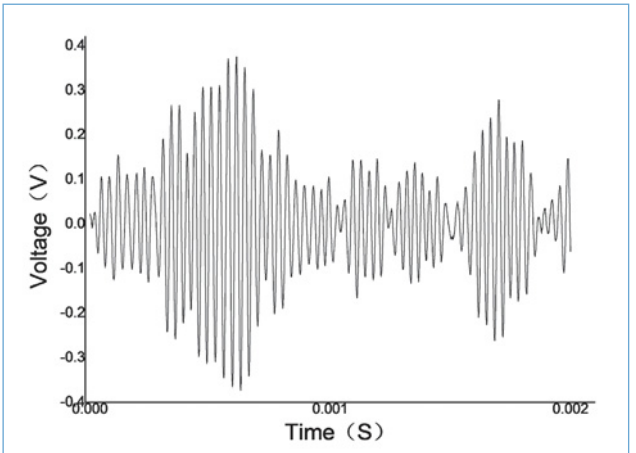

Figure 5: Partial discharge waveform of the sensor in position 3 at a 2m distance

PEAK® electronic design ltd

2 Year Warranty as standard

UK designed **and** manufactured

Atlas DCA - Semiconductor Analyser model DCA55

Now with backlit display and AAA battery!

- Connect any way round.
- Automatically identify component type.
- Automatically identify pinouts.
- Supports Transistors, MOSFETs, Diodes, LEDs and more.
- Measure transistor gain (h_{FE}).
- Measure V_{BE} (now 1mV resolution!)
- Measure leakage current.
- Measure MOSFET thresholds.
- Measure LED voltages.

Designed in the UK and made in the UK.

Just a few example screen shots:

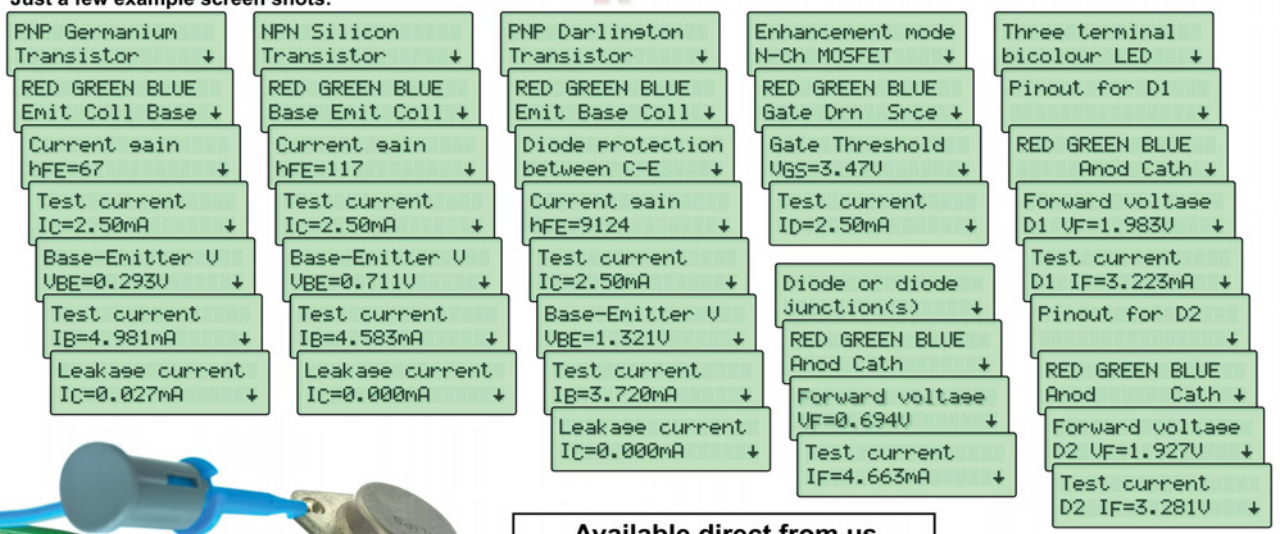

Universal test clips for large and small components

Available direct from us and from our distributors:

RS · maplin (The electronics specialist) · Rapid electronics · Farnell

Peak Direct Price:

£48.00

(£40.00+VAT)

Peak's insured UK delivery £3.00 inc. VAT

It's only possible to show summary specifications here. Please ask if you'd like detailed data. Further information is also available on our website. Product price refunded if you're not happy.

E&OE

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

Atlas House, 2 Kiln Lane
Harpur Hill Business Park
Buxton, Derbyshire
SK17 9JL, UK

Follow us on twitter for tips, tricks and news.
@peakatlas

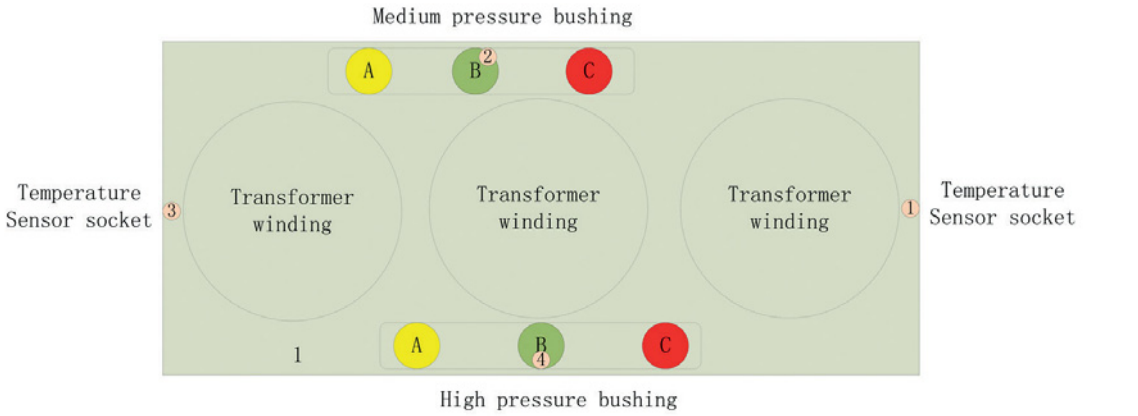

Figure 6: 110kV transformer top sensor optimal layout

separated by a distance of a few millimeters. When voltage is applied, an ultrasonic signal is generated.

The ultrasonic signal received by the sensor is analysed in Labview. Through spectrum analysis, we found that the sensor's centre frequency is 94kHz, close to the theoretical centre frequency of 101kHz.

We then repeated the test with oil. We placed the needle-plate electrode 25cm below the oil surface, and the F-P sensor 50cm from the power supply. We found the sensor's centre frequency in the transformer oil to be 30kHz, lower because the viscosity of oil is greater than that of air.

Transformer Test

Figure 2 shows a three-phase 110kV oil-immersed transformer, measuring 4.5m×2m×2.5m. The discharge model and sensor are placed on top of the transformer. During the test, we added a partial discharge detection device to measure the discharge magnitude, the unit of which is pC (10^{-12} Coulombs).

The experimental procedure (see Figure 3) is as follows: the needle-plate electrode model is placed deeply inside the transformer from the larger flange hole 1, with the sensor suspended from the smaller flange holes 2 and 3.

We record the signal waveform from the sensor. The partial discharge volume is controlled at about 100pC. Figure 4 shows the partial discharge waveform measured 2m deep at small cap 2, whereas Figure 5 displays the partial discharge waveform measured at 2m deep at small cap 3. Figure 3 shows the top of the transformer, with the specific locations of caps 2 and 3.

Figures 4 and 5 show that when the power supply depth is 2m, the sensor in positions 2 and 3 detects the partial discharge waveform with the signal amplitude of 150-400mV. Since the position-3 sensor also detects the waveform, this shows that the signal travels between the phase windings, so the sensor meets the requirements.

Flange hole 4 also holds a sensor for testing, but here no signal is detected, showing that the sensor is influenced by the coil and iron core. Installing multiple sensors at different places inside the transformer helps meet the detection requirements.

When the ultrasonic signal generated by the partial discharge propagates inside the transformer, it also reflects from the windings, core and other media. A 300mV signal can be detected 2m from a direct partial-discharge signal of 100pC. Thus, according to the transformer's internal structure and

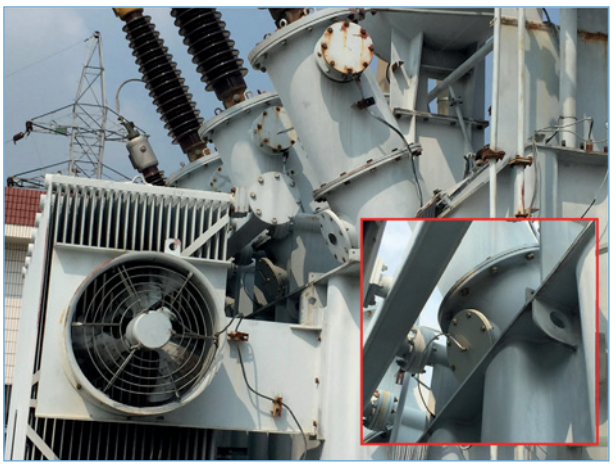

Figure 7: 110kV transformer top sensor location: (a) **(b)** **(c)**

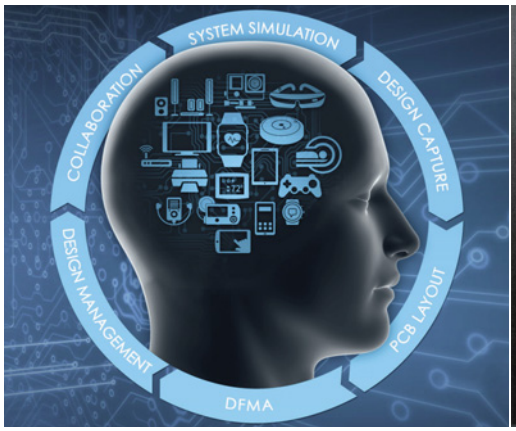

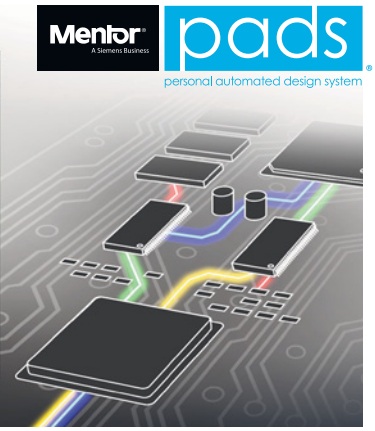

PADS® Strengthens DFT Capability with XJTAG® Boundary Scan Know-How

"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

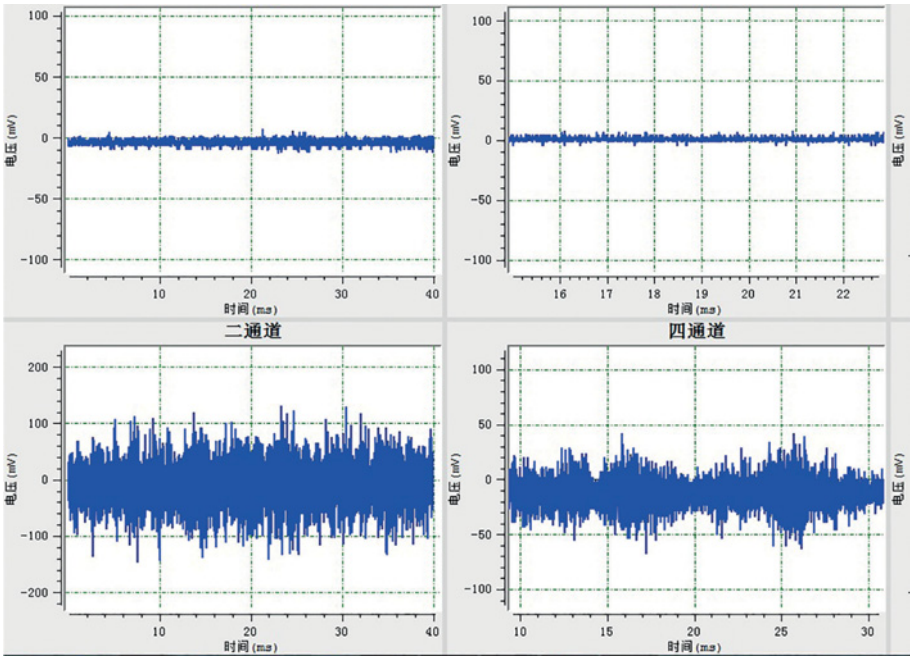

Figure 8: Background noise of each channel

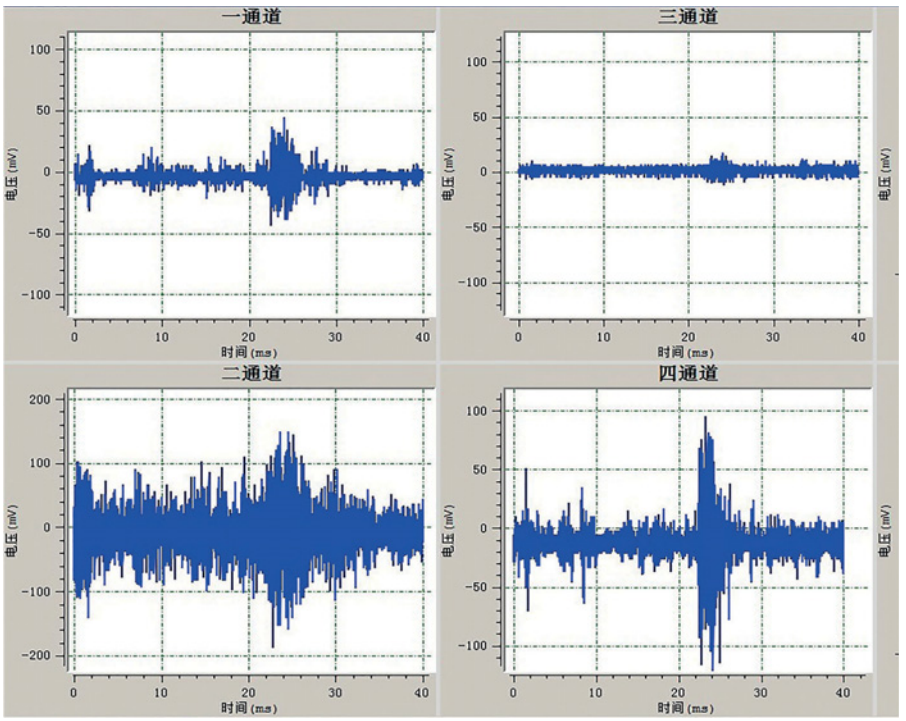

Figure 9: Detected partial discharge waveforms

the propagation rules for ultrasonic signals, the detection requirements can be achieved just by installing a sensor in the middle of the transformer's top; see sensors 2 and 4 in Figure 6.

Sensor Applications

An F-P optical fibre partial-discharge on-line monitoring system has been installed in a substation in Jiangsu, China; see Figure 7. The PD detection and collection devices are installed in an outdoor air-conditioned test cabinet. The site's monitoring device and its server are linked through an Ethernet cable via a LAN interface.

A smaller similar partial discharge signal is also observed during the test, as shown by the four channel waveforms in Figures 8 and 9, considered sporadic PD; the amplitude of the four-channel pulse is about 90mV.

So we have determined that in practical applications, the F-P fibre optic ultrasonic sensor can be used to measure low partial-discharge signals.

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

INDUSTRY'S FIRST MCU WITH INTEGRATED 2D GPU AND DDR2 MEMORY

Microchip announced the 32-bit PIC32MZ DA microcontroller (MCU) family, claimed to be the industry's first MCU with an integrated 2D Graphics Processing Unit (GPU) and up to 32MB of integrated DDR2 memory. This combination enables higher colour resolution and display size, up to 12 inches.

The PIC32MZ DA family bridges the graphics performance gap between MCUs and microprocessor units (MPUs) for customers who want to stay in the familiar design environment of an MCU. The devices provide MPU-like graphics capabilities with the seamless integration and programming model of Microchip's PIC32 and MPLAB IDE and Harmony software framework. These tools provide a visual graphics design environment, custom display driver creation, graphics libraries and an asset converter that can take a custom graphic and optimise it for their chosen display size.

The addition of DDR2 memory, a first for MCUs throughout the industry, enables 2x faster throughput and large graphics buffers and/or storage for increasingly complex communications protocol stacks and algorithms. The result is smooth, striking interfaces and fewer product variants in a crowded communications control market.

www.microchip.com

HIGH-END ANALOGUE RF AND MICROWAVE SIGNAL GENERATOR

The Rohde & Schwarz R&S SMA100B with a frequency range up to 20GHz is one of the most powerful analogue signal generator on the market. It delivers signals with the lowest possible phase noise and the highest output power with extremely low harmonics. Engineers no longer need to compromise between output power and a spurious free dynamic range. The R&S SMA100B is designed for the RF semiconductor, wireless communications and aerospace and defence industries.

The R&S SMA100B provides purest signals with the lowest possible phase noise at all offset frequencies (1GHz, -152dBc/Hz, 20kHz offset). A 6GHz instrument generates up to 38dBm RF output power, and a 20GHz instrument generates up to 32dBm in the microwave frequency range. Harmonics are extremely low across the entire frequency range; above 6GHz they are even significantly lower than 70dBc at 18dBm output power. Nonharmonics are below 110dBc at an output signal of 1GHz.

In addition, the R&S SMA100B is the world's only analogue signal generator that can simultaneously provide a second, independently configurable, extremely pure and synchronized clock signal up to a frequency of 6GHz.

www.rohde-schwarz.com

YOKOGAWA LAUNCHES NEW DL350 SCOPECORDER

The Yokogawa DL350 ScopeCorder is a comprehensive, fully-portable measuring instrument available for capturing, displaying, recording and analysing a wide variety of electrical and physical parameters in industry sectors including automotive, electronics, energy, transport and mechatronics.

Like other models in the Yokogawa ScopeCorder family, the DL350 combines features of a general-purpose oscilloscope and those of a high-performance data acquisition recorder in a single, portable instrument. Unlike alternative portable measuring solutions, such as oscilloscopes and combined oscilloscope/multimeters, this ScopeCorder adds very high levels of precision and accuracy to field measurements, isolated inputs for measurements at high voltage levels, and long-memory capabilities that allow long-term recording for many hours or even days.

Another key feature is its plug-in modularity, which allows configuration for a variety of applications. Whether carrying out straightforward high-precision voltage measurements or handling a blend of signals, the DL350 can deliver, without extra boxes or cables. This flexible input capability is achieved by incorporating two slots which are populated with any of 18 different types of user-swappable input modules.

www.tmi.yokogawa.com

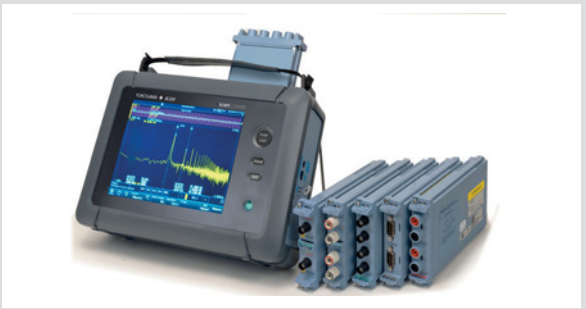

NEW LINE OF STRAIGHT WAVEGUIDE SECTIONS IN SIZES WR-10 TO WR-137

Pasternack, a provider of RF, microwave and millimeter wave products, has released a new series of straight waveguide sections in sizes ranging from WR-10 to WR-137. Typical applications include instrumentation, test benches, high-efficiency RF/microwave transmission, SATCOM, MILCOM, radar and telecom networks.

Pasternack's new line of straight waveguide sections consists of 62 models that operate in the frequency range of 5.85-110GHz and in 13 waveguide bands from C to W band. They also deliver VSWR as low as 1.03:1.

The waveguide straights are available in section lengths from 3-12 inches and are made of either painted copper alloy or gold-plated, oxygen-free hard copper (OFHC); they feature UG, CPR and UBR-style flanges.

Pasternack's straight waveguide sections are in stock and ready for immediate shipment with no minimum order quantity. These waveguides are ideally suited for aerospace, defence, industrial, test and instrumentation, telecom and medical industries.

"This full range of waveguide straights support over 13 bands. They perfectly complement our rapidly expanding waveguide component portfolio and deliver quality construction and consistent performance," said Steven Pong, Product Manager at Pasternack.

www.pasternack.com

NEW LINEAR HALL-EFFECT SENSOR ICS WITH ANALOGUE OUTPUT

Allegro MicroSystems Europe has announced two new linear Hall-effect sensor ICs with analogue output, targeted at the automotive and industrial markets. New applications for linear output Hall-effect sensor ICs, such as displacement and angular position, require higher accuracy and smaller package sizes. Allegro's A1308 and A1309 devices have been designed specifically to meet both requirements. Target applications include EPS – torque sensing, EPS – angle sensing, and transmission – fork position sensing. These temperature-stable devices are available in both surface-mount and through-hole packages.

The accuracy of each device is enhanced via end-of-line optimisation. Each device features nonvolatile memory to optimise device sensitivity and the quiescent voltage output (QVO: output in the absence of a magnetic field) for a given application or circuit. This A1308 and A1309 optimised performance is sustained across the full operating temperature range by programming the temperature coefficient for both sensitivity and QVO at Allegro end-of-line test. These ratiometric Hall-effect sensor ICs provide a voltage output that is proportional to the applied magnetic field. The quiescent voltage output is adjusted around 50% of the supply voltage.

www.allegromicro.com

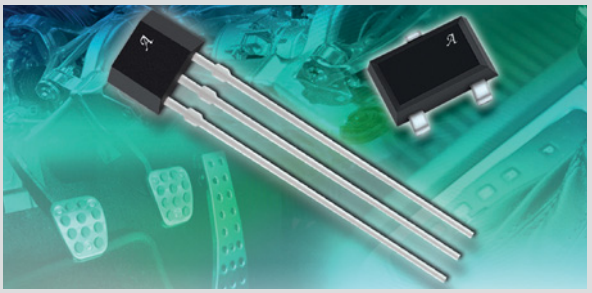

FISCHER CONNECTORS JOINS SOLARSTRATOS AS OFFICIAL PARTNER

The Fischer Connectors Group, based in St-Prex, Switzerland, joins the SolarStratos project as Official Partner to provide mission-critical connectivity to the solar-powered HB-SXA plane. This first manned solar plane to penetrate the stratosphere aims to demonstrate the potential of solar energy, and explore new territories for the future of aviation.

SolarStratos HB-SXA is propelled solely by solar energy, backed by a Li-ion battery. It will be the first piloted solar-powered aircraft to enter the stratosphere, and the first solar-powered twin-seater commercial aircraft in history. Designed and built in 2016 in Switzerland, SolarStratos took its maiden flight in Payerne, Switzerland, on May 5th. With a wingspan of 24.8m and a length of 8.5m, it weighs just 350kg. The project is led by SolarStratos's founder and pilot Raphaël Domjan, supported by an international team of experienced professionals and partners including engineers, a former astronaut, pilots, meteorologists and doctors.

Fischer Connectors's solutions were selected for their high reliability in extreme environments, optimal ratio of light weight and high performance, large volumes data transmission over long distances, and so on.

www.fischerconnectors.com

PARTNERS JOIN FORCES TO DEVELOP SMART-FACTORIES ARTIFICIAL INTELLIGENCE CHIP

ROHM Semiconductor and A*STAR's Institute of Microelectronics (IME), a research institute under the Agency for Science, Technology and Research (A*STAR), announced the joint development of an artificial intelligence (AI) chip to boost efficiency in predictive maintenance for smart factories.

The concept of 'Predictive Maintenance' has become widespread in the manufacturing industry as manufacturers begin to digitalise their production lines for increased productivity and competitiveness. Predictive maintenance forecasts machine failures and it involves monitoring the function and health of machines and identifying potential problems based on data received through device logs and sensors, eventually taking counter-measures such as repairing or replacing the affected machine.

Leveraging ROHM's original AI analytical algorithms and IME's capabilities in ultra-low power analogue/digital integrated circuit and systems, the collaboration will focus on a chip to filter volumes of data across multiple sensors, and analyse complex data patterns in real-time. This chip is expected to perform significantly faster than the conventional method for predictive maintenance, as well as reduce power consumption, paving the way for greater asset productivity and lower overall maintenance costs.

www.rohm.com/eu

SMARTKEM DELIVERS FIRST INDUSTRIAL GEN 2.5 OTFT PROCESS TO ASIA

SmartKem, a developer and supplier of high-performance organic semiconductors and organic thin film transistor (OTFT) technology, is collaborating with partners in Asia to deliver the world's thinnest and most robust industrialised display backplane technology on the market today.

SmartKem has now finalised an industrial grade process for the mass production of OTFT backplanes on large area glass or plastic. This has been achieved using its bespoke pilot line, based at the UK's Centre for Process Innovation (CPI) and its synthesis and formulation Technology Centre based at Hexagon Tower in Manchester. The approved process is now being transferred to its production partners in Asia.

SmartKem's industrialised Gen 2.5 POR is now undergoing transfer to customer production lines for the manufacture of product prototypes and scaled display production. It is anticipated that the first display products manufactured through this collaboration will hit the Asian market within 12-18 months.

SmartKem's groundbreaking truFLEX semiconductor platform offers display makers a fast and low risk route to market with highly differentiated product using existing production equipment on either glass or plastic substrates – delivering a new paradigm in display performance and user experience.

www.smartkem.com

DO YOU WANT THE BEST ELECTRONICS DESIGN SOFTWARE?

FEATURES

- Schematic Capture
- PCB Layout
- Gridless Autorouting
- 3D Visualization
- M-CAD Integration
- SPICE Simulation
- MCU Co-simulation
- Built in IDE
- Visual Programming

Labcenter Electronics

www.labcenter.com

Tel: +44 (0)1756 753440

More new products **in stock** than any other distributor.

Order now at
mouser.co.uk

Mouser and Mouser Electronics are registered trademarks of Mouser Electronics, Inc. Other products, logos, and company names mentioned herein, may be trademarks of their respective owners.

MOUSER ELECTRONICS

The Newest Products for Your Newest Designs®