PSSD EUROPE
Power Systems Design: Empowering Global Innovation
January/February 2013

Special Report: Energy Efficiency; Test & Measurement (pg35)

PSD EUROPE

Power Systems Design: Empowering Global Innovation

WWW.POWERSYSTEMSDESIGN.COM

Visit us online for exclusive content; Industry News, Products, Reviews, and full PSD archives and back issues

COVER STORY

Faster switching in inverters (pg 13)

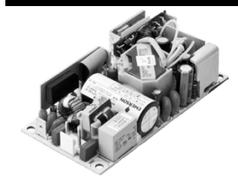

Highlighted Products News, Industry News and more web-only content, to:
www.powersystemsdesign.com

PSSD EUROPE

AGS Media Group
146 Charles Street
Annapolis, MD 21401 USATel:
+410.295.0177
Fax: +510.217.3608
www.powersystemsdesign.com

Editor-in-Chief
Joshua Israelsohn, Editor-in-Chief,
Power Systems Design
joshua@powersystemsdesign.com

Contributing Editors
Gail Purvis, European Editor,
Power Systems Design
Gail.Purvis@powersystemsdesign.com

Liu Hong, Editor-in-Chief, Power Systems Design China
powersdc@126.com

Ash Sharma, IMS Research
Ash.*sharma@imsresearch.com*

Dr. Ray Ridley, Ridley *Engineering*
RRidley@ridleyengineerng.com

David Morrison, How2Power
david@how2power.com

Publishing Director
Jim Graham
jim.graham@powersystemsdesign.com

Publisher
Julia Stocks
Julia.stocks@powersystemsdesign.com

Circulation Management
Kathryn Philips
kathryn.phillips@powersystemsdesign.com

Magazine Design
Louis C. Geiger
louis@agencyofrecord.com

Production Manager
Chris Corneal
chris.corneal@powersystemsdesign.com

ISSN number: 1613-6365

Volume 10, Issue 1

Simulation anyone?

One of the more interesting *automated design* developments recently at the University of Dundee is in the field of advanced drugs, which for complex conditions and multiple targets, cause complex and difficult chemistries in need of optimisation. Dundee took the US baseball's *Moneyball* approach by using advanced statistical and data analysis techniques.

Thanks to sensors, improved data capture and organised drug design databases, the developed algorithm has been able to learn from the human experience of drug design, mimicking it on a massive scale, and iterating it, to solve complex design problems. America's ITRS (international technology roadmap for semiconductors) estimates that the use of such simulation tools reduced chip development times and costs by about 40%.

It has certainly been the golden mean for GSS (gold-standard simulations) software that in its latest research where the 3D simulator ran simulations on random dopants, line edge roughness and metal gate granularity as main sources of statistical variability on an unprecedented statistical scale.

It emerges with the news that "If you can develop a *metal-gate-last* 28 nm FD-SOI technology you will be able to achieve an astonishing SRAM supply voltage, in the range of 0.5 to 0.6 V," says expert Professor Asenov. Small wonder GSS is part the Fraunhofer Institute's SUPERTHEME project for circuit stability under process variability and electro-thermal-mechanical coupling.

At the Fraunhofer Institute for Mechanics of Materials project manager Dr Dirk Helm and his scientists have use simulations to develop shape memory alloys, developing various objects, such as minuscule forceps for endoscopy, where its small dimensions and elasticity, can be thoroughly sterilized with no joints. Numerical simulation models allow researchers to calculate in advance the components, the strength and clamping force, and efficiently develop and manufacture the component.

"By using simulations, we can avoid producing most of these prototypes," said Dr Helm. " This saves costs because the raw materials for the shape memory alloys are very expensive and are sometimes difficult to work with." In addition, the researchers can estimate how durable the materials are through the simulations.

So smaller wonder, perhaps, that news from IT design-conscious Apple, is that it is to invest in WITNESS from Lanner—simulation software originally developed in the OR department of British Leyland Motors. The licence and maintenance agreements will put WITNESS at the heart of simulating and optimising business processes within Apple enable it to represent real world processes in a dynamic animated computer model and then experiment with alternative *what-if?* scenarios to optimise production across its business.

Best Regards,

Gail Purvis

European Editor, Power Systems Design
Gail.Purvis@powersystemsdesign.com

Emerson Network Power triple-output 45-W supplies carry safety approvals

Emerson Network Power is offering two new triple-output open-frame AC-DC power supplies, the NPT43-M and the NPT44-M. Extending Emerson's successful NPT40-M series, these new models are suitable for a wide range of IT equipment, medical, light industrial, instrumentation, and process systems, as well as low-power dental and laboratory equipment.

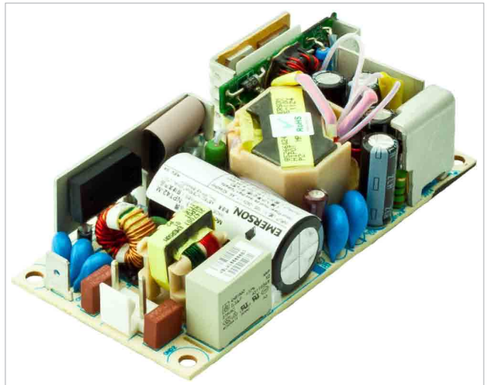

The NPT43-M and NPT44-M triple-output 45-W power supplies carry comprehensive worldwide safety approvals for IT equipment and non-patient-contact and non-patient-critical medical applications.

NPT40-M series power supplies are rated for 45 W output with convection cooling and up to 55 W with forced-air cooling. The NPT43-M offers 5, 15, and –15 V DC regulated outputs while the NPT44-M offers 5, 12, and 24 V DC regulated outputs.

The expanded NPT40-M series carries a comprehensive set of worldwide safety approvals for IT equipment and non-patient-contact and non-patient-critical medical applications. These power supplies can accommodate operating temperatures from -20 to +50 °C at full power and as high as 80 °C with de-rating. They feature an industry standard 2 x 4 inch (51 x 102 mm) footprint, and a height of less than 1 inch (25 mm), making them well suited for 1U-high portable and rack-mounting equipment.

Emerson Network Power's NPT40-M series has a wide-range universal input capable of accommodating any AC voltage in the range of 90 to 264 V AC. It can also operate from any DC input in the range 127 to 300 V DC, enabling it to be used virtually anywhere in the world. The NPT42-M power supply requires less than 74 W of input power, and inrush current is less than 50 A peak at 230 V AC input.

Emerson Network Power's NPT40-M series power supplies fully comply with the international EN 61000-3-2 standard for harmonic emissions. They feature built-in EMI filters (CISPR 22 Class B) and meet rigorous international EMC standards, including FCC Class B, EN 55022 class B, and VDE 0878PT3 Class B for conducted noise. Key safety design aspects include fuses in both the line and neutral connections, and a safety-ground leakage current which does not exceed 275 µA. Safety approvals include TUV/UL/CSA 60950 and 60601-1, CB certificate, CE mark (LVD) and CQC mark.

All three outputs on Emerson Network Power's NPT40-M series power supplies feature short-circuit protection. The main output is also protected against over-voltage conditions, and primary-side total power monitoring protects the overall power supply against overload. The power supply delivers a high demonstrated MTBF of 550,000 hours at full load and under 25 °C ambient operating conditions.

www.emersonnetworkpower.com

Efficiently powering an always-on world

By: John Collins, Global Segment Director, Data Centers, Eaton

A data center's top priority has always been maintaining uptime. From this perspective, it's understandable that efficient energy use has historically been an oversight.

However, the entire industry has become aware that this was changing due to federal mandates, social pressure, and a global energy crisis that has caused the cost of energy to rise dramatically. In fact, energy related costs can represent up to 1/3 of operating expenditures for data centers, which makes efficiency attractive from both competitive and societal perspectives.

The key to balancing business and energy requirements is not to sacrifice reliability in the name of efficiency, but to become smarter about how we use power and collaboratively share successful techniques as they evolve. For example, common data-center efficiency losses are a direct result of computing, storage, and networking underutilization. However, much innovation has gone into minimizing these losses with the most remarkable innovation being virtualization. Virtualization allows the utilization of servers to increase dramatically by decoupling IT hardware and software. One physical server now supports multiple operating systems running multiple applications. This allows data-center operators to reduce the number of servers and for processor power to match dynamically the varying demands of application workloads across a facility.

From a power perspective, what can we learn from the success of this computing model? The answer is the intelligent provisioning of power to the hardware that needs it, at the precise moment it needs it, rather than the unnecessary continuous feeding of energy into non-critical devices.

This approach is a modular power-infrastructure design. The core of this design, an intelligent central-software system, continuously monitors power requirements and ensures that only the modules necessary to supply the load at any particular time are delivering energy. The remaining modules hold in a low-power ready state, and the system is ready to react instantly to changes in load level to meet demand.

There are two primary benefits to this approach: The modules that are in service operate efficiently because they are highly loaded and the modules that are in standby consume very little power. As a result, large energy savings are possible without compromises in service quality.

It is a system that is at the heart of some of the most energy-efficient data centers operating in the world today. For example, Iliad Datacenter, one of France's largest co-location facilities, applied a modular system combined with an intelligent energy management system to boost energy efficiency at one of its sites from 92 percent to 98 percent.

If we are to continue meeting the needs of our always-on society without sacrificing the environment, we must take an intelligent, deliberate, and collaborative approach to managing power on an everyday basis.

www.eaton.com

Drive to energy efficiency continues to provide opportunities

By: Ryan Sanderson, Associate Director, IHS IMS Research

Most will agree that 2012 proved to be a difficult year for the power electronics industry. Floods in Thailand resulted in a shortage of hard-disk drives, restricting PC production in the first half of the year. Global economic confidence worsened owing to increased uncertainty surrounding the *Euro-Zone Crisis*, and many semiconductor markets plummeted despite already low inventory levels.

Despite this, some applications offered pockets of growth. Although market conditions remain unstable, this is projected to continue in 2013 and beyond. In nearly all of these, the common driver is the demand for higher efficiency, driven by either consumer needs or legislation. Several examples spring to mind:

Demand for LED lamps and luminaires continues to grow, driving opportunities for power-supply and -semiconductor manufacturers. A key driver is the LED lamp's much lower power need to provide luminance comparable to incandescent bulbs. It is estimated that lighting currently accounts for approximately 19% of the world's energy use. IHS IMS Research predicts that, in 2016, around 15% of all lighting will use LEDs, reducing the global energy consumption for lighting by around 20%—a saving of $100 billion over five years.

Power-supply efficiency in general is now a crucial factor for any OEM/ODM bringing a product to market, despite the majority of power supplies themselves being supplied by merchant vendors. The market for semiconductors in merchant power supplies is forecast to grow by more than $500 million from 2012 to 2016, driven by various methods to increase efficiency. One is demand for active PFC, including approaches which use SiC diodes. Another is the adoption of synchronous rectification at the output which is driving demand for additional MOSFETs. A third is the adoption of digital power techniques and components.

The overall market for digital-power components—power supplies and ICs—has grown rapidly over the last few years but this has been mostly accounted for by telecom and datacom applications. However, the situation is set to change with design wins in lower cost applications such as notebook and adoption by merchant power supply vendors. This, along with the existing OEM adoption is projected to drive growth in a number of applications over the next five years.

As well as generating, distributing, and monitoring power, it is also predicted that there will developments in the way we store energy efficiently over the coming years. Transitions from battery technologies such as lead-acid to lithium or other alternatives has already begun and is predicted to accelerate, particularly in applications such as UPS, energy storage for reserve power, and motive power applications such as electric vehicles and material handling equipment.

These few examples demonstrate just some of the recent developments in the power industry relating to energy usage. The full list is much larger and should provide many new opportunities in 2013 and beyond. Happy New Year!

www.imsresearch.com

Flyback power supply development: part 1

By: Dr. Ray Ridley, President, Ridley Engineering

This article is the first of a series in which Dr. Ridley shows the steps involved in designing and building an offline flyback converter. The first part of the series presents the power supply architecture and schematic, including the control and bias circuits.

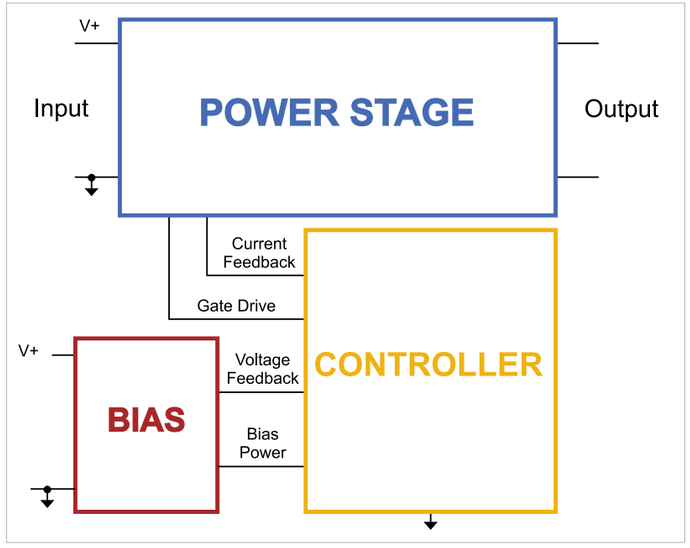

Figure 1: Three parts of a flyback converter—bias, control, and power stage with interconnect signals

Flyback power converter

The offline flyback converter is a topology that is found in many applications. It serves as a bias supply for other power converter topologies, or as the main supply for low-power electronics systems, typically 50 W or less. The flyback is usually selected for the following reasons:

1. Low parts count
2. Single magnetic element
3. Single ground-referenced switch
4. Ease of generation of multiple outputs

The disadvantages of the flyback converter are pulsating input and output currents, and a voltage stress on the switch that exceeds the input voltage. At low power levels, these issues are less important than the flyback advantages.

Power supply requirements

The specifications for the power supply were as follows:

1. Input Voltage: 185 to 265 AC
2. Output 1: 15 V DC @ 1.4 A isolated ±10%
3. Output 2: 15 V DC @ 50 mA nonisolated (bias and regulation supply)
4. Maximum power: 22 W
5. Power Topology: Flyback
6. Controller: 384x controller from Texas Instruments

The main output provides a relatively constant current, and the regulation range of 10% should be achievable through cross-regulation to the bias winding supply.

Block diagram and circuit schematic

Figure 1 shows the block diagram of the flyback power converter with the control, bias, and power stage blocks. Both the power stage and the bias circuit are fed directly from the rectified AC input line. The function of the bias supply is

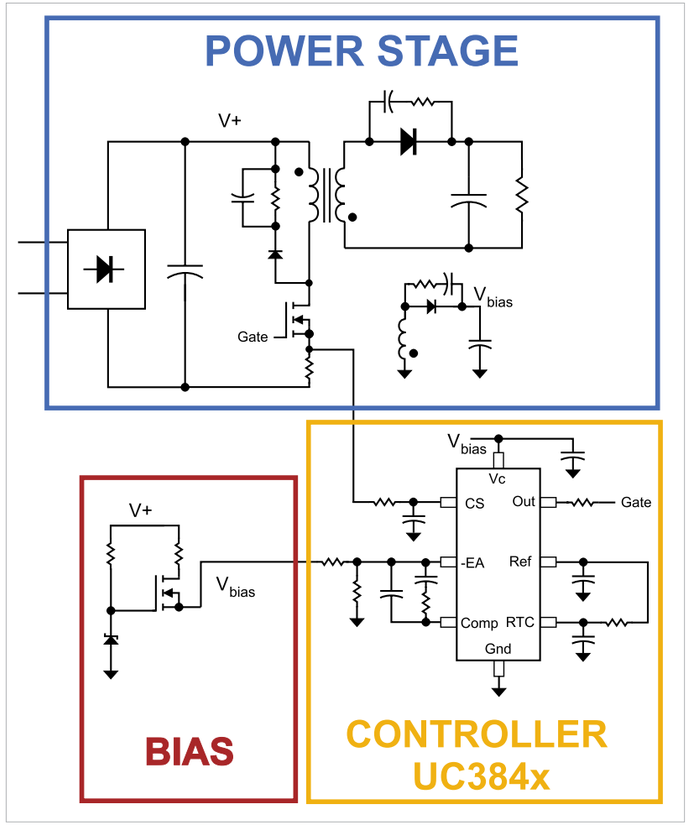

Figure 2: Schematic of an offline flyback converter with control and bias circuits

to provide the operating voltage to the control circuit, and to do this with minimal dissipation.

The controller sends a gate drive signal to the power stage, and receives a current-sense signal from the power stage, used to protect and regulate the converter. Voltage feedback for the control loop is derived from the bias circuit, so there is no need for isolated feedback from the power stage. The choice of the correct controller from the UC384x family of parts will be discussed in a later part of this series.

Figure 2 shows the complete schematic of the flyback power-supply system. In this series of articles, each of the parts of this circuit will be discussed in detail, including test procedures.

Bias Circuit

The bias circuit shown in figure 2 consists of just a few parts. A FET is used as a linear regulator to start the control chip and its gate is fed by a resistor and zener diode combination.

A secondary winding is added to the power transformer to provide the bias power once the main circuit reaches regulation, and the linear regulator is then automatically shut off. This simple circuit provides a rugged and efficient way to power the controller.

Control Circuit

The 384x series of current-mode controllers were selected since they are readily available from multiple sources and are low cost with high performance (**reference 1**). While this family of controllers may seem outdated, they are still used extensively by experienced power designers and power supply companies who wish to have full control over all of the waveforms of the circuit. This will be discussed in detail later in this series of articles.

Power Stage

The power stage shown in figure 2 includes a power FET with voltage clamp, current sensor, transformer, output diode with snubber, and output capacitor. A second output winding on the transformer, together with another diode and capacitor, feeds the bias circuit voltage.

Having a discrete FET design allows much greater design flexibility than working with integrated controllers where the power FET is on the same die as the control circuitry. Also, higher voltage parts are available in the discrete FET package, and this can

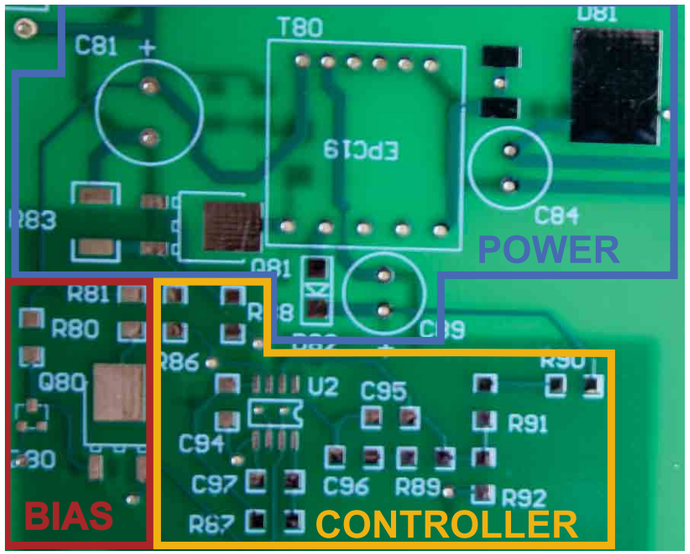

Figure 3: PC board layout of the flyback converter ready for parts population and testing.

be important for rugged design.

Printed Circuit Board

This may not be welcome news to those who are about to embark on their first power supply design, but the first thing that needs to be done is a board layout. It is not a good use of time to try and test controllers and power stage parts with a wired breadboard. There are many critical layout paths in the circuit, and proper placement of parts can only be achieved with a real printed circuit board.

During a power supply development, even a simple flyback, you can count on needing a minimum of two iterations of the PCB. Most projects will iterate the board at least three times by the time it is ready for manufacturing. It is very important to begin your initial testing on the PCB. The board layout creates parasitics and EMI issues that must be resolved as early in the design cycle as possible.

Figure 3 shows a photograph of the double-sided PCB designed for the flyback power supply. Notice that good partitioning of the control, bias, and power stage are maintained on the PCB, and this helps you work on the board later and to achieve good layout.

The details of the layout process are extremely important, and many of the rules and guidelines for proper layout were discussed in some detail in **reference 2**. If you are new to power supply design, the best thing to do is to dive into using a layout program and to learn as quickly as you can by making some boards and testing them. If you are fortunate, you may have someone to mentor you through this process, but most engineers have to learn board layout the hard way, by themselves. Once you become adept at PCB layout, it becomes the fastest way to build test circuits for your power supplies.

Summary

The architecture and schematic of a complete flyback power supply have been presented in the first part of this series of articles. A standard 384x current-mode controller driving a discrete FET is to be used for the power system to provide maximum flexibility of design at a low cost.

In the next article, the input rectifier design and bias circuit will be discussed and test results for this part of the circuit presented.

www.ridleyengineering.com

References:

1. *Current-mode PWM Controller*, UC384x datasheet, Texas Instruments, June 2007, http://www.ti.com/lit/ds/symlink/uc3842.pdf
2. *Power Supply Development Diary Parts IX-XII*, Ridley Engineering Design Center, Articles 42 through 45, http://www.ridleyengineering.com/index.php/design-center.html (note: These articles give design rules for PCB layout for power circuits.)

Faster switching in inverters

New EPCOS CeraLink capacitors benefit power-inverter's DC-link circuits for fast-switching IGBT modules

By: Wolfgang Dreipelcher, Senior Director, Global Reference Designs , EPCOS

The power-semiconductor switches used in power supplies and inverters are based on two technologies, namely MOSFETs and IGBTs. MOSFETs can operate at relatively high switching frequencies significantly above 30 kHz but, unlike IGBTs, they have a very large chip surface area.

A new generation of IGBT modules from Infineon Technologies operates at frequencies of up to 100 kHz. Line-based and turn-off losses are of about the same magnitude in both systems. Fast IGBTs, whose manufacturing complexity is significantly lower and chip areas are often smaller than superjunction MOSFETs, are the basis of an IGBT3 technology with high switching frequencies and an excellent price-performance ratio.

Fast-switching systems require circuit designs with minimized ESR and ESL values. Accordingly, the passive components—inductors, but especially capacitors—must also keep pace with the high switching frequencies. These in turn permit the use of more compact and lightweight passive components, producing lower losses and increasing efficiency.

Capacitors are at the focus of these developments. They must combine high switching frequencies with low ESL and ESR values as well as an extremely compact design.

Conventional capacitor technologies are only partially able to meet these requirements. The EPCOS CeraLink represents a completely new approach: This advanced component is a ceramic multilayer ripple-current suppressor, also known as a link circuit or DC link capacitor; moreover, it functions as a snubber.

New inverter designs possible

EPCOS developed the CeraLink in the company's Competence Center for Ceramic Components in Deutschlandsberg, Austria. The expertise that formed the basis for this advanced component was gained, among others, over many years in the volume manufacture of piezo actuators.

The EPCOS CeraLink offers the advantages of ceramic capacitors without their unfavorable characteristics. The patented multilayer component—based on antiferroelectric ceramic material with special copper internal electrodes—allows the use even more economically of both standard IGBTs and the new high-speed types with significantly higher switching frequencies. This naturally also applies to circuits with corresponding superjunction MOSFETs.

The innovative CeraLink combines high capacitance per volume with low ESL and ESR values, which thus enables significant improvements in efficiency, reliability, and space requirement for future IGBT and MOSFET inverter designs. Moreover, CeraLink is also available in a low-profile SMD design, making it suitable as a snubber for integration in power modules.

The EPCOS CeraLink components were adapted and continuously optimized for the first designs of special IGBT modules in close cooperation with Infineon Technologies, the market leader in IGBTs. This allowed the best results to be

Optimized parameters	
Insulation resistance	Typical values of 1 to 10 GΩ resulting in low leakage current, especially at high temperatures
ESL	<4 nH
ESR	Typically <4 mΩ. Extremely low even at low capacitances, resulting in low losses
Operating temperature	–40 to +125 °C (for short periods up to +150 °C), thus also suitable for SiC
Design benefits	
Internal copper electrodes	Low losses and extremely high current–handling capability
Internal busbar	Optimized for variable use
Various terminal configurations	Terminals for soldering and press–fit assembly technology
Compact case design	Case height optimized for widespread semiconductor modules
Rugged design	Designed for snubber and power applications in industrial and automotive systems
Compatibility	Special types for integration in power modules based on Si IGBTs and MOSFETs or SiC devices
Further benefits	
Suited for rapid rise times and high switching frequencies	
Positive DC bias effect on the capacitance	
Active cooling not always necessary	
Easy traceability thanks to QR codes	

Table 1: Optimized parameters and properties of the EPCOS CeraLink

achieved in terms of performance and energy efficiency. For both the EASY automotive series from Infineon Technologies and the corresponding types in industrial applications, all the relevant capacitor parameters and properties were optimized towards more economy and efficiency (**table 1**).

The first designs of an on-board inverter, the Infineon EASYKIT DCDC, were based on existing OEM specifications for rated voltages of about 400 V DC on the high-voltage side. The EPCOS CeraLink is currently available in several designs. Its capacitance range extends from 1 µF to 100 µF at rated voltages of 400 and 800 V DC.

CeraLink capacitors are available with various terminal. The SMD versions (LP and SMD) are designed for direct integration in semiconductor power modules in view of the restricted space available (**table 2**). They can be soldered, bonded, or sintered.

In close cooperation with EPCOS, Infineon Technologies has developed an HV/LV DC-DC demo board with an output of 2.7 kW. The requirements included a high-voltage input range from 200 to 400 V DC—depending on the HV battery used—and a low-voltage output range from 8 to 16 V DC that is typically standard in automobile electronics systems. Furthermore, the demo board had to cover a current range of up to 200 A DC.

Terminal	Low Profile (LP)	SMD	Soldering pin (SP)	Press–fit busbar (PFBB)
Package	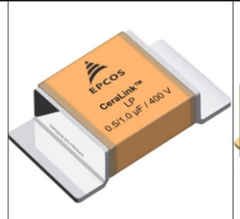	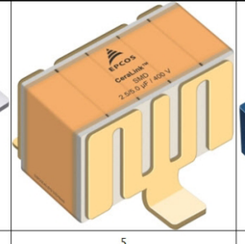		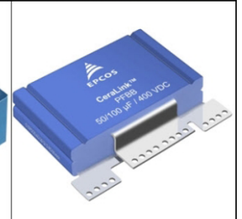
Capacitance [µF]	1	5	5 / 20	100
Rated voltage [V DC]	400	400	800 / 400	400
Dimensions [mm] (without terminals)	6.84 x 7.85 x 2.65	12.8 x 8.4 x 8.8	33 x 22 x 11.5	52.5 x 30.5 x 10.5

Table 2: Various design versions of the EPCOS CeraLink

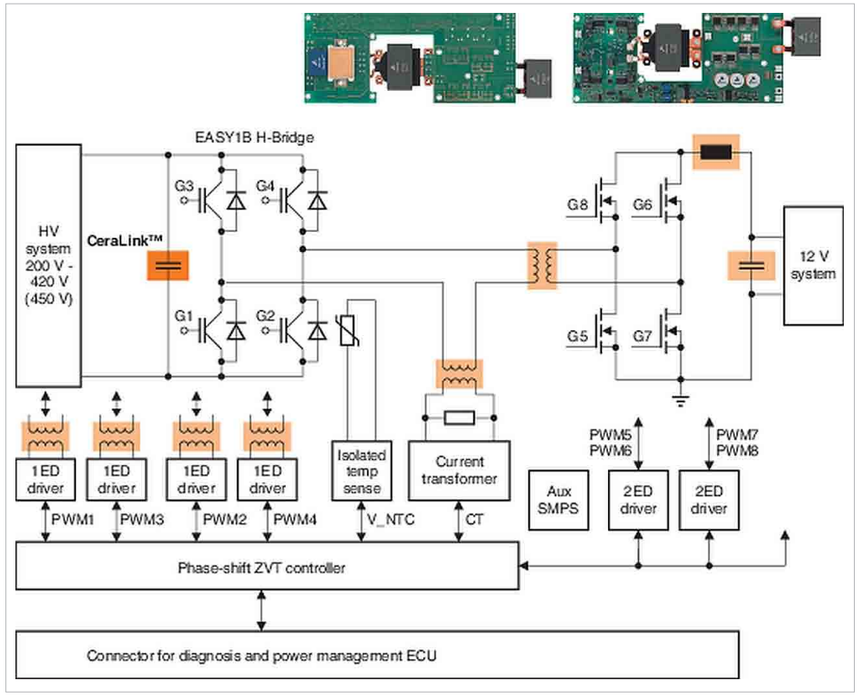

Figure 1: Infineon EASY 2.7 kW inverter

Nearly 100 EPCOS and TDK components

Various types of circuit topologies are available on the DC-DC converter market. However, the most widely used is the full-bridge circuit with a ZVT (zero-voltage transition) based on MOSFET transistors.

Infineon Technologies has redesigned these circuits with various EPCOS components and adapted them to its EASY series of fast-switching IGBTs (**figure 1**). They employ a large number of diverse EPCOS and TDK components (**Table 3**).

Component		Quantity
CeraLink with 20 µF for 400 V DC	EPCOS	1
Aluminum electrolytic capacitors	EPCOS	3
MLCCs	TDK	80
SMT power inductors	EPCOS	7
PCEM T7921 power choke from the Electromobility Platform with peak current of 225 A	EPCOS	1
PTEM T6973 power transformer from the Electromobility Platform	EPCOS	1
GTEM T7509 gate drive transformer from the Electromobility Platform	EPCOS	4
CTEM T7078 current sense transformer from the Electromobility Platform	EPCOS	1

Table 3: EPCOS and TDK components for fast-switching IGBTs from Infineon Technologies

The EPCOS CeraLink combines high capacitance per volume, low ESL and ESR values, and a minimum leakage current, and thus satisfies all the requirements of high-speed IGBT modules or MOSFETs. This system configuration also permits control of high current change rates (di/dt) of up to 10 kA/µs. Despite these extremely high potential rates, the generated voltage peaks (V = Ldi/dt) are extremely low thanks to the low ESL of the CeraLink.

The capacitor is not the only source of parasitic inductance. Noticeable stray inductances occur in a normal system configuration for several reasons, including the contacting inside the IGBT module and the feed

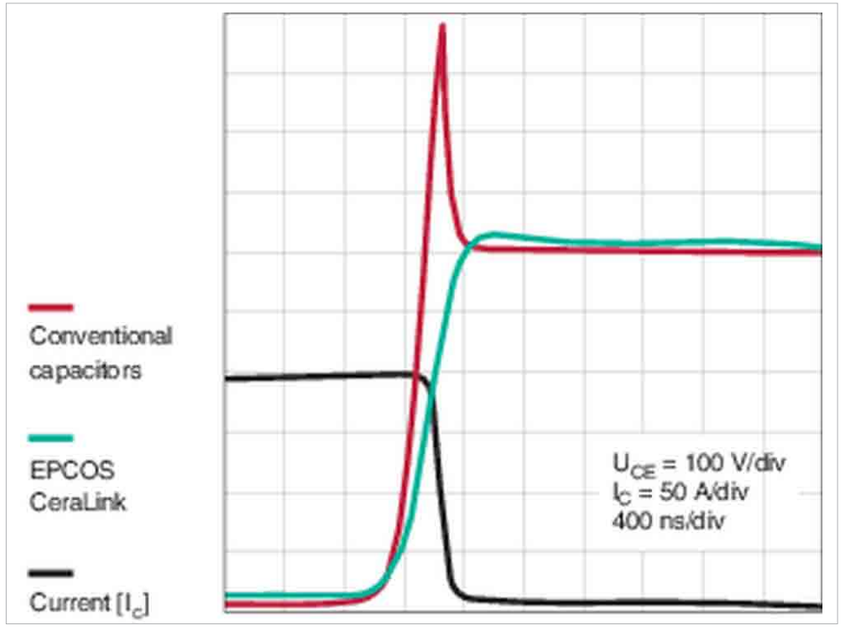

Figure 2: Voltage across the IGBT due to parasitic inductances when switching

line to the capacitor. The EPCOS CeraLink allows for dramatic reduction in the parasitic-inductance values for the feed line—to the same extent as the values for the capacitor itself, thanks to its compact design.

The compact link to the IGBT module simultaneously attenuates its over voltages, and a snubber capacitor is usually not necessary. **Figure 2** shows the voltage curve at turn-off of the IGBTs with and without an EPCOS CeraLink. The voltage rise is thus only minimal and is within the safe range for the IGBTs. In this case, the switching frequency is 100 kHz, meaning the capacitor sees a ripple-current frequency of 200 kHz. **Figure 3** shows the impedance and ESR curves as a function of the frequency.

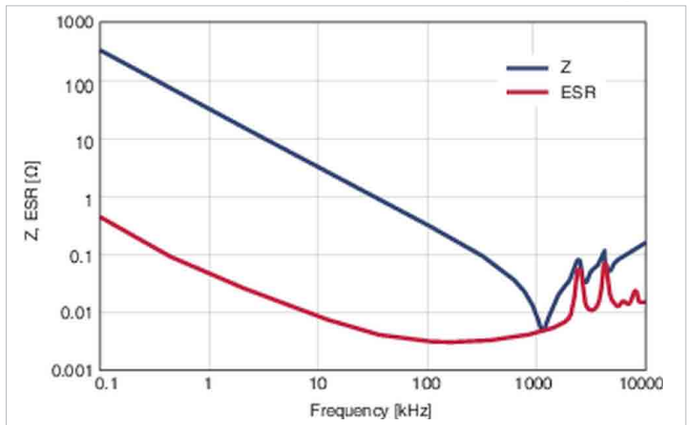

Figure 3: Impedance and ESR as a function of the frequency

Thanks to its very low ESR values, the EPCOS CeraLink attenuates the overvoltage peaks very effectively. As a rule, therefore, additional snubber capacitors are not necessary.

Although the capacitance of the EPCOS CeraLink is usually sufficient for pure DC-DC applications, it may be too low for motor operation, for example. This can be remedied by connecting aluminum electrolytic or film capacitors in parallel, as their high capacitance carries the low-frequency current component. The EPCOS CeraLink then handles the high-frequency component, including the snubber component.

www.epcos.com

90-W plug-n-play Power over Ethernet over 100-m CAT-5e cable

Proprietary PoE extension is fully backward compatible

By: Alison Steer, Product Marketing Manager, Mixed-Signal Products, Linear Technology

PoE (Power over Ethernet) has been increasing in popularity due to its flexible and cost-effective method of delivering both power and data over a single Ethernet cable. This allows equipment to be installed almost anywhere without the constraint of AC-power proximity or requiring installation by an electrician.

The original IEEE 802.3af PoE specification limited the power delivered to the powered device to just 13 W, which in turn limited the scope of applications to devices such as IP phones and basic security cameras. In 2009, the IEEE 802.3at specification increased this available power to 25.5 W. However, this was still not enough to satisfy the growing number of power hungry PoE applications, such as picocells, wireless access points, LED signage, and heated PTZ (pan-tilt-zoom) outdoor cameras.

In 2011 Linear Technology released a new proprietary standard, LTPoE++, which extends the PoE and PoE+ specifications to 90 W of delivered power, while maintaining 100% interoperability with the IEEE PoE standards. LTPoE++ provides a safe and robust plug-n-play framework that dramatically reduces engineering complexity in PSEs (power-sourcing equipment) and PDs (powered devices). The benefit of LTPoE++ over other power-extending topologies is that only a single PSE and PD is required to deliver up to 90W over 4 pair CAT5e cable, resulting in significant space, cost, and development

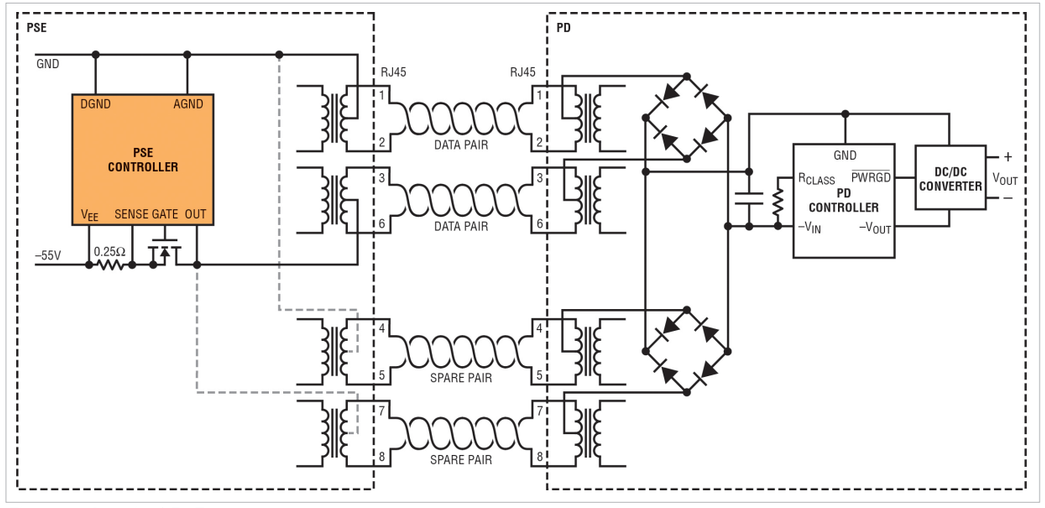

Figure 1: A typical PoE system

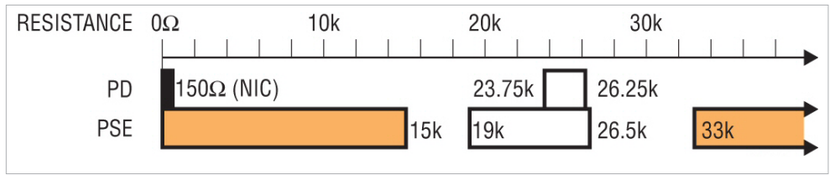

Figure 2: IEEE 802.3af signature-resistance ranges

time advantages (**figure 1**). Four different power levels are available—38.7, 52.7, 70, and 90 W—allowing system designers to size the power supply to the application's requirements.

Powered-device detection

Before a PSE can apply power to the line, it must check for an IEEE-mandated signature resistance with a power-limited probing source. To be considered a valid signature, the PD must look like 25 kΩ ±5% in parallel with a capacitance of 120 nF or less. The PSE, in turn, must accept a somewhat wider range of 19 kΩ to 26.5 kΩ to account for parasitic series and parallel resistances in the system (**figure 2**). The PSE must reject anything below 15 kΩ or above 33 kΩ, or anything with >10 µF across its terminals.

The PD signature impedance is allowed to have a voltage offset of up to 1.9 V—typically caused by up to two diodes in series—and a current offset of up to 10 µA—typically caused by leakage in the PD. These terms complicate the PSE resistance measurement, since a single V-I point measurement will not account for these errors. As a result, the PSE is required to take at least two different V-I points, separated by at least 1 V at the PD. It then must calculate the difference between the two points to find the true resistive slope, subtracting out voltage and current offsets.

Because CAT-5 cable typically runs in ceilings, walls, and other spaces where AC wiring is also present, 50 or 60 Hz noise can be significant. Linear Technology PSE controllers handle this by using a proprietary dual-mode, four-point detection method, which ensures immunity from false positive or negative PD detection.

Powered-device classification

Once the PSE has successfully detected a PD, it performs the power-classification step. The PSE must keep track of how many PDs are connected and what their power-classification levels are, and stop accepting PDs when its power budget is exhausted. The PSE checks a PD's classification signature by forcing between 14.5 and 20 V across the PD and measuring the current that the PD draws. The PSE uses the measured current to classify the PD.

LTPoE++ uses a three-event classification scheme to provide mutual identification handshaking between the PSE and PD while maintaining backward compatibility with the IEEE 802.3at standard. The LTPoE++ PSE determines if a PD is a Type 1 (PoE), Type 2 (PoE+), or LTPoE++ device by the PD response to the three-event classification scheme. The LTPoE++ PSE also uses the three-event classification scheme result to update the ICUT and ILIM thresholds.

On the other end, the LTPoE++ PD uses the number of classification events it receives to determine whether it is receiving power from a Type 1, Type 2, or LTPoE++ PSE. If the LTPoE++ PSE measures the PD's first classification event current as a Class 0, 1, 2, or 3 device, the LTPoE++ PSE will proceed to power on the port as a Type 1 device.

Otherwise, if the PSE detects a Class 4 PD during the first classification event, the LTPoE++ PSE will continue with a second classification event, as defined in the PoE+ specification. This informs the PD that it is receiving power from either a Type 2 or LTPoE++ PSE. The absence of the second classification event indicates the PD is receiving power from a Type 1 PSE that is limited to Type 1 power levels.

PD input power	Class pulse		
	1st Event	2nd Event	3rd Event
13	0	–	–
4	1	–	–
7	2	–	–
13	3	–	–
Invalid	4	0 to 3	–
25.5	4	4	4
38.7	4	4	0
52.7	4	4	1
70	4	4	2
90	4	4	3
Invalid	Overcurrent*	–	–
Invalid	4	Overcurrent*	–
Invalid	4	4	Overcurrent*
*Class current I_{CLASS} exceeds that specified for an overcurrent			

Table 1

The Type 2 PD physical layer classification is defined by IEEE as two consecutive Class 4 results. The LTPoE++ PD must also display two consecutive Class 4 results in the first and second classification events.

The LTPoE++ PSE will move onto the third classification event after valid Class 4 measurements in the first and second classification events. The third classification event must result in a classification other than Class 4 to recognize the PD as LTPoE++ capable. An LTPoE++ PSE will consider a PD that maintains Class 4 during the third classification event as a Type 2 PD. The IEEE 802.3at standard requires compliant Type 2 PDs to repeat Class 4 responses for all class events.

The third classification event informs the LTPoE++ PD that it is receiving power from an LTPoE++ PSE. **Table 1** shows the class-event permutations for the various PD power levels. The LTPoE++ PD presents a Class-0 through -3 classification current during the third classification event. The four different classes indicate to the LTPoE++ PSE the maximum power the LTPoE++ PD is requesting at its input. The four LTPoE++ power levels of 38.7, 52.7, 70, and 90 W at the LTPoE++ PD input correspond to the four classes, Class 0, Class 1, Class 2, and Class 3.

DC disconnect

Just as a PSE must only send power to valid PD, a PSE also must not leave power on after the powered device has been unplugged because a powered cable could subsequently connect to a device that doesn't expect power. LTPoE++ uses the DC disconnect method to determine the absence of a PD based on the amount of DC current flowing from the PSE to the PD. When the current stays below a threshold IMIN—between 5 and 10 mA—for a given time tDIS—300 to 400ms—the PSE assumes that the PD is absent and turns off the power.

Putting it all together

Once a PSE has successfully detected and classified a PD, it then makes the decision whether to supply power to it. If the PSE's available power is adequate to power the PD, the PSE powers on the PD and begins monitoring the port for the DC disconnect condition. The PSE now has the whole picture: The detection sequence tells it that there is a real PD attached to the port. The classification routine tells it how much power that PD will draw so it can allocate its power supply resources accordingly. Finally, the DC disconnect method tells it that the PD is still present and operating normally. The PD, in turn, has a straightforward way to communicate to the PSE what it is, how much power it wants, and whether or not it wants that power to keep flowing. All of this occurs without affecting the data stream in any way.

One important distinction with LTPoE++ is that it does not require the use of the LLDP (link-layer discovery protocol) that the IEEE mandated in its PoE+ specification for software-level power negotiation. LLDP requires extensions to standard Ethernet stacks and can represent a

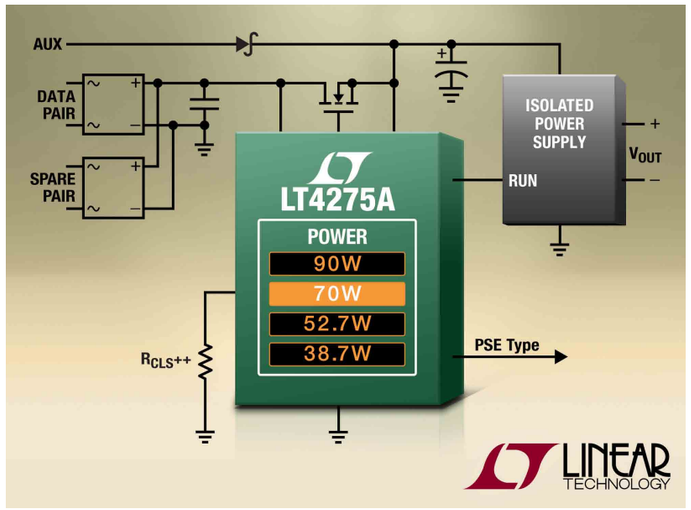

Figure 3: The LT4275 90-W PD controller uses an external MOSFET for increased power efficiency.

significant software development effort. LTPoE++ PSEs and PDs autonomously negotiate power level requirements and capabilities at the hardware level while remaining fully compatible with LLDP-based implementations. This gives LTPoE++ system designers the choice to implement or not implement LLDP. Proprietary end-to-end systems may choose to forgo LLDP support. This creates time-to-market advantages while further reducing BOM costs, board size, and complexity.

LTPoE++ plug-n-play support

Linear Technology offers single, quad, octal, and 12-port LTPoE++ PSE controllers with low power dissipation, robust ESD and cable-discharge protection, low component count, and cost-effective designs. When paired with the LT4275 PD controller (**figure 3**), a complete plug-n-play LTPoE++ system, with no LLDP required, can deliver up to 90 W while remaining fully compatible with PoE+ and PoE standards. The entire implementation uses external low-$R_{DS(ON)}$ MOSFETs to drastically reduce overall PD heat dissipation and maximize power efficiency, which is important at all power levels. High absolute maximum ratings on all analog pins and cost-effective cable-discharge protection ensure the devices are safe from the most common Ethernet line surges.

www.linear.com

Convert a buck regulator into a high-current LED driver

Coercing a voltage regulator to provide a constant current can extend your drive capability while maintaining low BOM cost.

By: Jon Kraft, Applications Engineer, Analog Devices

LEDs promise to change the world, and few doubt that they will. However, a key limiter to more rapid adoption is the cost of the LEDs themselves. The cost breakdown of LED luminaires varies, but it is safe to put the expense of the LEDs at around 25-40% of the total luminaire cost (**reference 1**). Projections suggest that LEDs will continue to represent a significant fraction of the total luminaire cost for many years (**figure 1**).

One way to reduce the total luminaire cost is to drive an LED at its highest possible current. If driven properly, these LEDs produce more lumens per unit cost. Doing so, however, requires higher current drivers. While there are many drivers available to operate LEDs at low currents (< 500 mA), there are fewer options at higher currents (700 mA to 4.0 A).

This is all the more surprising given that the semiconductor world is rich with DC-DC converters that can source as much as 4.0 A. The problem is that these DC-DC converters control voltages, not LED currents. Fortunately, there are some easy methods to take a readily available DC-DC buck regulator and convert it into an LED driver.

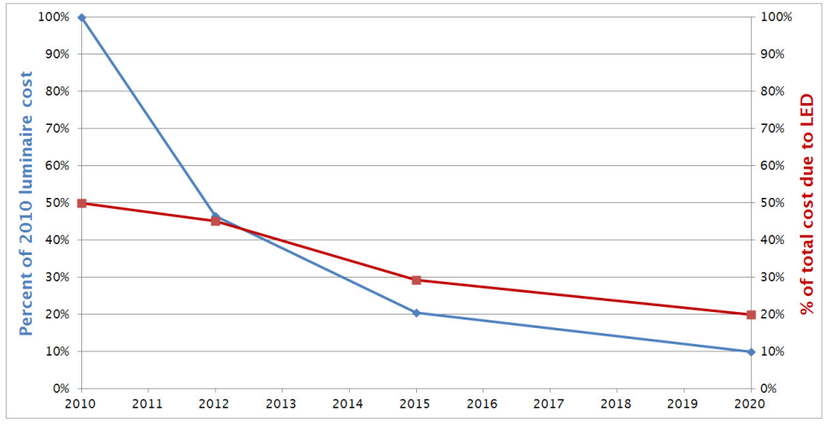

Figure 1: Projected relative cost of LED Luminaire and percentage due to the cost of LEDs (Source: DOE SSL 2011 Manufacturing Roadmap)

A buck regulator simply chops up an input voltage, passes it through an LC filter, and gives a stable output (**figure 2**). To do this, the buck employs two active elements and two passive elements. The active elements are a switch, *A*, from the input to the inductor and a switch, *B*, or diode from ground to the inductor. The passive elements are the inductor, L, and the output capacitor, COUT. These form an L-C filter, which reduces the ripple the active elements create.

If the switches are internal to the IC it's a *regulator*, otherwise it's a *controller*. If both switches are power transistors—MOSFETs or BJTs—then it's *synchronous*, otherwise it's called *asynchronous*. This gives several categories of buck circuits and each has its own merits and drawbacks. A discussion of which type to use, and all the tradeoffs, would be extensive. However, in a very general sense, the selection that often times gives an optimal efficiency, parts count, BOM cost, and board area is a synchronous buck regulator.

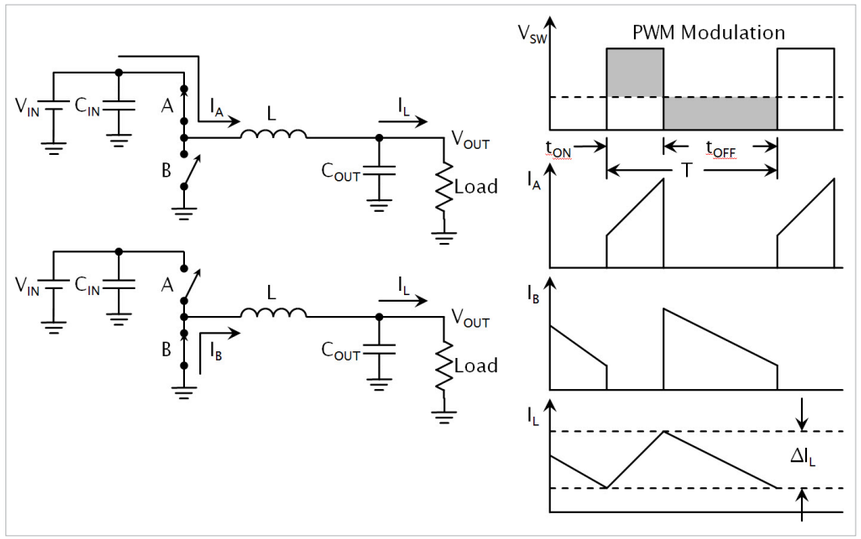

Figure 2: A basic buck-converter topology with waveform timing

However, synchronous buck regulators for driving high current LEDs—up to 4 A—are few and expensive. An alternative is to take a standard synchronous buck regulator and modify it to regulate LED current. Here, we'll use as examples two general-purpose synchronous buck regulators from Analog Devices: the ADP2441 and the ADP2384.

The ADP2441 is a high efficiency, 36-V input synchronous buck regulator, capable of producing an output current of up to 1.2 A. The ADP2384 is another high-efficiency synchronous buck regulator, but with an output current of up to 4.0 A, and an input voltage up to 20 V (**figure 3**).

For both the ADP2441 and the ADP2384, the output voltage is resistor divided down to the FB pin, compared against an internal 600-mV reference, and used to generate the proper duty cycle to the switches. In steady state, this FB pin regulates to exactly

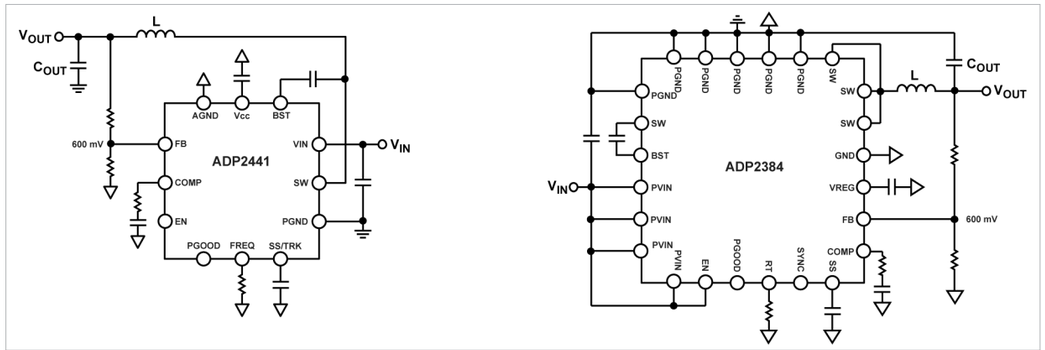

Figure 3: Examples of synchronous buck regulators, the ADP2441 (left) and ADP2384 (right) configured for voltage regulation

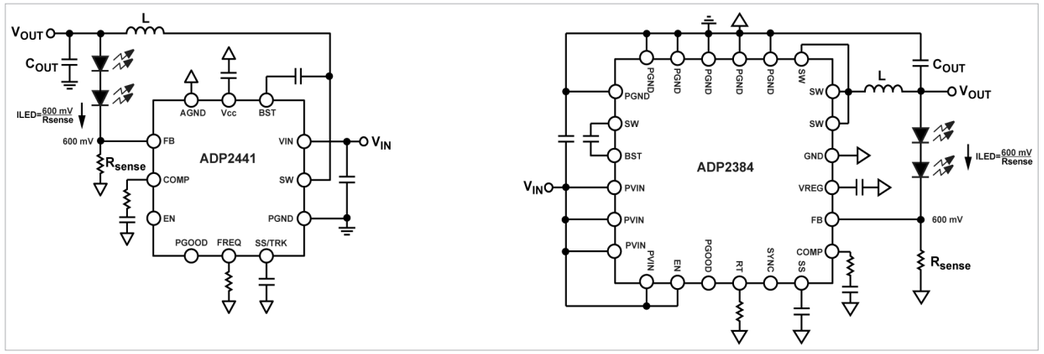

Figure 4: A basic but inefficient LED-driver configuration

600 mV. So instead of a resistor divider, it is easy to put the LEDs there, with a resistor, RSENSE, in series to set the current (**figure 4**).

Using a precision resistor from FB to GND sets the LED current to

$$I_{LED} = \frac{600\ \text{mV}}{R_{SENSE}}$$

This works nicely, but it produces a lot of power dissipation:

$$P_{diss} = 600\ \text{mV} \cdot I_{LED}$$

For low LED currents, the sense-resistor dissipation is not a big issue. But at high LED currents the impact to efficiency adds significantly to the heat that the luminaire must dissipate:

$$600\,\text{mV} \cdot 4\,\text{A} = 2.4\,\text{W}\ (!)$$

Fortunately, there are two methods to reduce the FB reference voltage for most buck regulators: use the SS (soft-start) or TRK (tracking) pin or offset the Rsense voltage.

Many general-purpose buck ICs

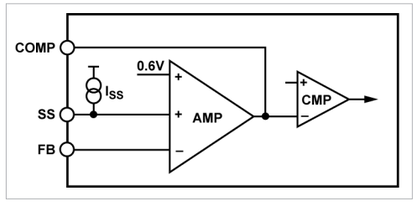

Figure 5: An example of soft-start or tracking pin operation

include an SS or TRK pin. The SS pin provides a controlled inductor current at startup. The TRK pin forces the buck regulator to follow an independent voltage. Often times the pins are combined into one SS/TRK pin. In most cases, the error amplifier will use the minimum of the SS, TRK, and FB_ref voltages to change the regulation point (**figure 5**)

For the current purposes, we want to just set the SS/TRK pin to a fixed voltage and use it as our new FB reference. A resistor divider from a constant voltage to this pin works nicely. Many buck regulator ICs include a controlled low-voltage output, like the VREG pin on the ADP2384 or the VCC pin of the ADP2441. For greater accuracy, use a simple two-terminal precision reference such as the ADR5040, for example. Either way, a resistor divider from the supply to the SS/TRK pin forms the new reference (**figure 6**). Setting the resistor divider to give a SS/TRK voltage around 100 to 200 mV generally offers the best compromise between power dissipation and LED-current accuracy. Another benefit to setting your own feedback reference voltage is that this method can easily accommodate any sense resistor value in the standard-value series. This eliminates the expense and inaccuracy of paralleling multiple RSENSE resistors to set the LED current.

Using an SS or TRK pin accommodates many, but not all buck regulators. Some ICs do not have SS or TRK pins. Additionally, in some buck ICs, the SS pin changes the peak inductor current, not the FB reference so carefully check your

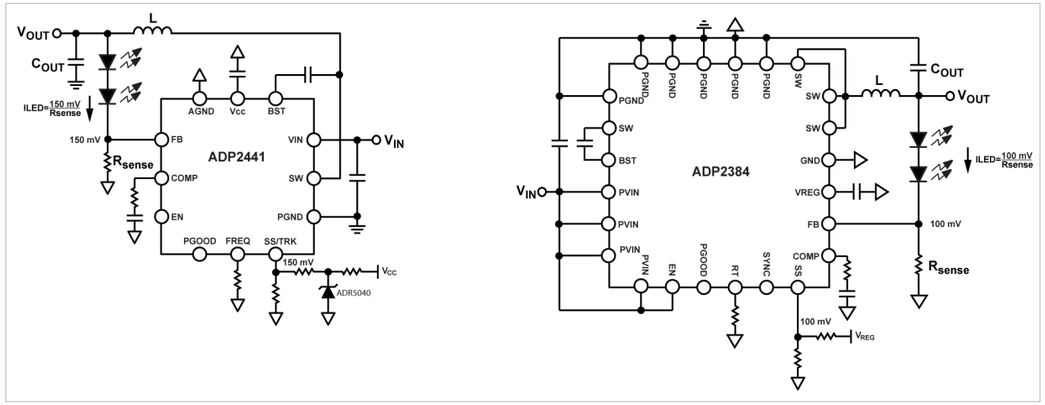

Figure 6: Using the SS/TRK pin to reduce the FB reference voltage

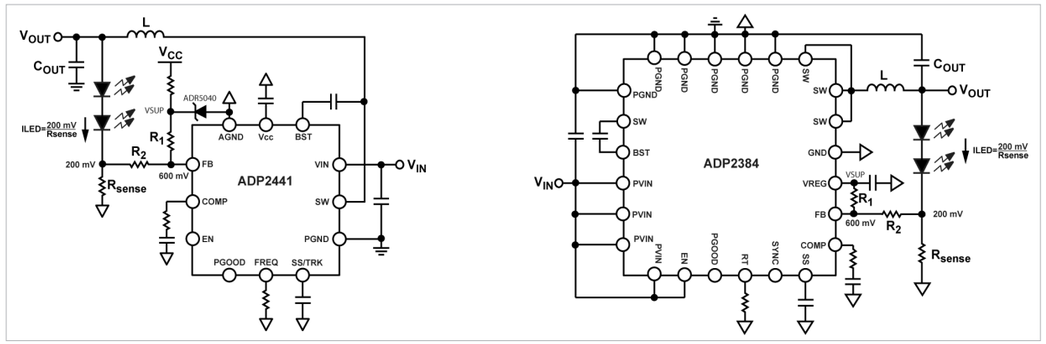

Figure 7: Offsetting the Rsense voltage

buck-regulator's datasheet.

If either of these situations is the case, you can use a different method: offset the RSENSE voltage. A resistor divider tied to an accurate rail between FB and RSENSE provides a fairly constant offset voltage between RSENSE and the FB pin (**figure 7**).

To implement the RSENSE voltage offset, calculate the necessary values for the resistor divider:

$$R1 = R2\frac{V_{sup} - FB_{REF}}{FB_{REF} - FB_{REF(NEW)}}$$

So to get an *effective* feedback reference, FBREF(NEW), of, say, 200 mV with R2 = 1 kΩ and VSUP = 3.0 V,

$$R1 = 1\,\text{k}\Omega\frac{3.0\,\text{V} - 0.6\,\text{V}}{0.6\,\text{V} - 0.2\,\text{V}} = 6.0\,\text{k}\Omega$$

With this arrangement, the LED current is equal to:

$$I_{LED} = \frac{FB_{REF(NEW)}}{R_{SENSE}}$$

Using this method does not require a SS or TRK pin. Additionally, the FB pin will still regulate to 600 mV, but the voltage at RSENSE regulates to the FBREF(NEW) voltage. This means that other functions of the chip, such as soft start, tracking, and power-good indication will still function normally.

The downside is that the accuracy of the supply can strongly influence the offset between RSENSE and FB. A precision reference like the ADR5040 has little problem. But if the supply in the above example had a ±5% accuracy, it would create a ±11% variation in the LED current. A comparison appears in **table 1**.

These tips should be taken as general guidelines for implementing comprehensive LED features into a standard buck regulator. Because

Option 1 Use SS/TRK to reduce FB reference	Option 2 Offset RSENSE Voltage
±5% supply voltage variation gives ±5% error on I_{LED}. This is not affected by the V_{SENSE} voltage, therefore this method has the best Rsense power dissipation.	Supply voltage variation severely degrades LED accuracy. Higher V_{SENSE} voltages improve this, but also increase the R_{SENSE} power dissipation.
P_{GOOD} will always remain low.	Since FB pin still regulates to 600 mV. The P_{GOOD} pin functions normally.
By keeping the SS/TRK pin lower than normal, some fault modes may not work properly.	All fault modes work normally

Table 1: Comparison of two methods to reduce RSENSE power dissipation

these features might be different from the intended application for the buck IC, however, it is always best to contact the semiconductor manufacturer to ensure that the IC can be used as described here.

www.analog.com

References:

1. *SSL 2011 manufacturing roadmap*, US Department of Energy, http://www.ssl.energy.gov
2. Marasco, Ken, *How to apply DC-to-DC step-down (Buck) regulators successfully*, Analog Dialogue, June 2011

Flash corruption: software bug or supply-voltage fault?

Answer: Both!

By: Shyam Chandra, Lattice Semiconductor

Embedded systems commonly use Flash memory to store firmware. Occasionally, the firmware stored in the Flash memory in some systems is accidentally corrupted, preventing the system from booting up after powering on. Flash corruption is commonly associated with a software bug.

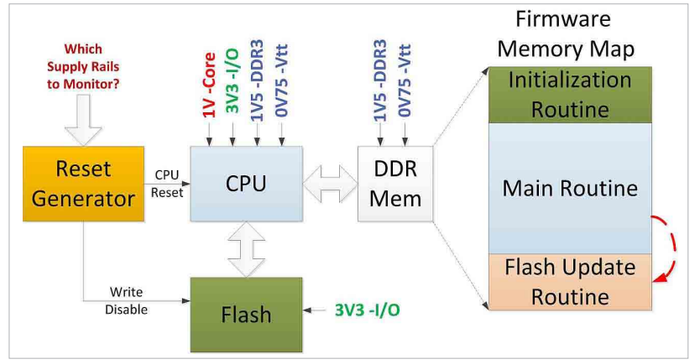

Figure 1: Typical CPU section and firmware memory map

However, it is also commonly understood that the probability of Flash corruption increases either during power-cycling tests or during margining tests. The Flash corruption problem tends to be more severe when the number of complex ASICs or SOCs a board uses increases. Fortunately, methods exist to minimize occurrences of Flash corruption that are *not* due to software faults.

Flash memory corruption

Figure 1 illustrates a typical circuit board's CPU circuitry. When the power turns on, the reset generator first activates the CPU reset signal. It then waits until the power to the CPU, the Flash memory, and the DDR memory each reaches its correct level, waits for an additional extended period of time—about 150 ms—and then deactivates the CPU reset signal. When the reset signal deactivates, the CPU begins to execute the initialization routine in the Flash memory, transfers the contents of the firmware stored in the Flash memory into the DDR memory and then executes the program from the DDR memory.

The procedure to load firmware into the Flash memory is: (1) Firmware is downloaded into the DDR memory through a communication interface. (2) Jump to the *Flash-update routine* to reprogram the Flash with the new firmware. (3) Power to the processor is recycled and the new firmware takes effect.

An event that causes the code execution to jump inadvertently to the Flash-update routine can corrupt the Flash memory contents. When the board power is cycled, the corrupt version of the code is loaded into the DDR and the board does not function as expected.

The code execution could jump to this Flash-update routine inadvertently due either to a software bug or to a faulty supply voltage rail (during the power-off event, for example). Normal debugging methods can detect a software bug. However, a faulty power supply voltage is hard to detect, as the supply voltage

error can occur anywhere.

Supply-fault-induced execution jumps
All ICs have both minimum and maximum operating voltage specifications. If the system exceeds an IC's maximum voltage specification, damage to the device results. If the supply drops below the minimum supply level, the device no longer operates as specified.

For example, the core-voltage specification of the CPU in figure 1 is 1.2 V ±5%. If the voltage drops below this level, the ability of the CPU's internal instruction execution pipeline to reliably transfer instructions and data is compromised and—depending on the CPU's process and operating temperature—the instruction can incorrectly execute.

As an example, a *move* instruction could execute as a *pop* instruction, and the code execution then jumps to a random memory location determined by the contents of the stack. Depending on the contents of that memory location and the error in execution, the processor can either hang or jump to the Flash-update routine, corrupting the Flash memory and overwriting the Flash memory contents.

A droop in DDR memory voltage or threshold voltage also introduces errors in the instructions and data transferred between the memory and CPU. This erroneous code execution can also cause a jump to the Flash-update routine, corrupting Flash memory.

Supply-voltage droop
The power-supply voltage droop can occur for either of the two following reasons.

Card power down: When the power to the board turns off, not all supplies on the board turn off at the same time because the turn off rate depends on a number of parameters including the supply capacity, load, and supply-output capacitance. Because the power-supply turn-off slew rate is slow in comparison to the processor's instruction execution speed, the processor can experience a supply fault, causing it to mis-execute instructions before the supply is fully off or before its reset signal activates.

Momentary ground-voltage rise: The power consumption of some processors can fluctuate dynamically, depending on the executing instructions. When such changes occur, the device draws large amounts of current for brief periods from the power source, and dumps these into the ground. As a result, the supply voltage can momentarily droop and the ground voltage may increase. The duration of such a condition depends on the inductance of the supply path.

Minimizing corruption
The probability of Flash corruption can be minimized by activating the CPU reset when any supply rail drops below its threshold level. This prevents code execution under faulty power supply conditions.

The reset generator activates both the CPU-reset signal as well as the write-protection signal to the Flash memory. In some cases, the reset generator output does not apply directly to the CPU. Instead, it connects to a CPLD, which executes a reset-distribution algorithm. In such cases, the write-protection signal for the Flash should be set because the CPU may not reset as soon as the power-supply voltage becomes faulty. The reset generator IC in figure 1 monitors all CPU rails—1 .0, 3.3, 1.5, and 0.75 V—and activates the reset signal and Flash write-protect signals when any one of them drops below their operating threshold levels.

Selecting a reset generator
The criteria for selecting a reset IC include the number of voltage-monitor inputs, glitch filtering, hysteresis, fault-detection accuracy (across the operating-temperature and –voltage ranges), and fault-detection speed.

Number of voltage-monitor inputs: The reset-generator IC must monitor all CPU voltage rails for faults—voltage excursions below corresponding operating threshold levels. In the case of figure 1, four inputs are required with thresholds set at 5% below the nominal operating voltage levels. For example, Lattice power management ICs support six to 12 voltage-rail-monitor inputs, and the reset generation threshold levels can be programmed from –0.5% to –20%.

Errors to avoid: Some designs use a single-rail reset generator that usually monitors only, for example, 3.3 V. This will not be sufficient, because the 3.3V rail may turn off at a different rate than the core voltage or the DDR voltage.

This arrangement could work only if all critical loads used the 3.3 V as their input supply. In most circuit boards, however, the power supply for the core and DDR use different input-voltage sources due to power dissipation, and so reset generation using only 3.3V cannot avoid Flash corruption. The same argument holds if the reset generator monitors only the core supply rail.

Glitch filtering: When the reset generator has single-ended, as opposed to differential, sensing of voltage rails, differences in the ground voltage between the reset IC and the CPU memory can gen-

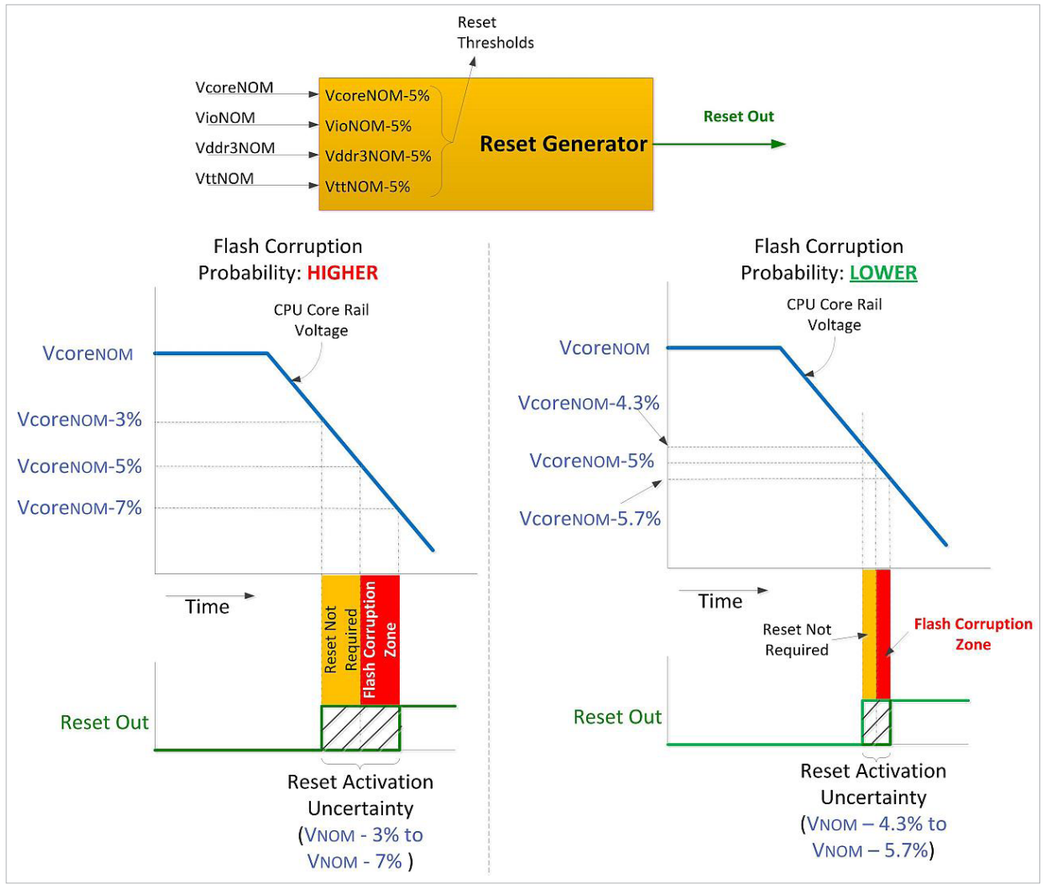

Figure 2: Flash corruption probability vs. reset-generator accuracy

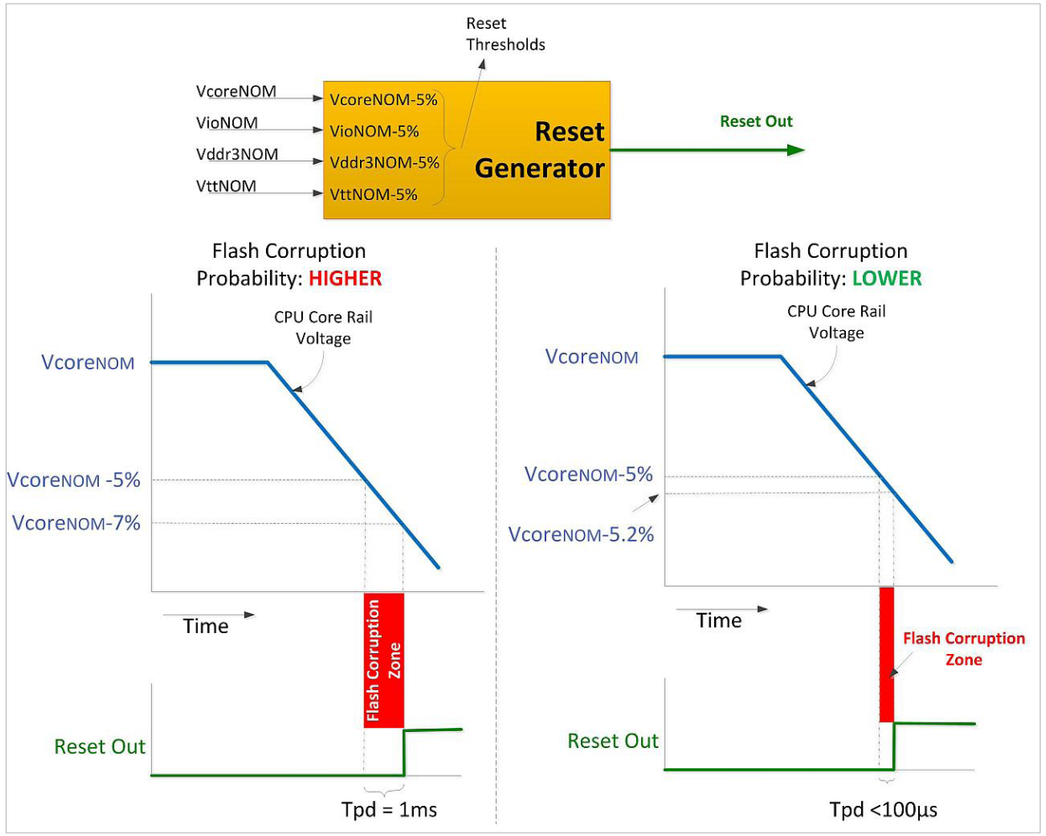

Figure 3: Flash corruption probability vs. reset-generator fault-detection speed

erate false reset signals. To make sure that the reset is actually due to a fault in the supply voltage and not a momentary ground voltage difference, reset ICs include glitch filters. For example, when their input glitch filters are enabled, Lattice power-management ICs ensure that the fault persists for 64 microseconds before activating the reset signal.

Reset generators using ADCs and microcontrollers to monitor voltages implement ADC sample-averaging algorithms to eliminate the effects of glitches, resulting in false reset activation. The averaging algorithm derives the actual ADC voltage by calculating the average of four ADC voltage samples.

Hysteresis: Most voltage rails source from switched-mode power supplies. The output of these supplies usually contain ripple. This ripple can cause a reset signal glitch when the supply level is close to the reset threshold. To avoid this, reset generators must have hysteresis voltage levels ranging from 0.5% to 1% of the voltage monitored.

Reset generators using ADC and microcontrollers to monitor voltages should implement hysteresis in software to prevent glitches in the reset output.

Fault-detection accuracy: For the purpose of this discussion, assume that the lowest operating voltage of the CPU is VNOM – 5%, where VNOM is the nominal core voltage, and that when the supply turns off, VNOM reduces linearly at a rate of 2%/ms.

The accuracy of a reset generator is a measure of uncertainty in its ability to detect a given voltage threshold. For example, a reset generator monitoring a VNOM – 5% threshold with an error of 2% can activate the reset output anywhere between VNOM – 3% to VNOM – 7%. The processor continues to execute instructions until the reset signal activates.

Consider two reset generators with assumed zero propagation delay (**figure 2**): one with an accuracy of 2% (left) and the other with an accuracy of 0.7% (right). As can be seen, the supervisor with 2% error has a much wider uncertainty range than that of the reset generator with 0.7% accuracy.

While the reset-output activation within the orange zone prevents the processor from executing even though the supply is healthy—an irritant—the activation in the red zone is its inability to prevent the processor from corrupting the Flash memory. Clearly, the narrower the reset generator's threshold uncertainty, the lower the probability of Flash corruption. The accuracy of Lattice power management devices is 0.7%.

Some designs use the *power-good* signal from DC-DC converters to determine the health of supplies and use a CPLD to generate reset signals. This method does not reduce the probability of Flash corruption because the accuracy of the power good signals from DC-DC converters ranges from 4% to 20%.

Also, some designs use low-cost comparators to monitor the voltages. In this case, one has to pay attention to voltage reference, resistor accuracy, and comparator offset errors. For example, for a voltage-monitoring circuit to maintain 1% accuracy across voltage and temperature, one has to use a comparator with < 1-mV offset error, a VREF with an accuracy < 0.5%, and 0.1% resistors to set the fault-detection threshold.

Fault detection speed (Tpd): For the purpose of this discussion, assume that the voltage-monitoring error of the reset generator is 0%.

Fault detection speed is a measure of the time required for the reset generator to activate the reset-output signal from the time the voltage crosses the fault threshold or the reset generator's fault propagation-time delay (**figure 3**). In the left side of the figure, the reset generator requires 1 ms to activate the reset signal. The voltage continues to droop and, by the time the reset signal is active, the supply voltage at the CPU is 7% below its nominal operating voltage. This allows about 1 ms for the CPU to corrupt the Flash.

When the fault detection speed is less than 100 µs, the voltage at the CPU is VNOM – 0.2%, and the probability of Flash corruption is exponentially less. Lattice power-management devices are able to activate the reset signal in about 64 µs with the glitch filter turned on, or 16 µs with the glitch filter turned off.

Some designs use a microcontroller with an ADC as a reset generator. In such arrangements, the voltage-monitoring routine activates the reset signal, based on a 10 to 50 ms real-time-clock interrupt. Consequently, the reset can activate with a delay of 10 to 50 ms. Because of this long delay, this method will not be able to prevent Flash corruption. Note that an ADC's errors and the error associated with the on-chip ADC voltage reference determine the voltage monitoring accuracy of an ADC. The number of ADC bits is not a measure of its accuracy.

Summary

The conventional thinking—that Flash corruption is due only to a software bug—results in engineers wasting time looking for one that does not exist. Flash corruption can occur after the power input to the board disconnects. The only way to minimize the chances of Flash corruption is by holding the processor in reset when there is a supply-voltage fault. The probability of Flash corruption can be significantly reduced by using a more accurate (< 1% voltage-fault-detection error) and faster (fault detection speed < 100 µs) reset generator.

www.latticesemi.com

String central inverter system

A novel system combines the benefits of both central inverters and string inverters.

By:Frank Hinrichsen and Mark Ahmling, Wind and Sun Technologies

Planners of photovoltaic power plants are no longer bound to string inverter concepts if multiple MMP(maximum power point)-trackers are required. Central inverter systems are, today, standard for the realisation of large photovoltaic power plants (**reference 1**). Plug-and-play systems that comprise the DC combiner cabinets, inverters, MV (medium voltage) transformer, and MV switchgear in one station or container are available from multiple companies. One station typically involves two or three inverters with a total output power of 1 to 1.5 MW. Each inverter uses an MPP-tracking algorithm to achieve the optimum working voltage for the photovoltaic panels.

Figure 1: System overview: Container with 2 × 1-MW central inverter and MV transformer, String Booster Box (one of up to 24) and four PV module arrangements that require individual MPP-tracking.

Therefore the electrical topology on the DC side has to be an extensive parallel connection of dozens of strings consisting of the same number of panels, same panel type, and same spatial direction. Additionally it has to be assured that partial shading caused by trees, wind turbines, and even by other panel rows is avoided for every position of the sun. That is because a partial shaded string has a significantly lower MPP-voltage and therefore delivers not much power, if connected in parallel to other non-shaded strings.

String inverter concept

Where such an arrangement cannot be realised, planners of PV power plants normally choose string inverters of typically 10 to 20 kW instead of central inverters, in order to achieve MPP-tracking for much smaller groups of paralleled strings (1 to 4). The disadvantage of this concept is that hundreds of inverters have to be installed, controlled, and maintained. Three-phase AC-cabling with four or five wires and a rather low voltage of 400 V has to be used instead of two wires (plus and minus) with about 450 to 900 V. Groups of inverters have to be connected in combiner boxes in order to increase the copper cross-section of the connection to the next MV transformer. The MV equipment has to be planned and installed additionally in the form of one or more transformer stations within the PV plant. This leads to higher costs of investment and operation.

Figure 2: String booster box SBB16-10 with 16 string inputs of up to 10 A partitioned to four independent boosters.

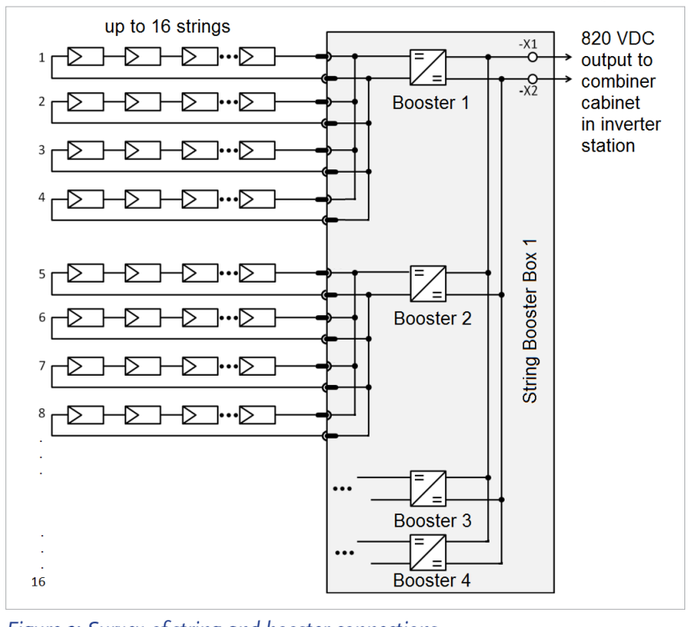

Figure 3: Survey of string and booster connections

String Central Inverter System

A novel system combines the benefits of both systems: The easy installation and control of a pre-assembled central inverter system and the flexibility in string arrangement and module-type choice known from the string inverter concept. The basic idea is to equip the string combiner boxes—that are required anyway in a central-inverter installation—with booster stages (**figure 1**).

The result is the so-called *string booster box*. One implementation, designated SBB16-10 appears in **Figure 2**. A simplified schematic appears in **Figure 3**. Every booster stage is responsible for four strings and features its own independent MPP-tracker. The common output voltage of the booster stages is limited to 850 V DC in case of open circuit. If connected to the DC-link of the inverter by automatic controlled motor-driven DC circuit breakers, this voltage is brought down to the controlled DC-link voltage of about 820 V. Thus, power transfer to the inverter station is performed on a constant high DC voltage level leading to very low power losses in cabling.

Because of the constant high DC-link voltage, the AC output voltage level of the inverter can be set significantly higher. Therefore the power rating of the inverter is also higher than if it was used in a system with conventional string combiner boxes where the DC-link voltage varies with solar irradiation. As a consequence, a central inverter system formerly rated to 1.25 MW is now able to convert 2 MW. Thereby, power loss rise is far less than linear because its main portion consists of conduction loss and the inverter output

currents remain the same. Naturally the booster stages have additional losses, but these are kept low by employing silicon carbide power devices. An efficiency of around 99 % is achieved for the box.

Figure 1 indicates the flexibility in arranging strings with different numbers of modules and different types of modules in different mounting positions. Of course, such a situation can be found in a medium-power installation with a 100 kW inverter rather than in a large PV plant that requires a 2 MW station. But that does not matter because the string booster box can be connected to small central inverters as well.
String booster boxes also fulfill the functions of fusing the strings and monitoring their currents like normal string combiner boxes do. Additionally, the string booster box integrates over-voltage protection. Communication between boxes and the inverter system is implemented by CAN-BUS while communication to the outside is via Ethernet.

Practical experience
After extensive field tests of a number of string booster boxes in an older PV plant with modified 200 kW inverters, the first PV power plant that was planned from the beginning for the novel String Central Inverter System is working well since September 2012. It features a total number of 36 string booster boxes that connect to four inverters in two containers.

www.windandsuntechnologies.com

References:

1. Hinrichsen, F.:"1 MW of Power from the Sun", Power Systems Design Europe, Sept. 2009

Special Report: Energy Efficiency; Test & Measurement

PSD EUROPE
Power Systems Design: Empowering Global Innovation

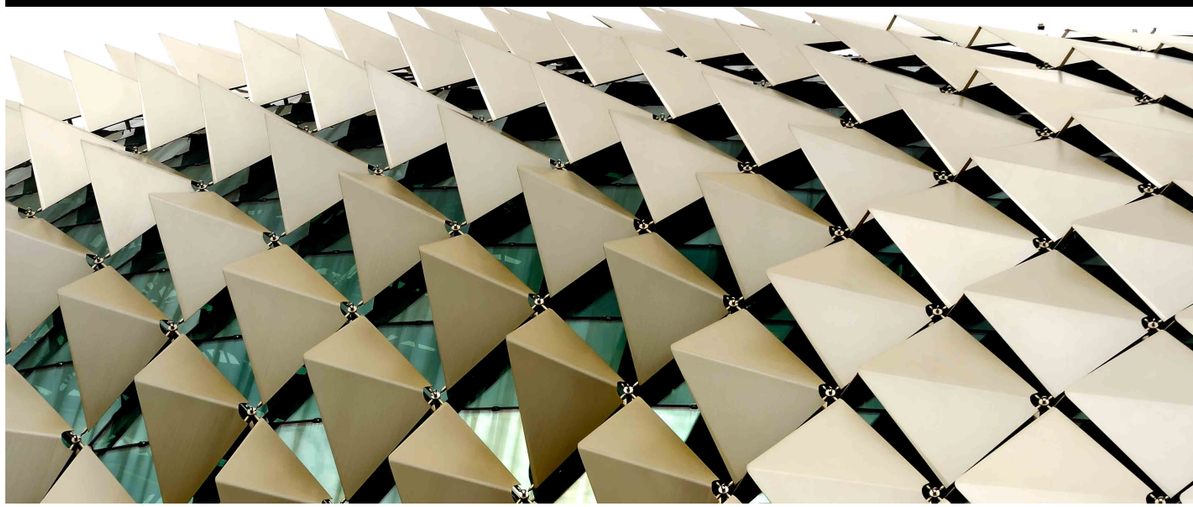

Inside:

Measuring wireless-power charging systems for portable electronics

Wireless-charging system development benefits from instrumentation that can bridge the time and frequency domains

By: Trevor J. Smith, Technical Marketing Manager, Tektronix

Mobile electronics can be found everywhere—homes, hospitals, schools, purses, and pockets. With the explosion in portability, consumers have come to expect and demand long battery life. Energy-efficient design techniques and improved battery technologies have helped to extend battery life but recharging devices is still a necessary and sometimes burdensome task.

Wireless-power charging systems look to simplify portable-electronics recharging by transmitting energy to the device without a physical connection. There are various implementations and standards for wireless charging but most use some form of electromagnetic induction. One example is the Qi (pronounced *chee*) interface as developed by the WPC (Wireless Power Consortium).

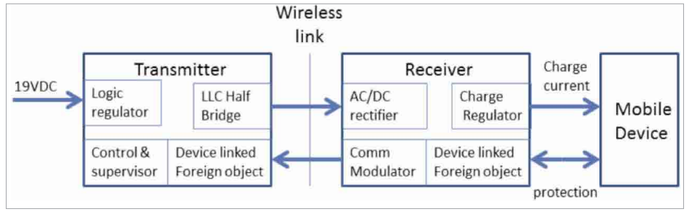

Figure 1: Qi wireless-charging system simplified block diagram

Unlike wireless-telecommunication systems like radio or cellular phones, wireless-power transmission depends more on the efficiency of transfer than signal to noise ratio. From a measurement perspective, the chargers present many challenges to the designer. The Qi wireless charger produces 5 W of charging power, while the Energy Star goals require high operating efficiency and low standby power. The efficiency of power transfer is dependent on system design including the transmitter, the receiver, and the interactions between them.

Designs typically target greater than 70% efficiency for a 5-W system. The selection of coils, shielding, components, and physical design influence the overall system efficiency This is more complicated in a wireless charging system than in a typical charger, because the wireless system requires both a transmitter and a receiver. Other complications exist due to the shielding requirements—necessary to protect sensitive electronics and the battery from the RF fields—and foreign-object detection to prevent heating of nearby metal objects.

System Overview

While both interesting and challenging, the Qi system includes low frequency modulated RF, digital, and analog circuits all on a single board (**figure 1**). The charging system uses digital communication, both for JTAG debugging and for transferring data between the secondary and primary circuits across the resonant link.

A secondary-side microcontroller monitors the charger's output voltage, generates signals, and uses modulation techniques to transfer information to the primary side. The primary-side microcontroller demodulates and interprets the information. The standard calls for organizing the modulated information into information packets comprising preamble bytes, header bytes, message bytes, and checksum bytes. Per the WPC specification, information packets can relate identification, configuration, control-error, rectified-power, charge-status, and end-of-power-transfer information.

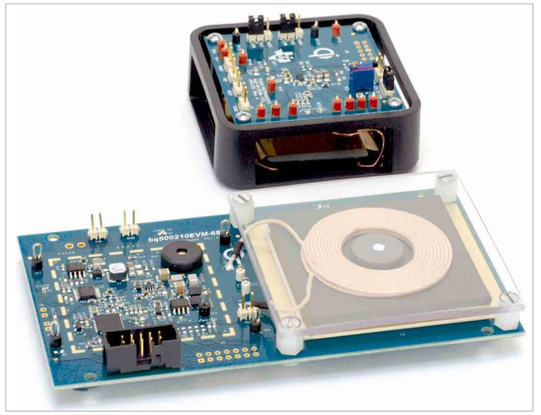

Figure 2: Texas Instrument wireless-power set, including a Bq500210EVM-689 transmitter and Bq51013EVM-725 receiver

Setting Up the Measurement

For this example, we tested a Texas Instrument wireless-power set, including a Bq500210EVM-689 transmitter and Bq51013EVM-725 receiver (**figure 2**). In addition to transferring power, the transmitter and receiver system also communicates controls signals modulated on top of the RF power signal from the secondary to the primary. The communication link allows the system to dynamically monitor and control power levels to prevent overheating due to nearby metal objects.

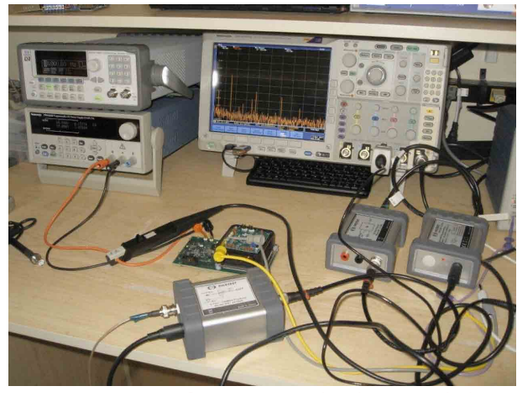

Figure 3: Test setup for measuring the charger including a mixed-domain Oscilloscope, voltage and current probes, laboratory-grade power supply, and signal injectors

For the test setup, a Tektronix PWS4323 programmable power supply provides the necessary 19-V and 500-mA DC input power (**figure 3**). The MDO4104-6 MDO (mixed-domain oscilloscope) makes the measurements. This instrument provides 1 GHz bandwidth across four analog channels and includes serial packet decode for many industry standard control buses such as I²C and SPI. The MDO4104-6 includes an RF input to measure the RF signals up to 6 GHz.

The MDO monitored control signals via the analog input and the RF signal and analyzed their respective harmonics with the RF input channel. Note, due to the planar coupling and close proximity of the primary and secondary coils it is difficult to measure the energy transfer wirelessly. Consequently, we monitored the power on the secondary side with a high impedance voltage probe directly at the output.

The Qi charger's power stage is based on a half-bridge LLC topology resonant converter controlled directly from the primary-side charger controller. The resonant converter is frequency modulated over a range of approximately 110 to 205 kHz, as required by the WPC standard, to regulate the charger output power.

There are several ways to look at the LLC resonant waveforms. **Figure 4a** shows a facsimile of the resonant-link current, measured

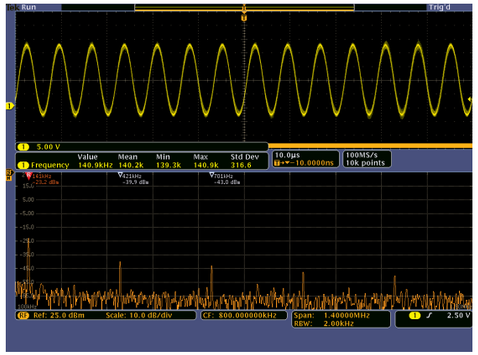

Figure 4a

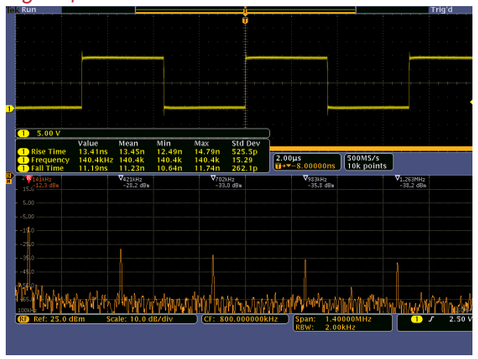

Figure 4b

Figure 4: (A) LLC half-bridge resonant-link (upper trace) current and (lower trace) switching spectrum (B) LLC half-bridge switching (upper trace) voltage and (lower trace) switching spectrum

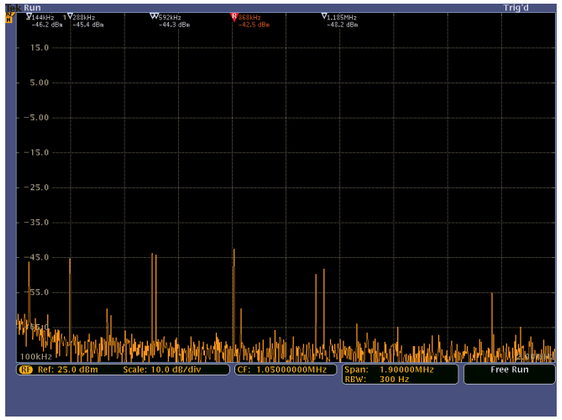

Figure 5a

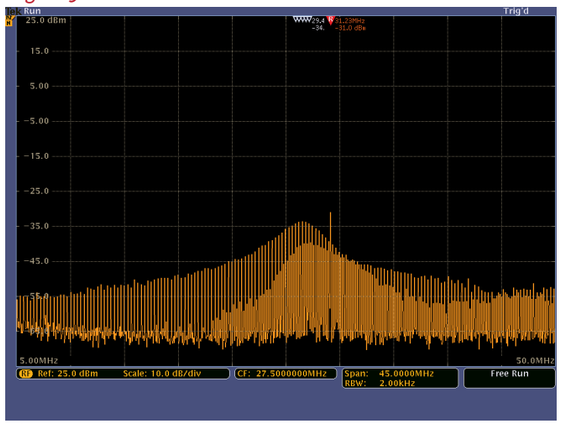

Figure 5b

Figure 5: (A) Radiated EMI signals from the LLC half-bridge switches and the 600-kHz buck regulator. (B) Radiated EMI signals from the LLC half-bridge switches, 600-kHz buck regulator and 31-MHz microcontroller.

with a voltage probe connected to the resonant capacitor link at TP2 and also shows the resonant frequency. This current is also measureable directly using a current probe or by means of the I_Sense testpoint on the transmitter board.

The RF power signal on the transmitter side is measureable with a voltage probe connected to TP1 via a Picotest J2180A preamplifier. The preamplifier provides a high input impedance and a 50 Ω output impedance, facilitating a voltage probe connection to the LLC half-bridge switch node.

The RF channel displays the fundamental operating frequency of 141 kHz and is also rich with the odd harmonics associated with the 50% duty cycle switch voltage. The lack of even harmonics provides assurance that the duty cycle is precisely 50% (**figure 4b**).

The MDO also measures higher frequency signals, such as EMI. **Figure 5a** shows the radiated EMI over a frequency range of 100 kHz to 2 MHz from the LLC half bridge switching, as well as from a low power 600 kHz buck regulator that efficiently converts the 19-V DC input voltage to 3.3 V required by the transmitter controller. **Figure 5b** shows the

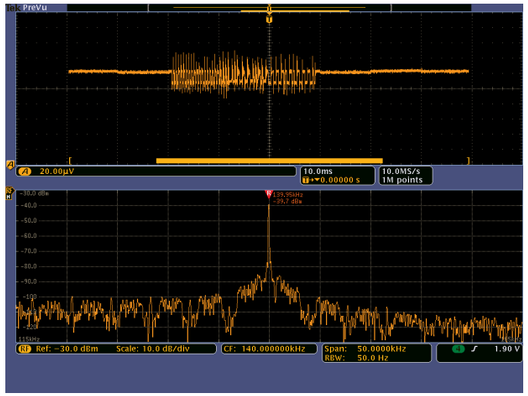

Figure 6: RF input showing (lower trace) the spectrum of the LLC half-bridge converter and (upper trace) the amplitude modulation, using the spectrum time amplitude vs time function to show the communication signals.

radiated EMI signals from the LLC half bridge switches, 600-kHz buck regulator, and 31-MHz microcontroller over a frequency range of 5 to 50 MHz.

The digital communications across the resonant link implement using either resistive or capacitive techniques. Either method results in an amplitude modulation of the primary voltage. An MDO4000's spectrum-time capabilities show the time-varying nature of the modulation behaviour, specifically with the amplitude vs. time waveform (**figure 6**).

The upper trace shows the amplitude modulation signal while the lower half shows the resonant link signal in the spectrum view. The MDO can extract the digital information using either a voltage probe or a near-field H probe, connected to the RF input via the Picotest J2180A preamplifier. This set up uses a near-field probe set from Electro-Metrics.

The MDO can also display the communication signals in the time domain on the receiver modulation-control pin and transmitter-primary winding voltage (**figure 7**). Both of these signal measurements use a TDP0500 differential voltage probe for maximum clarity and minimum circuit loading. The differential probe is more important in the measurement of the primary voltage than it is for the receiver controller voltage. This is due to the receiver control signal being ground-referenced and of relatively low impedance. The primary voltage is floating and, being part of the resonant tank

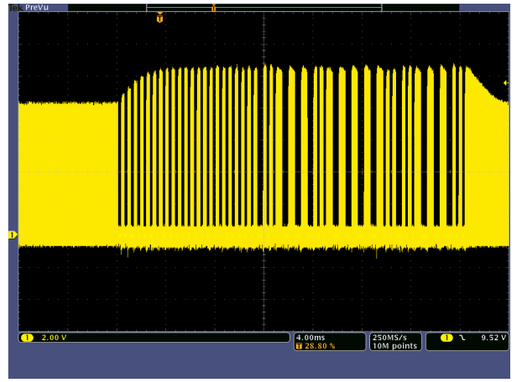

Figure 7a

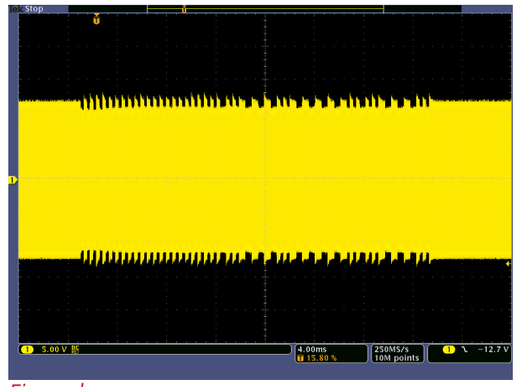

Figure 7b

Figure 7: (A) The receiver modulation-control pin and (B) the primary-coil voltage measured with a TDP0500 differential probe.

circuit, the primary voltage is more sensitive to loads such as the input capacitance of voltage probes.

The Qi charger ultimately provides a 5-V 1-A output to charge portable electronic devices' batteries. The LLC half-bridge converter coarsely regulates the secondary side voltage. A 5-V LDO precisely regulates the output.

A small-signal step load current

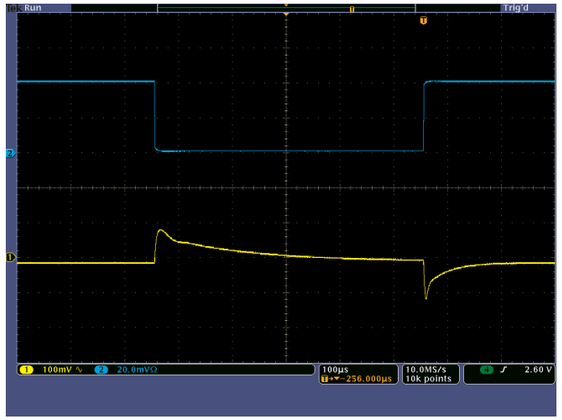

Figure 8: Dynamic-load response of the 5-V output: The blue trace is current (20 mA/div) and the yellow trace is voltage (100 mV/div).

applied to the 5-V output, using a Picotest J2111A current injector, serves as the stimulus to measure the dynamic response and control loop stability of this final output regulator (**figure 8**). The current injector is in place of an electronic load to allow faster rise and fall times.

The primary value of a mixed-domain oscilloscope is its ability to make time-correlated measurements across two domains: time and frequency. In addition, it can make these measurements between multiple analog, digital, and RF signals.

Time-correlated measurements allow the mixed-domain oscilloscope to measure timing relationships between all of its inputs. It can measure, for instance, the time between a control signal and the beginning of a radio transmission, the rise time of a transmitted radio signal, or the time between symbols in a wireless data stream. The MDO can analyze a power-supply-voltage dip during a device state change and correlate it to the impact on the RF signal. Time correlation is critical for understanding the complete system operation: cause and effect.

Using the MDO4000 we were able to measure the logic, RF, and analog functions of the Qi wireless battery charger, aided by a few probes and accessories. In the case of the communication signal, the MDO measured the signals in more than one domain. This included monitoring control signals, RF received output with spectrum and time-domain views. This capability allowed us to see the signal at its point of origin, within the RF link signal, and at the point of receipt across the transmitter winding. We also measured the performance of the final LDO output regulator and captured the EMI signals.

www.tek.com

Hall-effect sensors deliver higher efficiency in brushless-DC motors

Multiple sensor-performance characteristics influence commutation timing and, ultimately, motor efficiency

By: Josh Edberg, Global Marketing Manager, Honeywell Sensing and Control

As energy efficiency and cost savings become the biggest drivers in the design of electronic equipment, BLDC (brushless-DC) motor manufacturers must adapt to these new requirements, and deliver motors that operate more efficiently. One way BLDC motor designers can increase efficiency is by selecting the right bipolar latching Hall-effect sensor for electronic commutation. Higher efficiency translates into less material required to deliver the same amount of power in a BLDC motor, which helps reduces cost.

Playing a big role in motor efficiency, Hall-effect sensors can significantly affect the reliability, performance, and life cycle of many critical applications, ranging from robotics and portable medical equipment to HVAC fans and machine tools. These applications generally require a highly efficient and quiet motor, which are two key advantages of BLDC motors.

Unlike brush DC motors, which use mechanical commutation, BLDC motors employ electronic commutation to control the power distribution to the motor. This entails the use of Hall-effect sensors and control circuitry to detect the position of the rotor. A magnetic field from a permanent magnet or an electromagnet drives these Hall-effect sensors. They respond to a positive magnetic field (south) to operate and a negative magnetic field (north) to release.

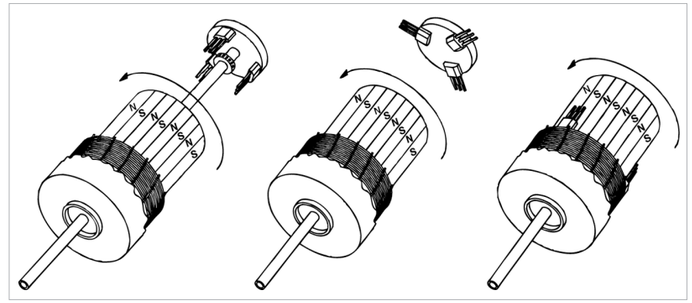

Figure 1: Hall-effect sensors can mount in three locations: inside the motor, at the end of the motor's shaft, and around the rotor shaft's ring magnet.

The Hall-effect sensors can mount in any one of three locations in the motor to measure correctly the motor's position. They may mount inside the motor, at the end of the motor's shaft, and around the rotor shaft's ring magnet (**figure 1**). Their job is to determine when to apply the current to the motor coils. This data communicates to the electronic controller to rotate the magnet at the right time and orientation.

When selecting a bipolar latching Hall-effect sensor to commutate the motor, BLDC-motor manufacturers need to evaluate several key design characteristics to ensure that the motor will operate efficiently. These include sensitivity, repeatability, stability over temperature, and response time. Hall-effect sensors should offer consistent repeatability and stability and provide a fast response to changes in the magnetic field to deliver enhanced motor efficiency.

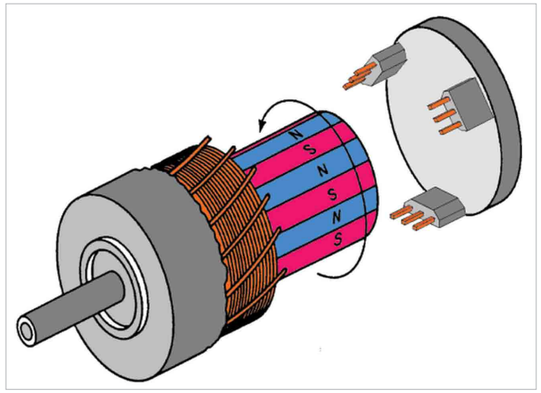

Figure 2: BLDCs use electronic commutation to control the power distribution to the motor, using Hall-effect sensors and electronic controllers to rotate the magnets at the right time and orientation.

Here's why BLDC-motor manufacturers should evaluate these characteristics when selecting Hall-effect sensors:

Sensitivity

Three factors determine a Hall-effect sensor's sensitivity level: The placement of the sensor with respect to the magnet, the air gap, and the magnet's strength. The sensor is activated by a magnetic field, and changes state—operates and releases—based on magnetic field strength, measured in Gauss.

By selecting a high-sensitivity sensor, which is typically rated at less than 60 Gauss, BLDC motor manufacturers can use smaller magnets or less expensive magnetic materials, allowing the BLDC motors to be used in compact designs while reducing costs. This is increasingly becoming an issue as the price for rare earth magnets continues to rise and supply tightens.

Another benefit of high sensitivity is that it allows for a wider air gap. This provides greater design flexibility by allowing the motor designer to place the sensor further away from the magnet and still maintain reliability. If the designer does not opt for a wider air gap, the sensor exhibits higher sensitivity, translating into higher reliability and repeatability. The upshot: A Hall-effect sensor with high sensitivity or a lower magnetic switch point delivers more efficient motor performance.

As an example, Honeywell Sensing and Control's SS360ST, SS360NT, and SS460S high-sensitivity bipolar latching Hall-effect sensors operate from only 30 Gauss, typical, at 25 °C, and 55 Gauss, maximum, over the full operating temperature range of –40 to 150 °C. This allows BLDC manufacturers to build smaller products while eliminating the need for higher-cost magnetics (**figure 2**).

Repeatability

High sensitivity allows a Hall-effect sensor to be more repeatable. Repeatability is based on a Hall-effect sensor's latching time. The sensor's output directs current through the coil windings in the motor. This produces a magnetic field that interacts with the field from the permanent magnets on the shaft, causing the shaft to spin. A highly repeatable sensor changes state at the same angular position each time the magnet passes by the sensor.

As an example, one complete revolution of the motor shaft is 360 degrees. If the sensor turns on at five degrees when the shaft spins the first time, does it change state at five degrees at the next rotation? A highly repeatable sensor will maintain all of the angular measurements close to the same value, which in this case, is five degrees.

This is critical because the timing between the current flow through the coil and the shaft's position must be highly accurate to produce the maximum amount of torque on the shaft. Any delay in a sensor's response to changes in the magnetic field typically leads to lower bandwidth and accuracy errors. An error in the switching point of the Hall-effect sensor reduces the motor's torque, translating into lower motor efficiency (**figure 3**).

Stability

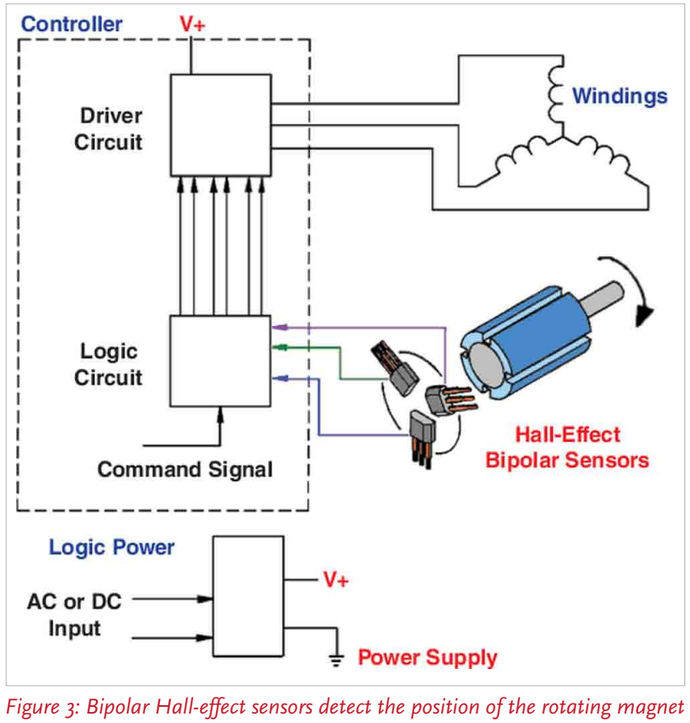

Figure 3: Bipolar Hall-effect sensors detect the position of the rotating magnet in BLDC motors. In this example, an eight-pole motor with a three-phase winding uses three bipolar latching Hall-effect sensors.

Stability refers to how much the angular position changes over temperature or voltage. A highly stable Hall-effect sensor will require the same Gauss level to turn a part on whether it's at 25 or 125 °C. So if a part operates at 30 Gauss at 25 °C, does it operate around 30 Gauss at 125 °C or does the operating point shift?

This is an important characteristic because stability-over-temperature, coupled with high sensitivity, is required for precise position detection. In addition, magnetic stability improves jitter performance—a critical factor in BLDC efficiency, which contributes to less speed variation.

Response time

Response time is the time it takes for the sensor to change state. As an example, when a 30 Gauss magnetic field level is applied to the sensor, which has an operating point of 30 Gauss, the response time is measured from the point when the Gauss field is applied to when the sensor's output changes state.

If a sensor has a slow response time—switching at a different magnetic field level than what is required—it can result in lower bandwidth, accuracy errors, and lower repeatability. Ultimately, this translates into lower motor efficiency and issues commutating a motor at high frequencies.

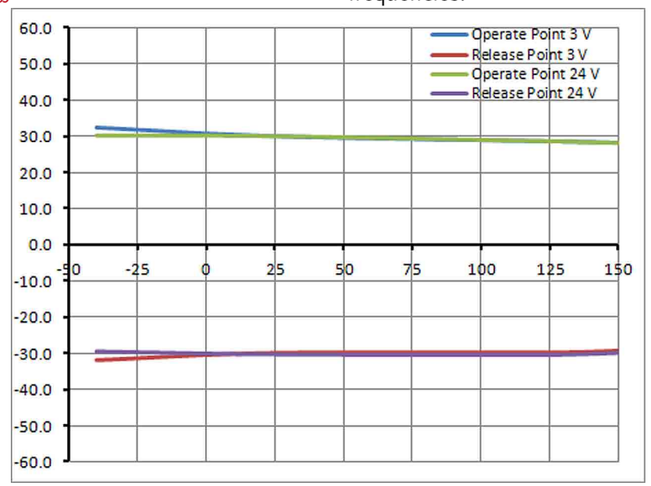

Figure 4: The Honeywell SS360ST, SS360NT, and SS460S bipolar latching Hall-effect sensors offer reliable switching points over its operating temperature range of –40 to +150 °C.

Miscommutation also results in lower effective torque constant (Kt) and higher torque ripple, causing additional electrical noise, which affects efficiency and system performance. Thus, a faster response time to a change in the magnetic field provides higher efficiency in commutating a BLDC (**figure 4**).

Typically, sensor manufacturers have used chopper stabilization to achieve high stability and high sensitivity. This averaging process on the hall die allows for greater stability over temperature while improving the sensitivity by mitigating the inconsistencies from sensor to sensor due to packaging stresses. This is especially the case with single and dual Hall elements, which are very susceptible to packaging stresses and are often used for high sensitivity Hall-effect sensors.

Although chopper stabilization offers high sensitivity, this technique does have some drawbacks, which negatively affects motor efficiency. Because chopper stabilization is an averaging process, it will always increase the response time when compared to a sensor that is not chopper stabilized. This delay can be in the range of 10 to 30 µs, which can be very significant to motor efficiency and, in the case of high-frequency motors, it may not be possible to use chopper-stabilized parts. In addition, chopper-stabilized parts often have issues with repeatability and increased electrical noise.

Although chopper stabilization has traditionally produced highly stable and sensitive Hall-effect sensors, motor designers should evaluate new technologies that can eliminate some of those drawbacks. As an example, Honeywell's Hall-effect sensors use a quad Hall element and proprietary software to account for any drift in the switch point, eliminating the need for chopper stabilization. These sensors provide a fast response time, reduce sensitivity to packaging stresses, and generate less noise, resulting in improved motor efficiency.

Hall-effect sensors used for electronic commutation in BLDC motor designs play a big role in a motor's efficiency. Motor designers need to look closely at several specifications including sensitivity, repeatability, stability, and response time to achieve a highly efficient motor that can also be designed into space-constrained applications. They also need to look at Hall-effect sensors utilizing new technologies to deliver highly stable and sensitive parts without the side effects of increased electrical noise and slower response time often associated with chopper stabilization.

sensing.honeywell.com

Designing a 75-W SEPIC

A combination of analysis, simulation, and response measurement leads to a design with excellent gain and phase margins

By: Florian Mueller, Reference Design Engineer EMEA, Texas Instruments

Engineers are often afraid of designing a SEPIC (single-ended primary-inductance converter). In particular, the compensation of this topology is more complicated compared to other topologies. The SEPIC's LCL arrangement adds an additional double pole and a second RHPZ (right half plane zero) to the transfer function. However, the SEPIC has some advantages like low input ripple current, low EMI, and, of course, the possibility to step down or step up the input voltage. An awareness of some of the SEPIC's characteristics can simplify the design process and reduce development time.

SEPIC versus flyback

The SEPIC capacitor clamps the voltage of the MOSFET and therefore limits the ringing and the overshoot on the switch node. This reduces EMI and improves overall efficiency. The SEPIC also provides better cross regulation in multiple-output applications. The disadvantages of the SEPIC compared to a flyback converter is the missing isolation between the primary and secondary side and the more complicated compensation.

Dimensioning the power stage

For the power stage, TI's *power-stage designer* tool is a great help to the engineer (**reference 1**). It is available as a free download and simplifies component selection.

The capacitance value of the SEPIC capacitor is not particularly critical, but the circuit designer must ensure that the capacitor can handle the AC and RMS current. The tool simplifies the process of determining these values. Increasing the capacitance will decrease the undesired circulating current, but this is only necessary if the design uses a well-coupled inductor with low leakage inductance (**reference 2**).

The minimum necessary output capacitance depends on the maximum load step on the output. At the very first moment of a load step, the additional current draws from the output capacitance. At this time, the controller is starting to change the duty cycle to deliver more energy to the output. At the time $t_{crossover}$, which equals the inverse of the converter's bandwidth, the controller delivers all of the current. This behavior leads to an assessment of the minimum output capacitance, $C_{OUT(MIN)}$:

$$C_{OUT(MIN)} = \frac{1}{2}\frac{\Delta I_{LOADSTEP}}{BW \cdot \Delta V}$$

where ΔILOADSTEP is the step change in output current, BW is the converter's bandwidth, and ΔV is the resulting output voltage ripple.

As the equation suggests the size of the output capacitance depends on the system's bandwidth. If the converter's bandwidth doubles, the converter requires only half of the output capacitance. Therefore it is important to design a proper compensation network to get a high bandwidth with a reasonable phase and gain margin at the same time.

Unfortunately, the switching frequency, error amplifier, and the RHPZ limit the maximum achievable bandwidth of the SEPIC. In systems without an optocoupler the RHPZ is often the limiting factor. Topologies which are operating in CCM and delivering energy during the off-time own a RHPZ.

Compensation

Unfortunately the compensation of a SEPIC is more complicated

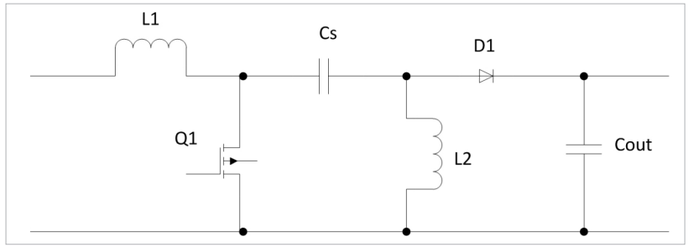

Figure 1: SEPIC power stage

than of a flyback. There is a resonant circuit consisting of the inductor L1, capacitor Cs and inductor L2 (**figure 1**). This circuit adds an additional double pole and a second right half plane zero to the power stage's transfer function. In many cases this is not a problem, because the affect depends on the quality factor of the LCL-network. If the quality factor is low, then the designer can compensate the SEPIC like a flyback converter.

If the converter is operating in voltage mode, the transfer function shows two double poles. The first double pole results from the output filter of the power stage, mainly from the output capacitor (plus the capacitance of the SEPIC capacitor) and the two inductors. The second double pole results from the resonant circuit mentioned above.

If the converter is operating in current mode or in DCM (discontinuous conduction mode) the first double pole degenerates into a single pole. The second double pole does not change.

In voltage and in current mode the power stage shows two RHPZs when operating in CCM (continuous conduction mode). The first RHPZ only depends on the input inductor L1, the duty cycle, and the load resistor. The second RHPZ depends on the output inductor L2 and the values of the parasitic resistance in the circuit. There will be a high phase shift in the system. It is impossible to reach a bandwidth beyond the frequency of the RHPZ.

Simplified Analysis

The following formulae result from a simplified analysis of the circuit. The analysis does not include all the parasitic elements in the circuit. Designers can complete the more complex analysis, including parasitic elements, with Vorperian's circuit analysis techniques (**reference 3**). For further simplification of the formulae, the duty cycle is set to 50%.

$$L_{TOTAL} = L_1 + L_2$$

$$C_{TOTAL} = C_{SEPIC} + C_{OUT}$$

where COUT is the converter's total output capacitance.

In voltage mode, the first double pole occurs at

$$f_{LCdpole1} = \frac{1}{2\pi\sqrt{L_{TOTAL}C_{OUT}}}$$

$$Q_{LC\_pole1} = \frac{R_{LOAD}}{2\pi f_{LCdpole1}L_{TOTAL}}$$

In current mode, the first pole occurs at

$$f_{RCpole1} = \frac{1}{2\pi R_{LOAD}C_{OUT}}$$

The first RHPZ depends on the input inductor, load resistance, and duty cycle:

$$f_{RHPZ1} = \frac{R_{LOAD}(1-D)^2}{2\pi L_1 D}$$

where D is the duty cycle

The second double pole is due to the resonant circuit:

$$f_{dpole2} = \frac{1}{4\pi}\sqrt{\frac{L_{TOTAL}C_{TOTAL}}{L_1L_2C_{OUT}C_{SEPIC}}}$$

$$Q_{LC\_pole1} = \frac{R_{LOAD}}{2\pi\frac{f_{LCdpole1}^2}{f_{dpole2}}L_{TOTAL}\frac{C_{SEPIC}}{C_{OUT}}}$$

The second RHPZ depends on the output inductor and the SEPIC capacitor:

$$f_{RHPZ2} = \frac{1}{2\pi}\sqrt{\frac{D}{L_2C_{SEPIC}}}$$

A 75-W SEPIC power reference design appears in **figure 2**; the power stage transfer function in **figure 3**.

Power stage transfer function

A 75-W SEPIC power reference design serves as an example for the transfer-function analysis (figure 2). Using the component

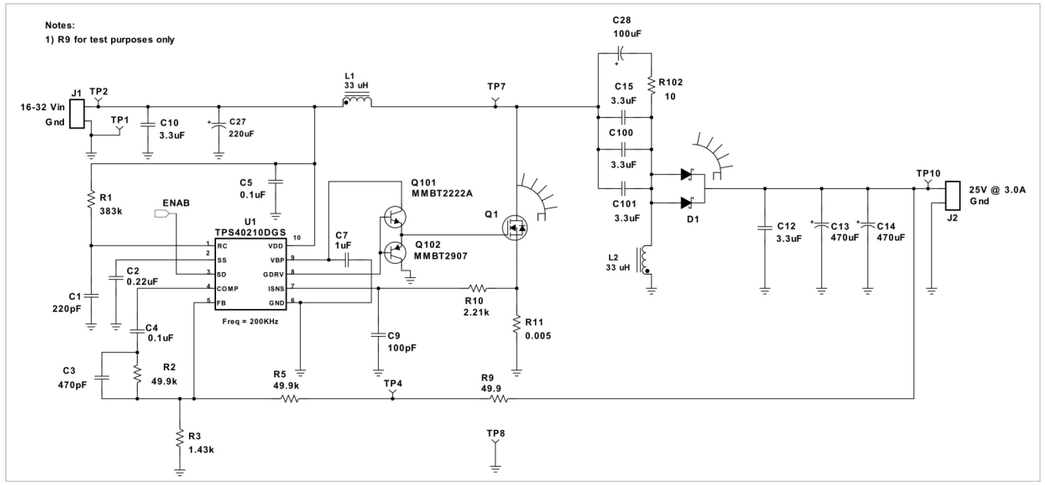

Figure 2: Schematic of a 75-W SEPIC

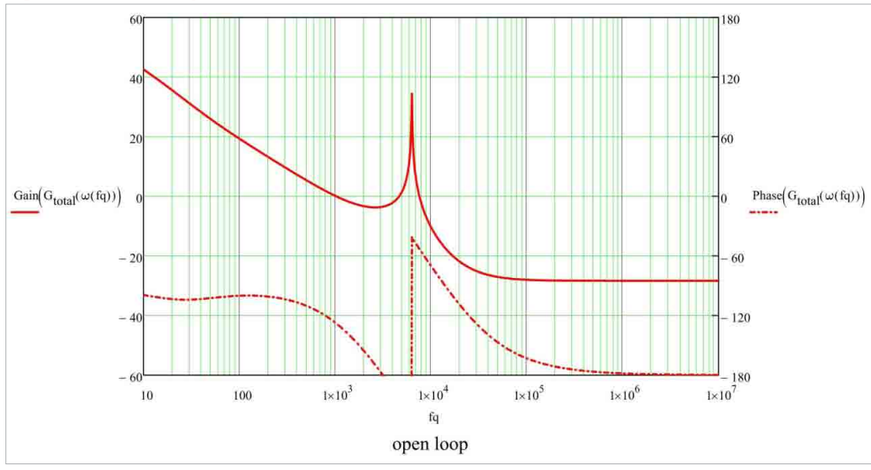

Figure 3: The power stage's calculated bode plot

values from the reference design, the loop of the power stage shows a gain increase between 5 to 10 kHz due to the second double pole, which limits the reachable bandwidth (figure 3).

From the Bode plot, $f_{ESRzero}$ = 12.1 kHz, $f_{RCpole1}$ = 0.02 kHz, f_{dpole2} = 6.3 kHz, f_{RHPZ2} = 6.7 kHz., and $f_{ESRzero}$ = frequency of the zero (output capacitor). To get a stable system the gain must be reduced by more than 40 dB. Therefore the crossover frequency (bandwidth) will be low, resulting in a slow system.

There is another approach: the resonant circuitry can be damped by adding a capacitor and a resistor in parallel to the SEPIC capacitor (C28 and R102 on schematic). The capacitance of C28 must be much higher than the capacitance of the SEPIC capacitor—the sum of C15, C100, and C101. The value of resistor R102 should be in the range of 10 Ω. Providing an additional damping network makes the compensation quite easy, because the impact of the second double pole will be lower. As the TPS40210 in the reference design is a current-mode controller, a type-II compensation network is sufficient. For a stable system with good phase and gain margin, set the frequency of the compensation

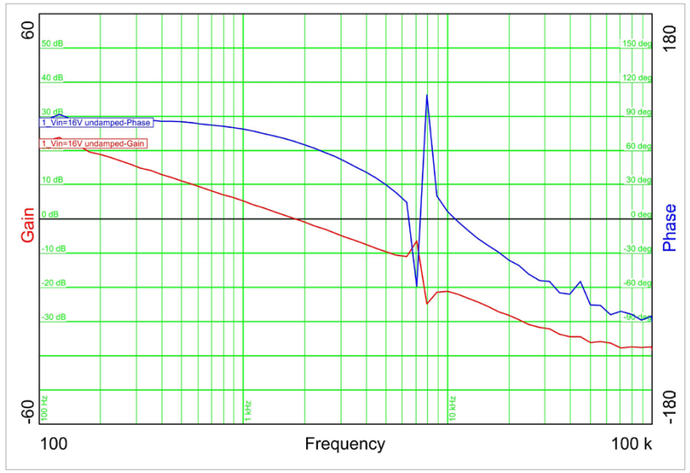

Figure 4: Bode plot of the undamped system

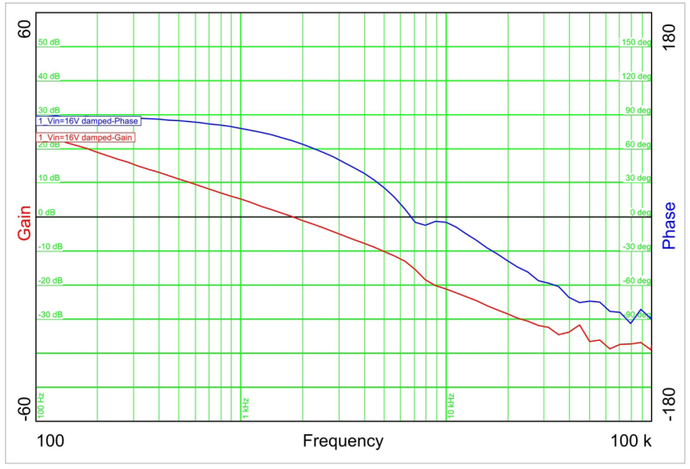

Figure 5: Bode plot of the damped system

network's zero to the frequency f_{RCpole} and set the frequency of the compensation network's pole to half of the switching frequency.

The gain, defined by the ratio R2:R5, should be set low at the first power on to ensure a stable system. After the first measurement of the loop with a frequency response analyzer, the gain can be increased. The phase margin should be around 50 to 60 degrees and the gain margin around 10 to 20 dB.

Figure 4 shows the measured open loop of the undamped system. The gain increases at around 6 kHz, which causes problems for the stability of the system. The switch node shows jitter.

Adding the damping circuitry—a 100 uF electrolytic in series with a 10 Ω resistor—changes the Bode plot (**figure 5**). Now the phase and gain margin is in a range which guarantees stable operation.

In many cases a SEPIC can be compensated like a flyback converter. The quality factor of the LCL-circuitry depends on the parasitic elements and the duty cycle. Therefore it is not always necessary to damp the circuitry, especially for high-input-voltage applications. Take care if the input voltage decreases, as the gain of the double pole frequency might rise and the undamped system could become unstable. Check the compensation with a frequency response analyzer, as a detailed simulation is not always possible due to the parasitic capacitances and inductances.

www.ti.com

References:

1. *Power-stage designer* download, Texas Instruments, http://www.ti.com/tool/powerstage-designer
2. Kollmans, Robert, *Power Tip number 32 and 33*, http://www.youtube.com/watch?v=ydkoaoDPIZU
3. Vorperian, Vatche, *Fast analytical techniques for electrical and electronic circuits*, Cambridge University Press, 2002

Rethinking battery management for EVs and grid-tied energy storage

By: David G. Morrison, Editor, How2Power.com

Large-format, high-power battery packs are expected to play a central role in two technological revolutions. One is the electrification of the automobile through the further development and adoption of electric and hybrid electric vehicles. Another is the deployment of battery-based energy storage for the power grid. But before either of these technological transformations can occur on a wide scale, battery technology must advance to provide the performance, operating life, and low cost required in the targeted applications.

Much of the news regarding battery development concerns improvements in cell chemistry or construction, but a recent article by Eric Macris and John Bendel of Element Energy explains how innovation in battery management can have a dramatic impact on the critical aspects of battery performance and cost.

In their article, *Hardware-Software Architecture Improves Performance Of Large, Battery-Powered Systems*, which was published recently in How2Power Macris and Bendel describe an integrated hardware and software-based approach to battery management for large, multi-cell battery packs that promises to double battery cycle life, while also extending battery run time (per charge) as much as 10%.

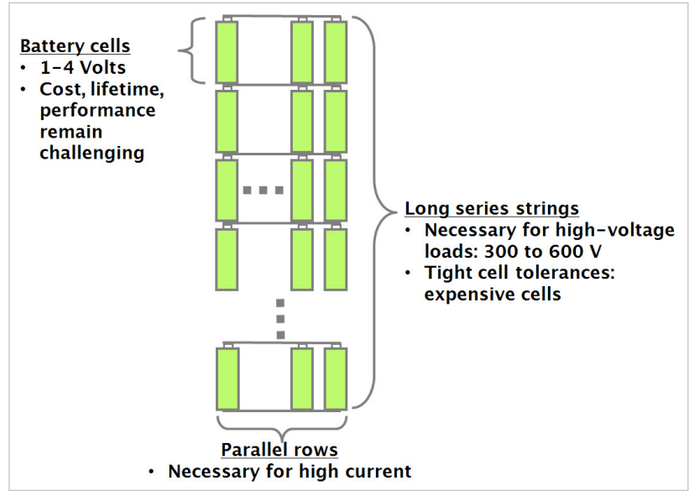

Figure 1: Typical battery packs like those being developed for electric vehicles and grid-tied energy storage systems are expensive, short-lived, big, complex, and limiting. Cell variations drive the degradation in pack capacity over the life of the pack. Ultimately, the weakest cells constrain both pack performance and operating life.

Other benefits of Element Energy's battery management system—which is agnostic with regard to cell chemistry—include a 30% reduction in charging time (with no penalty), integration of voltage regulation within the battery pack, and active control of cell-level isolation from the pack for improved pack safety and reliability.

As Macris and Bendel explain,

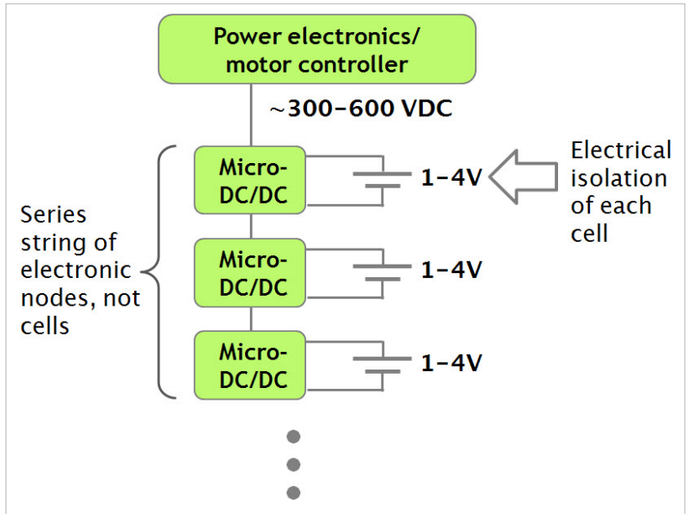

Figure 2: Element Energy's Integrated Energy Management System employs a string of dc-dc converters to buffer each cell in a large, multi-cell battery pack, which enables the voltage and current of each cell to be managed independently across variations in pack-level voltage and current.

the main challenge in controlling charging and discharging of large, multicell battery packs stems from the variability in the characteristics of the individual cells, which are typically connected in long series strings to generate high voltages (**figure 1**). One important effect of cell variability is that some cells charge faster than others. Without intervention from the battery management system, these fast charging cells could become overcharged to the point of thermal runaway. In the conventional approach to battery management, shunt resistors are switched in parallel with the fast charging cells to prevent their overcharging, while giving other cells the chance to reach full charge.

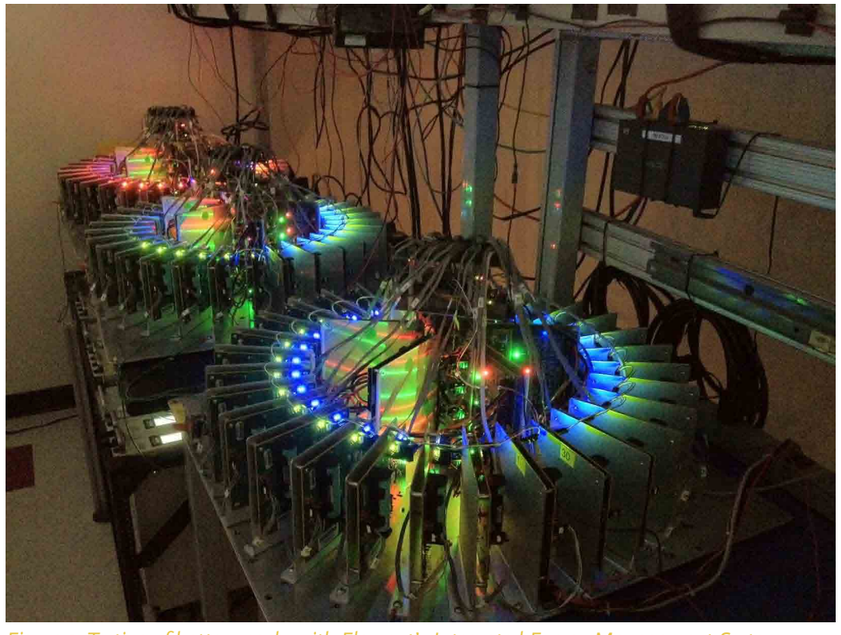

Figure 3: Testing of battery packs with Element's Integrated Energy Management System.

As the authors explain, this conventional approach to battery management has major drawbacks. "It is a compromised solution to the problem of variability among cells. It provides only a limited mechanism for managing the uniformity of charge balance among the cells. Moreover, it essentially forces dissimilar cells to perform as if they are identical—a goal that can be achieved for a time, but at the expense of pack lifetime. In the long run, this architecture shortens pack lifetime by exacerbating the differences among cells, making weak cells even weaker over time by forcing them to provide performance similar to strong cells. These weak cells ultimately govern pack lifetime."

With its novel system architecture, Element Energy departs from the conventional approach using shunt resistors. Instead, it employs a combination of hardware and software to optimize charge and discharge of each cell continuously. In their architecture, a dc-dc converter buffers each individual cell from the rest of the system. This allows the voltage and current of each cell to be managed independently across variations in pack-level voltage and current (**figures 2** and **3**). By continuously adjusting the electrical load for each cell, Element's Integrated Energy Management System maximizes each cell's contribution of energy to the battery pack over the operating lifetime of the pack. This leads to the doubling of battery life mentioned previously.

According to Macris and Bendel, this system can be implemented with a modest number of relatively inexpensive components. And while the hardware itself does not heavily impact the cost of the battery management system, it also enables a newfound flexibility. Control of individual cell loading permits cell management to be abstracted into the software layer of system architecture, where according to the authors, "charge and discharge profiles can be managed truly independently and continuously for each cell in a battery pack."

Element Energy put its battery management system to the test in a series of experiments, testing large battery packs assembled with off-the-shelf, 18650 cylindrical, Li-ion cells. One control group of battery packs was built using conventional battery management methods, while an experimental group of battery packs was built using Element's Integrated Energy Management system. Full details on the design and construction of these battery packs, test conditions, results, and analysis are presented in Macris and Bendel's article in How2Power Today, December 2012 issue, which is available online at www.how2power.com/newsletters.

These experiments represent the first large-scale tests the company has conducted on its battery management system. And while these experiments support the claims of improved battery performance described above, the authors claim that these initial results are just a start in harvesting the benefits of their new technology. As the authors state, "With improvements in cell behavioral modeling and charge time uniformity, along with various other software advances, one can reasonably expect an additional 30% extension in cycle life, even shorter charging times, and various other performance gains."

www.how2power.com

About the author:
When not writing this column, David G. Morrison is busy building an exotic power electronics portal called How2Power.com. Do not visit this website if you're looking for the same old, same old. Do come here if you enjoy discovering free technical resources that may help you develop power systems, components, or tools. Also, do not visit How2Power.com if you fancy annoying pop-up ads or having to register to view all the good material. How2Power.com was designed with the engineer's convenience in mind, so it does not offer such features. For a quick musical tour of the website and its monthly newsletter, watch the videos at www.how2power.com and www.how2power.com/newsletters/.

Test and measurement of a straw in the wind

By: Gail Purvis, Europe Editor, Power Systems Design

As any industrial journalist of a few decades ago knew, the way to discover if a depression was lifting from the manufacturing sector was to go and talk to the gas suppliers. If industrial gas sales were on the up, manufacturing was also looking up.

Some of this may have changed. And it begins to be feasible that there may be different bellwethers for differing sectoral-manufacturing recoveries.

Recently, an intriguing press release arrived from Livingston, the test-and-measurement rental company. It took the cheerful approach of promoting means where customers only pay for equipment for as long as it is actually of use, rather than being left with redundant equipment after deployment has been completed, or when the test and measurement requirements change.

Customers are not exposed to on-going running costs (such as insurance, maintenance, repair, recalibration, downtime cover, transportation, and disposal) either. Furthermore, peaks or troughs in equipment demand can also be handled easily.

Building on success in Europe, India, South East Asia, Australia and New Zealand, test and measurement sourcing specialist, Livingston is now building an explorer's global reach. The company announced the establishment of new operations that focus on addressing customer requirements in the Middle East and Africa and it will be able to offer models from EXFO (test and service communications) of Anristu, Tektronix, Aeroflex, Fluke, JDSU, and Spirent for renting, as well as a large stock of used equipment.

"With more advanced telecom infrastructure being rolled out and test activity increasing as a result, there are very clear opportunities for equipment rental in both the African and Middle Eastern regions," said Basel Shubair, MD of Livingston's new Middle East & Africa business, with a Dubai HQ.

"Telecom operators and their subcontractors need to have access to test tools that will allow them to implement higher capacity networks and to expand coverage, but in many cases cannot deal with heavy financial outlays that purchasing these items would entail. Partnering with Livingston will prove invaluable to them," he continued.

"The number of mobile subscribers in Africa is expected to reach the 1 billion mark within the next three to four years, calling for upgrading of their networks to 3G technology. In the meantime states across the Middle East are beginning to deploy 4G base stations to deal with the rising popularity of mobile data services.

As the UK's MOD announces it will auction some of the radio frequency it owns below the 15 GHz frequency, (the most useful for a wide range of communications) in a sale to take place in 2014. It gives private operators a chance to buy more spectrum for 4G mobile services. Test and measurement instrumentation hire or sales may also be one place to judge interesting this market's real potential growth.

www.powersystemsdesign.com

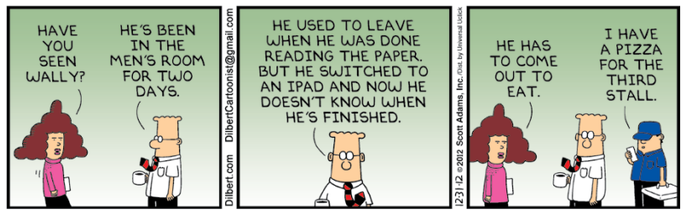

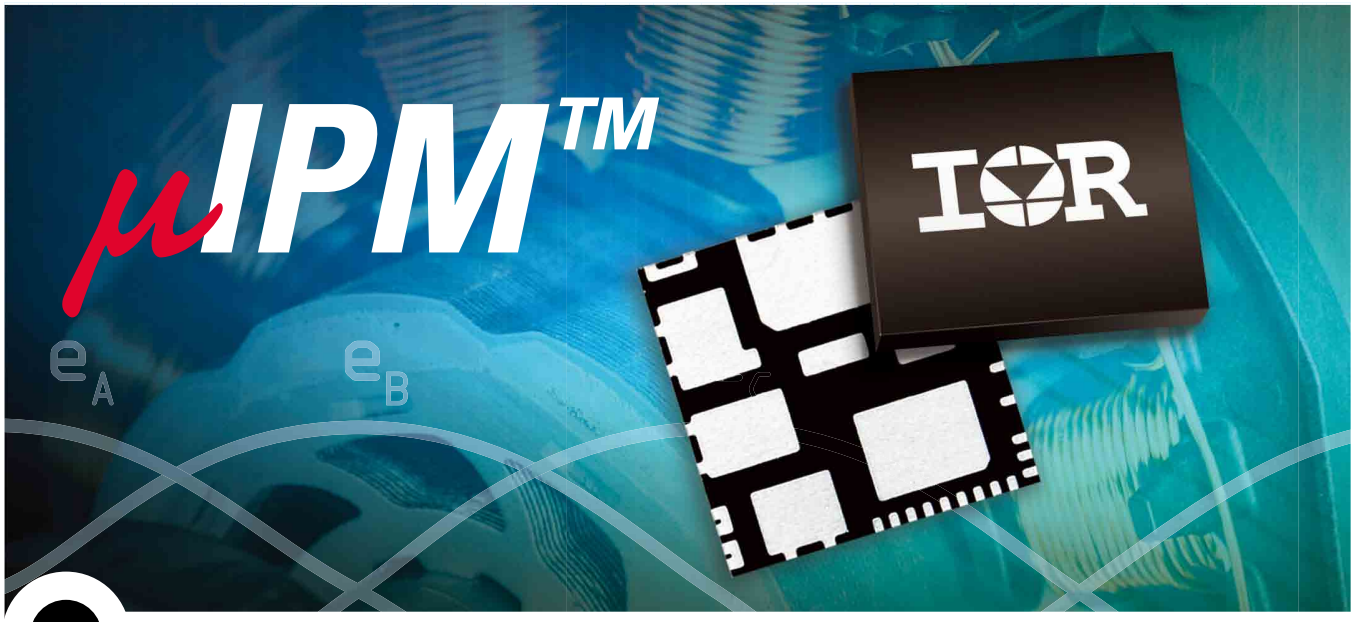

Innovative Power Module Reduces System Size

µIPM™ Power Modules Deliver up to 60% Smaller Footprint

Specifications:

Part Number	Size (mm)	Voltage	IO (DC@ 25°C)	Motor Current**		Motor Power VO=150/75VRMS	Topology
				w/o HS	w/HS		
IRSM836-024MA	12x12	250V	2A	470mA	550mA	60W/72W	3P Open Source
IRSM836-044MA	12x12	250V	4A	750mA	850mA	95W/110W	3P Open Source
IRSM836-025MA	12x12	500V	2A	360mA	440mA	93W/114W	3P Open Source
IRSM836-035MB	12x12	500V	3A	420mA	510mA	108W/135W	3P Common Source
IRSM836-035MA	12x12	500V	3A	420mA	510mA	100W/130W	3P Open Source
IRSM836-045MA	12x12	500V	4A	550mA	750mA	145W/195W	3P Open Source

Features:

- 3-phase motor control IC
- 12x12x0.9mm PQFN package offers up to 60% smaller footprint
- Eliminates the need for heat sink
- DC current ratings from 2A to 4A
- Voltage range of 250V – 500V

µIPM™ Advantages:

- Shortens design time
- Shrinks board space requirements
- Simplicity - Eliminates Heat Sink
- Replaces more than 20 discrete parts to deliver a complete motor drive stage
- Slashes assembly time and cost
- Simplifies procurement and inventory management
- Reference design kits available for quick evaluation on any 3-phase motor

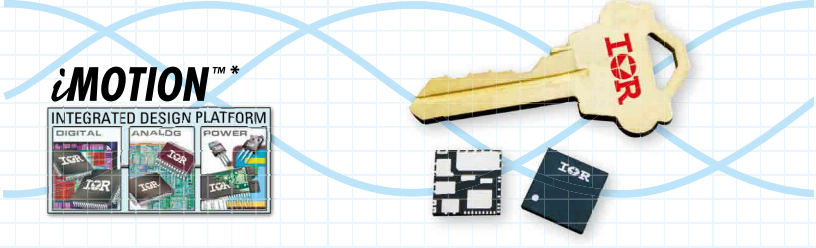

For more information call +49 (0) 6102 884 311
or visit us at www.irf.com

International
IOR Rectifier
THE POWER MANAGEMENT LEADER

* IR's iMOTION™ (ai mo shan), representing the intelligent motion control, is a trademark of International Rectifier
** RMS, Fc=16kHz, 2-phase PWM, ΔTCA=70°C, TA ≈ 25°C