PSSD EUROPE
Power Systems Design: Empowering Global Innovation
November 2012

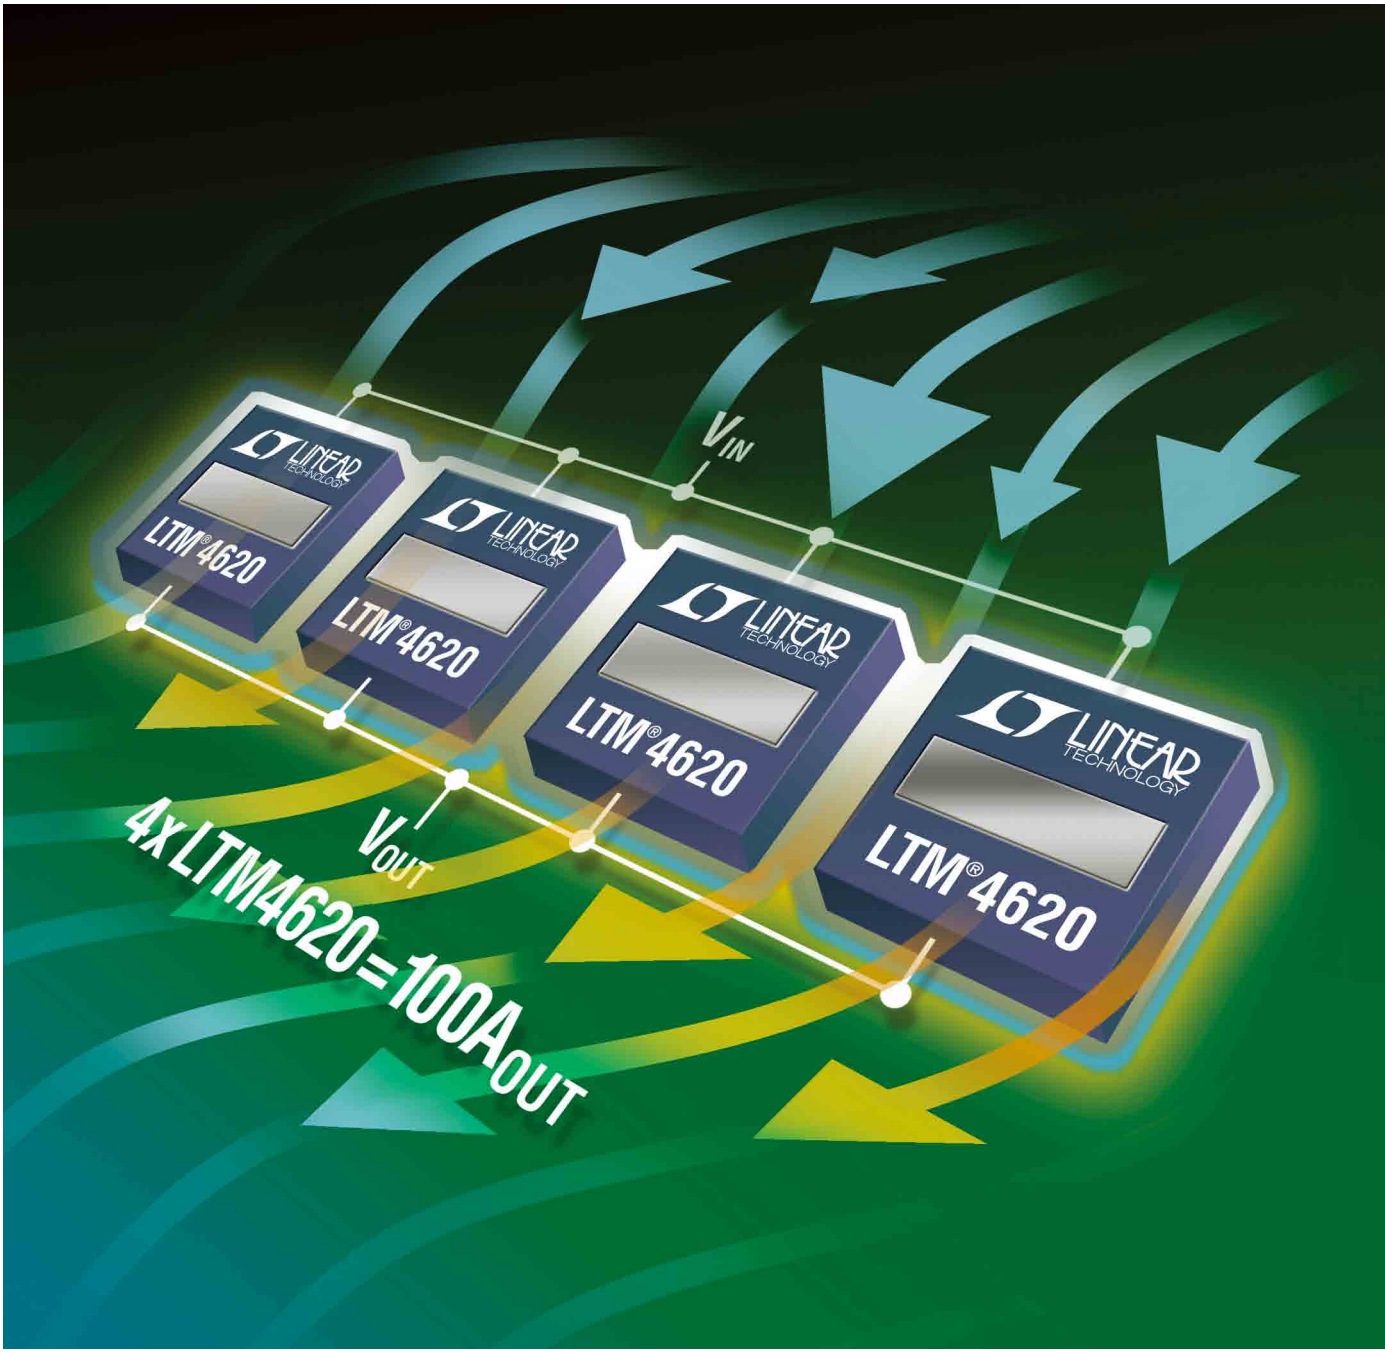

Special Report: Communication, Data, and Computing (pg37)

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

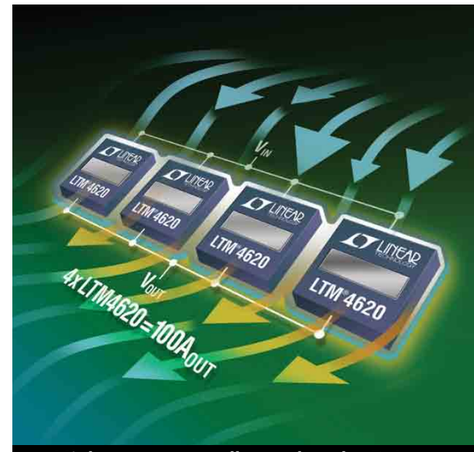

Higher power, smaller, and cooler PoL DC-DC regulation

Title (pg 14)

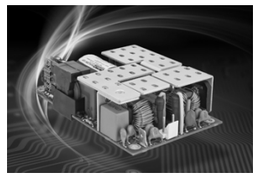

Highlighted Products News, Industry News and more web-only content, to:

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

Registration of copyright: January 2004
ISSN number: 1613-6365

Volume 9, Issue 9

Things of the Internet are power players too

One of the more intriguing approaches to the increasing intertwine of IT and electronics is being undertaken by Finland's VTT Technical Research Centre which together with its Japanese partner, the University of Tokyo, is developing uID or a universal identification technology.

Research Professor, Heikki Ailisto says this will enable the identification and tracking of individual products, components, and food items that will allow the information on origin, manufacture and history to be attached to the most commonplace items over their entire life cycle. A timber plank, for instance, can be tagged with information on which forest the timber was cut from, where it was sawn, how many times it has been painted, and with what paints."

Ailisto says that the Internet of Things (IoT) is the third of identifiably big telecommunication waves. The phone connected 500 million places. The mobile phone connects 5 billion people. IoT will connect 50 billion devices, machines and objects. Objects do not require an actual data connection, but can be named and connected to background systems, with identifiers.

Ubiquitous computing and the IoT will revolutionise technology and business and VTT has been developing its ubiquitous computing applications and basic technology in the OPENS (Open Smart Spaces) programme. The achievements to date include the implementation of interoperability platform Smart M3, which enables various appliances and objects in the home or office to *converse*, understand each other, and share information.

US technology research consultancy, Gartner, has recently named its top 10 strategic technologies for organisations in 2013, ranking IoT the 5th slot. It is preceded by after the mobile device battles and the need to support a variety of form factors with Windows 8, Google' Android, and Apple iOS; more Web apps as HTML5 becomes capable with enterprises facing complex app stores models in turn delivering mobile apps through their own private application stores, and providing an *apptrepreneurs* support ecosystems.

The PC (personal cloud) is becoming the location of choice for individual personal content and how consumers will connect with the numerous web of devices they choose to use during different aspects of their daily lives.

And of course the Cloud is an essential element for IoT, which Gartner simply sees as enabling a wide range of new applications and services, while raising many new challenges. Mobile, points out Gartner, no longer refers only to cellular handsets or tablets but is being embedded among things, as in pharmaceutical containers and automobiles.

Objects like smart meters, now *communicate* via NFC, Bluetooth, LE, and Wi-Fi to a wide range of devices and peripherals, such as wristwatch displays, healthcare sensors, smart posters, smart shirts and clothing, as well as home entertainment systems.

The IoT is a power player and an enabler, offering a wide range of new applications and services, but in turn raising many new challenges.

Gail Purvis

Europe Editor

Power Systems Design

Murata introduces new power supplies with medical approvals

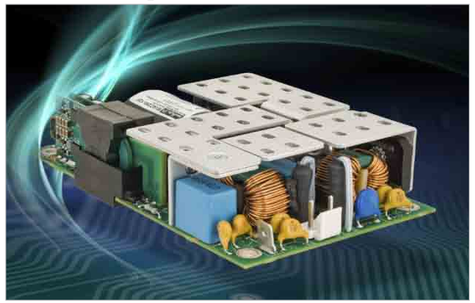

The MVAC250 and the MVAC400 comply with the UL60601-1 3rd-edition medical safety standard, the IEC60950/UL60950 safety specifications, and EN55022 standard for conducted emissions.

Murata has introduced the MVAC series of highly efficient, 3" x 5" open frame AC-DC power supplies from Murata Power Solutions. Suitable for use in a wide variety of industrial factory automation, communications, and medical equipment, the series comprises two single output models, the 250-W MVAC250 and the 400-W MVAC400 units.

Both models accommodate the universal AC input voltage range from 90 to 264 V AC with active PFC (power factor correction) and active inrush-current control. Each model is available with a nominal output voltage of 12, 24, or 50 V DC. In addition, a 12-V-DC 1-A fan supply and an auxiliary 5-V-DC 2-A output are provided across the range.

With typical efficiency ratings of 94%, Murata Power Solutions packages the units in an industry standard 1U open frame format measuring just 127.00 x 76.20 x 35.56 mm (5 x 3 x 1.4 inches). These low-profile power converters have significantly better convection-cooled performance characteristics than other competing designs. The MVAC250 can deliver up to 170-W output and the MVAC400 up to 250-W without the need for any forced airflow.

With their high convection-cooled ratings, the use of the MVAC series allows OEMs to eliminate the fan and fan controllers in their traditional power designs. In addition to the cost savings, removing cooling fans also allows designers to increase their product MTBF while decreasing system noise.

The MVAC250 and the MVAC400 comply with the third edition of UL60601-1—the harmonized standard for medical electrical equipment recognized by public-health authorities in most countries—making them suitable for use in medical applications in multiple geographies. The units also meet the IEC60950/UL60950 safety specifications and EN55022 standard for conducted emissions.

Each power supply is suited for use in a broad range of operating temperature environments. They are capable of operating on full power from -10 to +50 °C with a start-up temperature down to -20 °C and derated operation up to 70 °C.

Protection features include over voltage, over current, and over temperature. A remote-sense input and power OK signal are provided across the range. Optional droop current sharing allows multiple supplies to operate in parallel.

With lead-times of stock to 12 weeks, pricing at 1,000-pcs for MVAC400 series products is approximately $130.00 per unit, while pricing for MVAC250 units depending on the model ranges from approximately $102.00.

www.murata-ps.com

DBV architecture saves power in datacenters

By: Patrick Le Fèvre, Marketing and Communication Director, Ericsson

The increasing demand for more internet services and cloud computing is driving both the expansion and building of new datacenters around the world. A key challenge for operators is the minimization of energy expenditure at the board level.

A significant aid to this process is the use of advanced DBV (dynamic-bus-voltage) architectures. Advanced board-power-consumption monitoring and control technologies enable significant energy savings at the board level.

Static bus goes dynamic
Today, the standard power architecture increasingly being used is the IBA (intermediate-bus architecture), which uses IBCs (intermediate-bus converters) to convert a traditional 48-V-DC distribution-level power line used in telecoms typically to a static 12 V DC. This first down-conversion 12-V level feeds a number of POL (point-of-load) DC-DC regulators, which supply the final load voltages at a chip's logic supply levels of 3 V or below.

The choice of 12 V DC has been made to ensure a high enough voltage to deliver all the power required by the board, or load, in times of high data traffic. However, this approach becomes highly inefficient when the traffic demand is low.

The DBV architecture is an evolution of the IBA and provides the possibility to adjust dynamically the power envelope to meet load conditions. It achieves this by adjusting the intermediate bus voltage, previously the 12-V-DC fixed bus voltage. As implemented by Ericsson, the DBV architecture accomplishes this through the use of advanced digital power control and optimized hardware combined with an energy-optimizer series of algorithms. This can lead to reductions in both energy consumption and power dissipation, which in turn contributes to a reduction in the amount of cooling that is required.

DBV is a technology that makes possible to reduce board power consumption from anywhere between 3% and 10%, depending on the board application. The potential for energy saving is a very substantial one, especially when taking into account the fact that 1 W saved at the board level can result in a 3 W saving at the power grid level.

Advanced DBV Conversion Technology
Converters for the DBV architecture can be adjusted within an operation range of 13.0 to 8.2 V and adjusted further down to 4 V to power below-5-V sleep modes, which today requires an extra power-module. The energy optimizer not only optimizes switching parameters to reduce energy consumption, but also offers features including the ability to handle input voltage transients with slew-rates up to 0.5 V/µs, while keeping the output voltage within ±10%. This ensures that the output voltage does not trigger over-voltage protection. It also very efficiently manages pre-bias start-up operation and shut down is fully controlled avoiding voltage spikes that cause avalanche conditions in the secondary-side synchronous rectification MOSFET, making a contribution to further improve reliability.

www.ericsson.com

Power products a big piece of the data-center infrastructure pie

By: Jason dePreaux, Associate Director, IMS Research, IHS

In a recent report, *Data Center Infrastructure* Yearbook, IMS Research forecast that the market for power and cooling products supporting data centers will grow to more than $15 billion by 2014. Power back-up and distribution equipment will comprise over 2/3 of this figure.

Fastest growth is forecast for products that improve data-center efficiency. For example, the need to monitor electricity use is reshaping the power-distribution market with intelligent hardware, which commands higher prices than do their dumb counterparts. Cooling equipment is changing to cope with high-density computing environments. Even enclosures are evolving to facilitate increased airflow and power cabling.

In power, one of the biggest shifts is in rack-level power distribution. Traditionally an afterthought, the trend has been to bring monitoring and switching to these products. PUE (power-usage effectiveness) is becoming a ubiquitous, if not somewhat flawed, metric to compare total data-center power use to server power. Deploying intelligent rack PDU help data-center managers get a better handle on their PUEand what they can do to lower it.

Single-phase units make up more than half of the UPS market. At the large end, static UPS modules are creeping up above 1 MVA to satisfy growing power demands. Transformerless UPSs using IGBTs now account for more revenues than SCR-based units do. These new UPSs offer better efficiency, especially under partial load, where UPSs spend most of their lives.

Future growth in the data-center-infrastructure market will be a balance between opposing forces (**figure 1**): Digitization continues unabated. Mobile data, electronic health records, and internet connectivity drive data-processing and storage needs, demanding more data centers. However, new servers offer greater efficiency, which reduces the need for additional infrastructure. Today, economic conditions often serve as tiebreaker. The cloudy economic picture dampens companies' willingness to spend on capital-intensive projects like data centers.

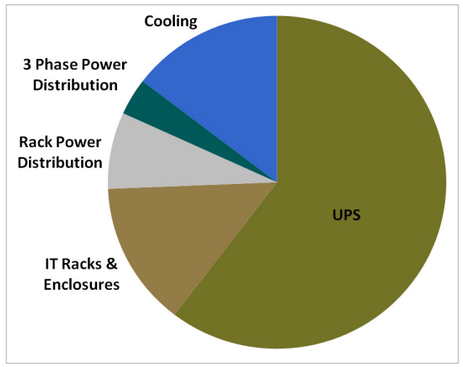

Figure 1: Data Center Infrastructure Market: Revenue Breakdown—2014 (Source: IMS Research)

Still, big multinationals are expanding their reach into data-center products and services. Over the past five years, three vendors have consolidated their positions by acquisition. Schneider Electric, Emerson, and Eaton combined to hold 47% of the 2011 data-center-infrastructure market. These *big three* have each made multiple acquisitions to enhance product portfolios and extend geographic reach. Beyond this are hundreds of vendors with smaller niches around the world.

www.imsresearch.com

Input impedance measurements and filter interactions: Part II

By: Dr. Ray Ridley, President, Ridley Engineering

Introduction

In this article, Dr. Ridley continues the discussion of power supplies with input filters. He shows how the output impedance of the input filter is measured, and demonstrates the importance of input filter damping.

Input Filter Measurements

As discussed in the last article of this series [1], an input impedance measurement gives information about the characteristics of the power supply's input terminals. We use this information in conjunction with measurements of the output impedance of the input filter to assess whether a system interaction is likely to occur. This is important since it can lead to instability of the power supply control loop.

Figure 1 shows a block diagram of a switching power supply connected to an input filter. In order to assess filter interactions correctly, all filter components should be considered to be part of the input filter. This includes any bypass capacitors at the input of the switching power supply, and any other filter components that may be included at the front end of the power supply. In the last article of this series [1], the measurement of the input impedance of the power supply was demonstrated, and the effect of a small bypass capacitor was shown.

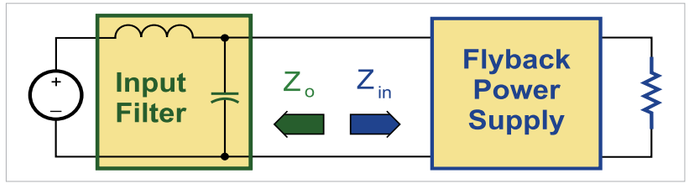

Figure 1: Power supply with input filter module. All filter components, including input bypass capacitors, should be considered part of the input filter.

In this article, the techniques for measuring the input filter itself are presented. The output impedance of the filter can be measured as shown in **Figure 2**. The source from the frequency response analyzer is isolated through a wideband transformer [2], and connected to the input filter in series with a current-sense resistor. If a 1-ohm resistor is used, no scaling is needed on the gain-phase measurement from the analyzer to convert to impedance values.

The node between the sense resistor and the filter under test is connected to a ground reference. Channel A of the analyzer measures the voltage across the resistor (current signal) and Channel B

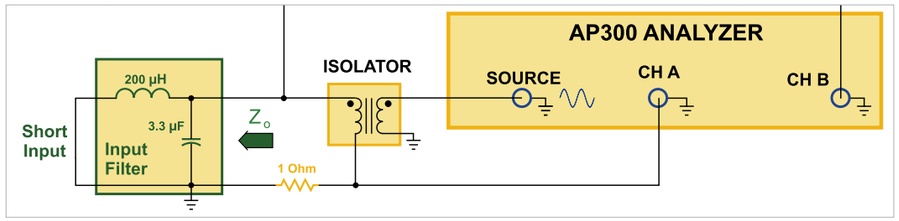

Figure 2: How to measure the output impedance of the input filter. Notice the input of the filter must be shorted for a proper measurement.

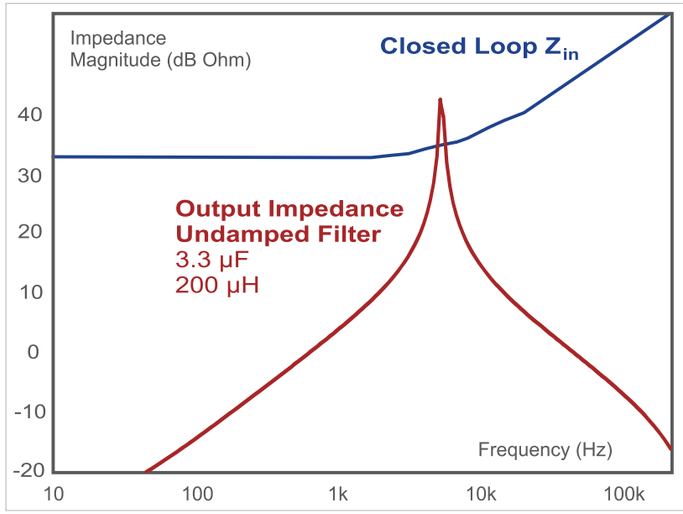

Figure 3: Comparison of measurements of power supply input impedance and input filter output impedance. With an undamped filter, the output impedance exceeds the input impedance of the power supply.

measures the voltage across the filter (voltage signal). The ratio B/A of these two test signals gives the impedance of the input filter.

Notice that a short-circuit is placed across the input terminals of the input filter to properly make this measurement. Notice also that the filter impedance test is usually done without any power applied to the circuit. This greatly simplifies the setup and makes it a safe and straightforward test.

Figure 3 shows the results of impedance measurements for a switching power supply and the input filter.

The blue curve of Figure 3 shows the closed-loop input impedance of the power supply, measured as described in [1]. The effect of the 3.3 µF bypass capacitor has been removed from this measurement, and this component is included in the filter measurements.

The red curve of Figure 3 shows the output impedance of the input filter with the 3.3 µF capacitor, and a 200 µH inductor. The capacitor has a very low ESR, and this produces a filter network with very little damping. At the resonant frequency of about 6.2 kHz, there is a very sharp peak in the output impedance, and its value exceeds that of the closed-loop input impedance of the power supply. There will be dramatic changes in the loop gain of the power supply, and this will be discussed in detail in the next article of this series.

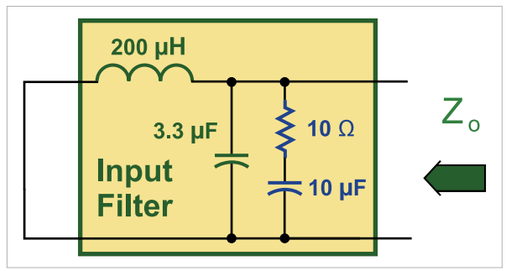

Figure 4: One possible network for damping the input filter. This network applies damping at the filter resonance without affecting the attenuation or dissipation of the filter.

It is not advisable to have such an undamped filter at the input of the power supply, and the filter should be modified to avoid this. **Figure 4** shows one possible network that can be used to damp this input filter effectively. A 10 µF capacitor in series with a 10 ohm damping resistor is connected in parallel with the original 3.3 µF capacitor. The new capacitor branch carries very little switching current and has minimal dissipation. (It is, however, a sizable component when compared to the 3.3 µF MLC capacitor.) This particular damping arrangement is chosen since it provides the same attenuation as the undamped filter at higher frequencies, and filter performance is not compromised for the sake of damping.

Figure 5 shows the effect of damping on the impedances of the power system. If a 50 ohm resistor is used, the output impedance of the filter is reduced to a value just a little less than the input impedance of the power supply. A 10 ohm resistor provides much better damping, and greatly reduces the Q of the input filter. This is a good design choice.

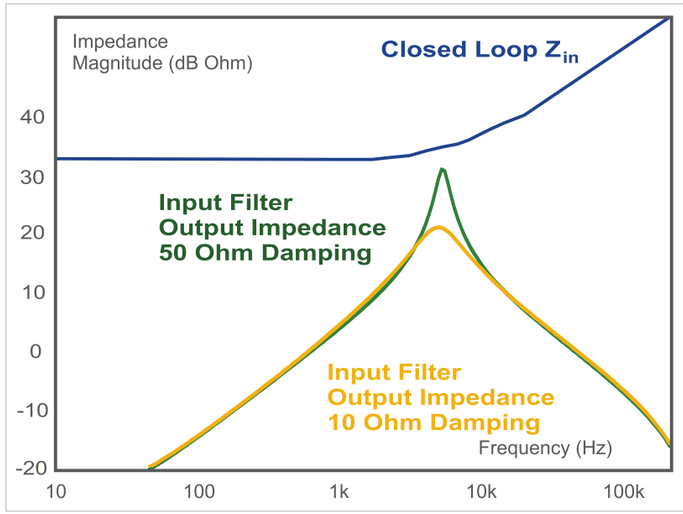

Figure 5: Input filter impedance with added damping network.

The filter configurations used to meet conducted emissions standards of modern switching power supplies will be much more complex than the simple LC filter discussed here. However, damping networks can be applied in a similar manner, and the technique for measuring remains the same. Many power supply designers overlook the needs of damping the input filter, and often underestimate the amount of space needed for damping networks.

Summary

This article discusses how to measure and damp the output impedance of an input filter. This is a critical step for power supply design, and it must be done carefully. It is especially crucial for high density power supplies where capacitors have very small values of ESR, and high-Q filters are often created.

In the next article of this series, measurements of loop gains with and without the input filter in place will show the dramatic effects that result when the filter is not designed properly.

References

1. Ridley Engineering Design Center, www.ridleyengineering.com/index.php/design-center.html, Article [60] "Input Impedance Measurements and Filter Interactions."
2. Ridley Engineering Design Center, www.ridleyengineering.com/index.php/design-center.html, Article [63] "Why You Only Need One Injection Isolator."

www.ridleyengineering.com

Higher power, smaller, and cooler PoL DC-DC regulation

Point-of-Load regulators benefit from innovative packaging

By: Eddie Beville, Power Module Design Manager and Afshin Odabaee, Product Marketing Manager, Linear Technology Corp.

Each generation of high-end processors, FPGAs, and ASICs burdens power supplies with heavier loads. System designers, however, rarely allocate precious additional system-board space to correspond to the power inflation.

Compounding the squeeze on power supplies is the widespread requirement for greater numbers of dedicated board-mount power supplies, which provide POL (point-of-load) regulation for multiple voltage rails. Individual rails must increasingly support from tens to over a hundred amperes at low voltages—as low as 1 V or less—requiring an initial accuracy of ~1% and superb load-transient deviation of less than a few percent. The challenge, therefore, is to find POL supplies that are accurate, can deliver high load currents at low voltages, while taking little system board space.

Once system designers find a suitably powerful regulator, they must evaluate it for power loss and thermal resistance. These two parameters can break an otherwise good regulator design if it can't meet system heat requirements, especially when the system must operate at an elevated ambient temperature.

Obviously, conversion efficiency must be high in order to limit power loss, and the package design must feature low internal thermal resistance and a low thermal resistance connection to the ambient environment. As POL packages shrink, the thermal path between the regulator and the board decreases in area, making it increasingly difficult to keep the board cool because the power regulator usually conducts most of the power loss back into the system board, increasing the internal system temperature.

Heat and cost of cooling

System and thermal engineers spend a lot of time modeling and evaluating these complex electronic systems in order to develop designs that remove power loss in the form of heat. Airflow and heat sinks are typical means to remove this unwanted heat. Compounding the problem, modern processors, FPGAs, and custom ASICs usually dissipate significantly more power as the internal system temperature increases.

This characteristic, unfortunately, results in greater power demand from the power regulators, increasing their internal power loss and the system temperature even further. So high-density power components must limit their losses and remove heat effectively. However, most compact-packaged power devices either dissipate too much power or cannot effectively remove the heat and therefore cannot operate at elevated temperature without significant de-rating.

It's no surprise that to keep the temperature of a high-power design to reasonable levels attention to cooling methods is crucial. System-design constraints have forced designers to implement cooling methods that rely on fans, cold plates, heat sinks, and sometimes submerging the system in special liquids. All are costly, but necessary. However, a high-power POL regulator that can deliver the required power while dissipating heat evenly and efficiently reduces the requirements for cooling that portion of the circuit, saving on cooling size, weight, maintenance, and cost.

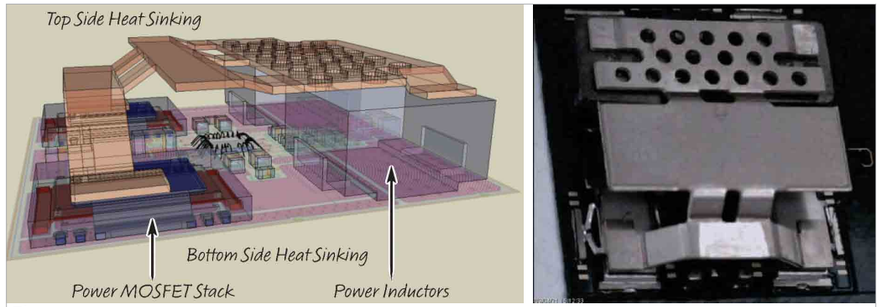

Figure a

Figure b

Figure 1: Side-view rendering (a) and an unmolded LTM4620 (b)

Power density is misleading

The topic of high-power-density DC-DC regulators is misleading because it does not address the device's temperature behavior. System designers should seek more information from the device's data sheet once they decide on a product that meets the system's electrical, physical, and power requirements for a DC-DC regulator. For example: If a DC-DC regulator in a 2 x 1 cm package delivers 54 W to a load, its power density is 27 W/cm². This number may impress a few designers and satisfy their search with their desired power, size, and price. However, what's forgotten is heat, which finally translates into temperature. The key data to study are the DC-DC regulator's thermal impedances: junction to case, junction to air, and junction to PCB.

Continuing with the above example, the device has another attractive attribute. It operates at an efficiency of 90%. It dissipates 6 W while delivering 54 W output in a package with 20 °C/W junction-to-air thermal impedance. Multiply 6 W by 20 °C/W and the result is 120 °C rise on ambient temperature. At 45 °C ambient temperature, junction temperature of the package of this seemingly impressive DC-DC regulator is calculated at 165 °C, which is not acceptable for two reasons: It's above the maximum temperature of most silicon ICs—120 °C—and it requires special attention to keep the junction temperature at a safer value below 120 °C.

The DC-DC regulator that seemed to address all the electrical and power requirements failed to meet thermal guidelines of the system or proved too costly to use due to additional measures to operate at a safe temperature. It's important, therefore, to study the thermal performance of a DC-DC regulator.

Alternatives include, for example, high-density scalable regulators such as the LTM4620 µModule, that combine excellent electrical performance and low power loss with a thermally enhanced LGA package designs to help resolve high-power-density challenges. The LGA package measures 15 x 15 x 4.41 mm. The µModule is capable of driving two independent 13-A outputs, or a single output at 26 A. The package supports both top and bottom heat sinking for excellent thermal performance.

The LTM4620 consists of two high-performance synchronous buck regulators. The input voltage range is 4.5 to 16 V; the output voltage range is 0.6 to 2.5 V and 0.6 to 5.5 V for the LTM4620A. The regulator provides ±1.5% total output accuracy, 100%-tested accurate current sharing, fast transient response, multiphase parallel operation with self clocking, programmable phase shift, frequency synchronization, and an accurate remote-sense amplifier. The regulator provides output over-voltage protection, fold-back over-current protection, and internal temperature-diode monitoring.

Package Design

A side-view rendering and top-view photo of an unmolded LTM4620 shows that the package design comprises a highly thermally conductive BT (bismaleimide triazine) substrate with adequate copper layers for current-carrying capacity and low

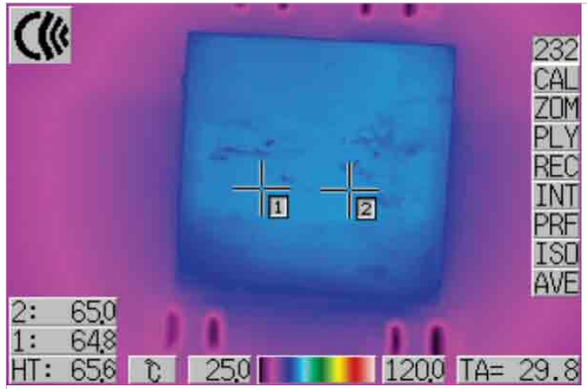

Figure a

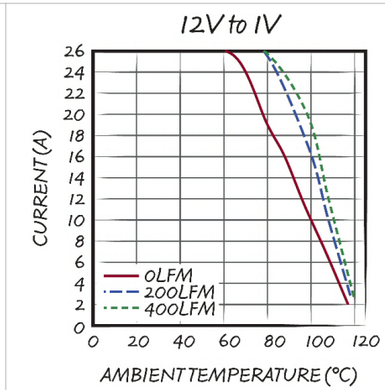

Figure b

Figure 2: Thermal image (a) and de-rating curve (b)

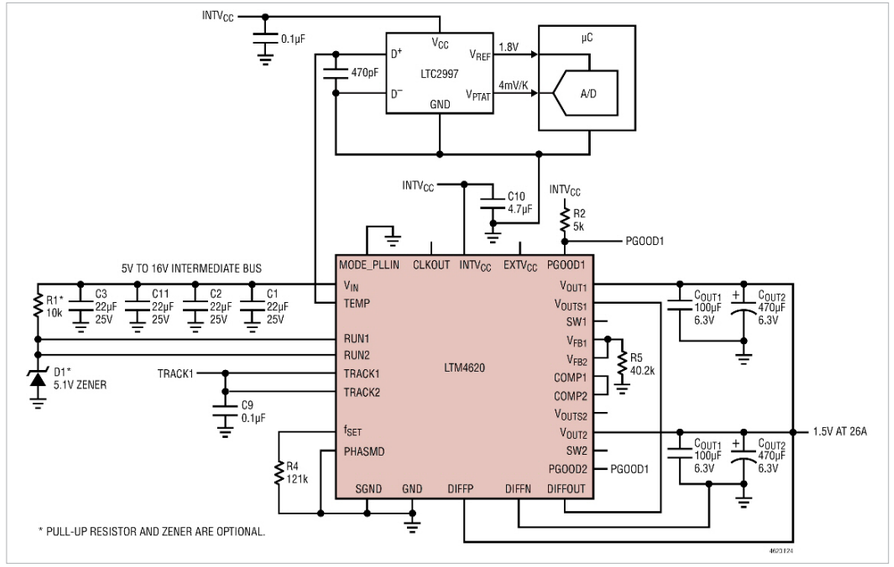

Figure 3: Two-phase parallel outputs combine to deliver 26 A at 1.5 V.

thermal resistance to the system board (**figure 1**). A proprietary lead-frame power-MOSFET stack provides high power density, low interconnect resistance, and high thermal conductivity to both the top and bottom of the device.

The proprietary heat sink design attaches to the power MOSFET stacks and the power inductors to provide effective topside heat sinking. System designs can add an external heat sink to the topside exposed metal to remove heat with airflow. Airflow alone with no heat sink removes heat from the topside due to construction of the integral heat sink and the mold encapsulation.

A thermal image and de-rating curve for a 12- to 1-V, 26-A design shows that the temperature rise is only 35 °C above ambient with 200 LFM of airflow (**figure 2**). The de-rating curve shows that the maximum load current requires no de-rating to ~80 °C. The thermal data confirms the merits of a thermally enhanced high-density regulator.

Electrical Performance

Operating the LTM4620 in current-sharing mode results in a high-density 1.5-V, 26-A design (**figure 3**). This configuration ties the pairs of Run, Track, Comp, V_{FB}, Pgood, and V_{OUT} pins together to implement the parallel operation. The design shows one way of monitoring the device's internal temperature diode using an LTC2997 temperature-sensor monitor. Many different devices that monitor diode-connected transistors can monitor the temperature-sense diode.

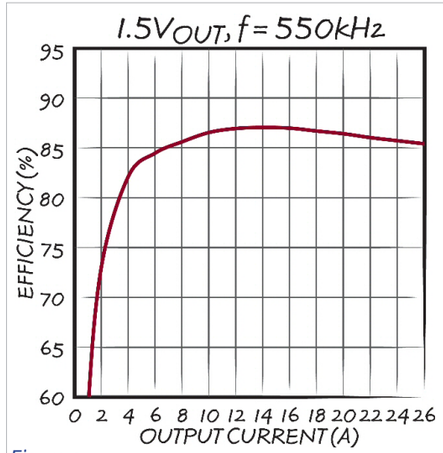

Figure a

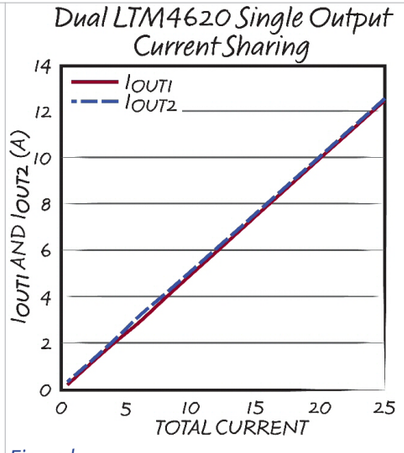

Figure b

Figure 4: Two-phase 1.5-V efficiency (a) and current sharing (b)

The 86% efficiency for the two-phase parallel output and the two-channel current-sharing performance is very respectable for such high-density design (**figure 4**). As shown in figure 2, the temperature rise is modest due to a low Θ_{JA} thermal resistance after board mount. The effective top and bottom heat sinking enables the regulator to operate at full power with low temperature rise.

The high-efficiency and fast-transient-response current-mode architecture fits well with the low-voltage core power-supply requirements for high-performance processors, FPGAs, and custom ASICS. The output-voltage initial accuracy and the differential remote sensing provide proper DC-voltage regulation at the load point. The thermal capabilities and the excellent current sharing allows for scaling the output current capability to > 100 A.

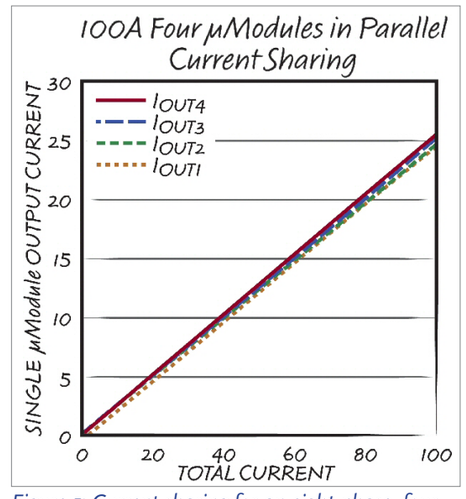

Figure 5: Current sharing for an eight-phase, four-µModule 100-A regulator design

Multiphase operation does not require external phase-shifted clock sources for each regulator channel. Each LTM4620 has Clock-in and Clock-out pins with internal programmable phase shifting for clocking the paralleled channels. Designs can select either external frequency synchronization or internal on-board clocking. These clocking features further enable power scaling.

Extending this design to an eight-phase, four-µModule regulator design delivers 100 A with excellent current sharing (**figure 5**). The 100-A design requires about 1.95 in² of board space. A heat sink across all four modules can remove heat with airflow to minimize dissipation into the system board.

www.linear.com

Making the switch to silicon carbide

Silicon-carbide MOSFETs outperform silicon IGBTs in high-frequency switching

By: Bob Callanan, Applications Manager and Julius Rice, Applications Engineer, Cree

To demonstrate the advantages of silicon-carbide MOSFETs versus silicon IGBTs (insulated-gate bipolar transistors) when used in high-power, high-frequency circuits, it is necessary to employ a test platform that will showcase each device's power handling characteristics—especially switching losses—under hard-switched conditions. A SEPIC (single-ended primary-inductance converter) test platform provides a means to emulate a set of established circuit parameters, such as switching frequency, voltage, current, and temperature. It also provides a means to calculate the differences in device losses between the two different power-switching technologies.

The results of this test platform demonstrate that SiC MOSFETs have up to three-times higher operating frequency with the same efficiency, only 10% of the turn-off losses, and significantly lower relative switching losses as comparably rated Si IGBTs at 30 kHz.

Comparison of switching performance

This demonstration compares a Cree CMF20120D1 1.2-kV, 20-A silicon-carbide MOSFET (**reference 1**) and an Infineon IGW40N120H32 1.2-kV, 40-A silicon IGBT, which is a trench-stop field power device (**reference 2**). The IGBT exhibits a forward voltage at 20 A that closely matches the SiC MOSFET. Both devices were packaged in TO-247 plastic packages and mounted on identical connectorized daughter-boards with equivalent heat sinks, which enabled the two components to be swapped out quickly and efficiently for the test procedure. **Table 1** compares the specifications for the two devices.

Parameter	SiC MOSFET CMF20120D	Si IGBT IGW40N120H3
Breakdown Voltage	1.2kV	1.2kV
Max Current	17A, T_C=100°C	40A, T_C=100°C
Forward Voltage at 20A	1.68V,T_J=25°C 1.98V,T_J=150°C	1.70V,T_J=25°C 1.79V,T_J=150°C
Total Gate Charge	91 nC, V_{GS} = 20V	185 nC, V_{GC} = 15V
Total Gate Energy (V_G*Q_G)	1.82 µJ	2.78 µJ

Table 1: SiC MOSFET and Si IGBT comparative specifications

Selection of the SEPIC platform for circuit comparison

Developing a single test platform to compare the switching performance of these two device technologies required a transformerless DC-DC converter design that re-circulated the load current back into the input link. This avoids the hardware issues that arise with transformers operating at high frequencies, such as proximity and skin effects in the windings and excessively high switching losses. The SEPIC platform delivers the flexibility to buck or boost without inverting the output voltage, making it a good choice to compare the relative switching performance and device losses between SiC MOSFETs and Si IGBTs.

With the SEPIC platform, it is possible to demonstrate the switching performance of the power devices without re-designing and retrofitting each potential application (motor drives, power-factor-correction circuits, grid-interfaced solar and wind inverters, and uninterruptible power supplies) with SiC-based devices and then comparing the overall system performance for each application.

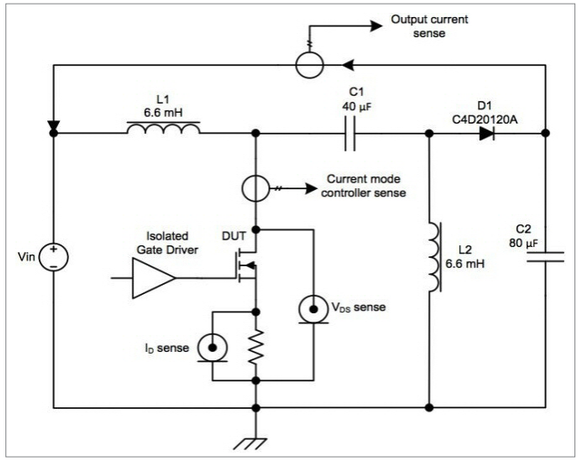

Figure 1: SEPIC demonstration platform schematic

The SEPIC demonstration platform features a simple buck-boost design with the ability to re-circulate the converter's output current back to the input side, while maintaining the devices under test (the power switches) at a duty cycle of approximately 50% (**figure 1**). This results in the voltage across the switch being twice the output current. Therefore, the input DC supply is only required to deliver half the desired switch voltage. Also, by referencing the switches to ground, accurate measurements of their voltage and current are easy to perform.

The SEPIC schematic consists of the DUTs (device under test), which are the swappable power switches being tested; a blocking capacitor, C1; two inductors, L1 and L2; and a diode, D1. As shown by the directional arrows, the output current feeds back into the input source, VIN. A peak-mode current controller feeds from the current transformer to the DUT's drain connection. Meanwhile, the input for this controller is supplied from an amplifier that regulates the re-circulated output current. A Hall-effect sensor monitors this output.

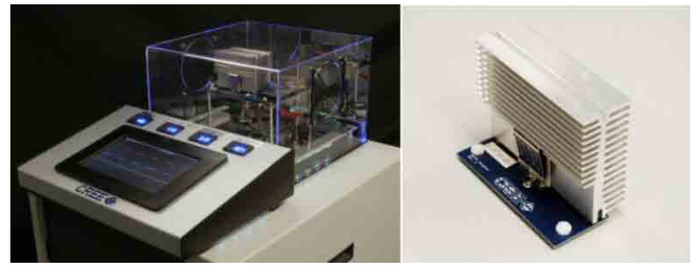

Figure 2: SEPIC demonstration platform display and daughter card assembly

Other diagnostic elements of the SEPIC platform include a high-frequency current-monitoring resistor that senses the DUT current, and a voltage probe featuring a Kelvin connection to measure the DUT voltage. Isolating the gate driver eliminates the potential to form a ground loop around the resistor.

The high-voltage components and logic power supplies used in the SEPIC demonstration platform and the inductors share an enclosure, with the control panel and converter hardware on top of the assembly protected by an acrylic safety barrier (**figure 2**). Test results display on a monitor that shows voltage and current waveforms for the devices under test and meters display total delivered power, total delivered current, total system losses, and input power.

Performance comparison at 30 kHz

Using the input power versus peak switch current for the SiC MOSFET as a baseline, the comparative switching losses are shown by subtracting this data from the input power versus peak switch current of the Si IGBT. The devices were tested under identical conditions of voltage, current, and frequency using the same circuit. The difference in input power between the two devices, therefore, is a direct measurement of their respective switching losses.

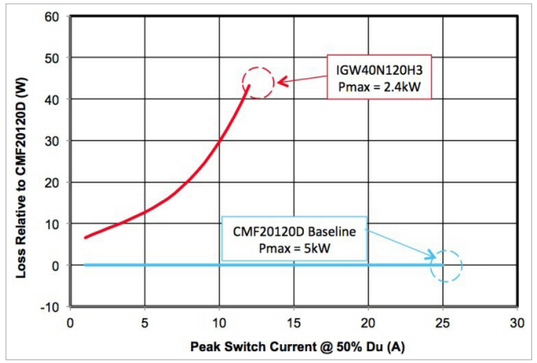

Figure 3: Relative switching losses of 1.2 kV SiC MOSFET versus Si IGBT at 30 kHz; VDS = VCE = 800 V

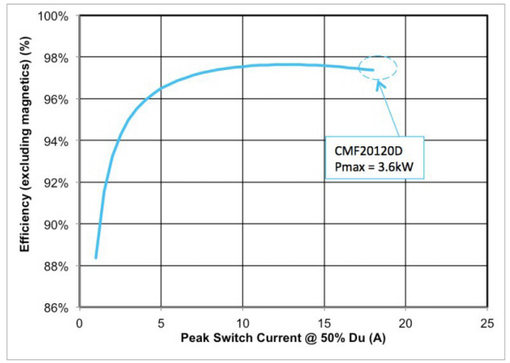

Figure 5: 1.2 kV SiC MOSFET efficiency (without magnetics losses) at 100 kHz; VDS = VCE = 800 V

The Si IGBT's input power loss relative to the SiC MOSFET was plotted against the peak switch current (**figure 3**). For this portion of the testing at 30 kHz switching frequency, the input voltage to the SEPIC circuit was 400 V, resulting in a switch voltage of 800 V. The test terminated when it exceeded the DUT's thermal limits.

The SiC MOSFET achieved the highest switch current—25 A—with 5 kW of delivered power; while the Si IGBT reached a maximum switch current of 12 A with 2.4 kW of delivered power; with the IGBT's switching losses being the limiting factor.

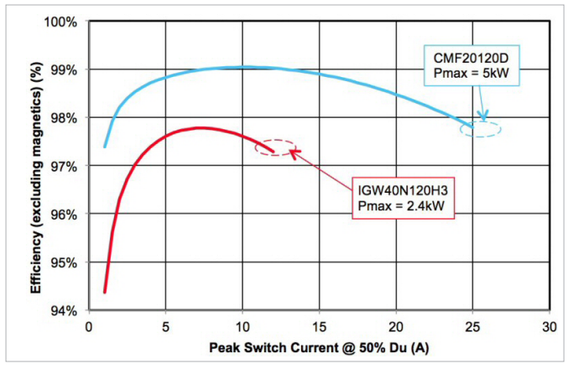

Figure 4: Relative efficiency (without magnetics losses) of 1.2 kV SiC MOSFET versus Si IGBT at 30 kHz; VDS = VCE = 800 V

It is also possible to show the devices' relative efficiency as opposed to relative power losses. The switching devices themselves have low-loss characteristics, although the inductors (as part of the SEPIC circuit) have much larger losses (I^2R); however, these losses are predictable. Therefore, by subtracting out these inductor losses from the total system losses, it is possible to see the relative switching efficiency of the devices under test.

Plotting the switch efficiency versus peak switch current demonstrates that the SiC MOSFET is the higher efficiency device, delivering 5 kW of power; while the Si IGBT delivers just 2.4 kW, with an efficiency about 1.2% lower than the SiC device (**figure 4**).

Switching performance at 100 kHz

While testing at a frequency of 100 kHz, the silicon IGBT exhibited switching losses that are so high that hard-switched operation was impractical. Conversely, the SiC MOSFET achieves a switching current of 17 A, with delivered power of 3.6 kW at this higher frequency operation, achieving both high efficiency and low losses (**figure 5**).

www.cree.com

References

1. http://www.cree.com/products/pdf/CMF20120D.pdf
2. http://www.infineon.com/cms/en/product/PSLPopup.html?productType=db3a3044252fod6f01257e72d2163909

Designing a DiSEqC-antenna phantom power-supply

A simple phantom power supply design allows power, a control link, and the received radio signal to coexist on the coax

By: Andrea Longobardi, Product Definer and Corporate Applications Engineer, Maxim Integrated

The DiSEqC (digital satellite equipment control) standard is a communication protocol developed by Eutelsat used between a satellite receiver (decoder), defined as the master, and satellite peripheral equipment such as dish switchers, LNBs (low-noise blocks), and dish positioners, defined as slaves. The DiSEqC communication system uses only the existing coaxial cable, thus making DiSEqC well suited to reduce cost and improve reliability. The DiSEqC is an open standard with nonproprietary commands.

To allow one-way DiSEqC communication on the antenna cable, the radio head unit must transmit a 22-kHz tone burst, which the remote antenna must receive. The voltage amplitude of this tone burst is 650 mV, as the DiSEqC standard requires. Note that the antenna coax cable also feeds the LNA (low-noise amplifier) and carries the received radio signal. For this reason, the DiSEqC receiver must be able to reject the radio signal on the cable.

Application Circuit

In a DiSEqC application circuit, a dual high-voltage, current-sense LDO/switch is used in LDO mode and the regulated output voltage is dynamically changed to generate the DiSEqC pulses (**figure 1**). The blue block is the radio head unit, which includes the remote antenna power supply used also as DiSEqC tone-burst transmitter and the tuner. The red block is the remote antenna formed by the physical antenna, the LNA, and the DiSEqC receiver (a low-power comparator, such as a MAX9311).

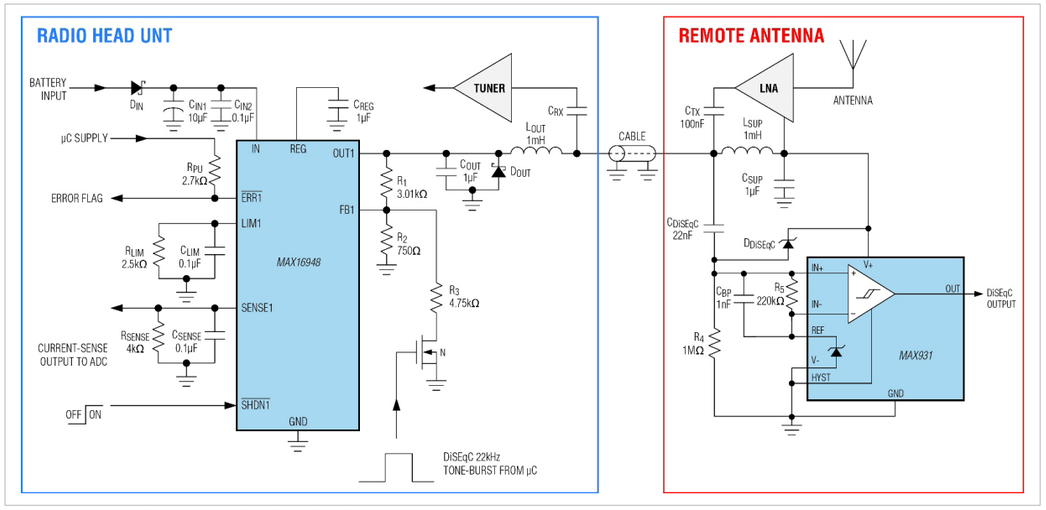

Figure 1: The DiSEqC application circuit

The coax cable enables the communication between the radio head unit and the remote antenna (radio signal and DiSEqC tone-burst) and is also used to feed the remote LNA, saving cost and cable weight.

The automotive LDO/switch is configured in LDO mode with a 5-V voltage output when the external NMOS is turned off (DiSEqC tone-burst off). Choosing resistors R1 and R2 determines this output voltage, as indicated in the device's data sheet (**reference 1**) and related application note (**reference 2**). If a different remote antenna feed voltage (V_{OUT}) is required, use the following equation to choose R1 based on R2:

$$R_1 = \left(\frac{V_{OUT}}{V_{FB}} - 1\right) R_2$$

where V_{FB} is the voltage at the feedback pin in regulation (1 V, nominal) and R2 must be less than or equal to 1 kΩ.

When the external NMOS turns on, resistor R3 connects in parallel with R2. This brings the LDO regulator output voltage to 5.65 V. With this circuit configuration, a user can easily generate a DiSEqC 22-kHz tone burst by turning the external NMOS on and off through the microcontroller. If a different remote-antenna feed voltage is required, choose a value for resistor R3 using the equation

$$R_3 = \frac{R_1 R_2}{\left(\frac{V_{OUT} + 0.65}{V_{FB}} - 1\right) R_2 - R_1}$$

R_{LIM} and R_{SENSE} set the output current limit to 200 mA and the ADC full-scale range to 4 V (**reference 3**). For reasons of clarity, only one channel of the LDO/switch appears on the schematic, but the same considerations are valid for the second channel.

The output inductor, L_{OUT}, is necessary to filter out the radio signal and to not conflict with the LDO regulator. Considering the AM band's lower frequency of 148 kHz, a 1-mH output inductor is sufficient. The tuner extracts the radio signal from the coax cable with a bypass capacitor, C_{RX}. The remote antenna power supply used to feed the LNA is obtained through a lowpass filter built with inductor, L_{SUP}, and capacitor, CSUP. To a first approximation, the power-supply filter is an RLC lowpass filter (**figure 2**). The -3dB passband must be below the frequency used for DiSEqC communication.

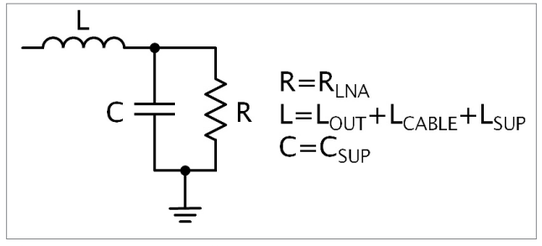

Figure 2: Power supply filter

The low-power comparator acts as the DiSEqC receiver and is supplied with the same voltage supply as the LNA. The negative comparator input, IN-, is connected to the REF voltage provided by the comparator itself, while the positive comparator input, IN+, is polarized with a resistor divider, R4 and R5, in order to have zero voltage output on the DiSEqC output in the absence of a DiSEqC tone-burst.

To sense the DiSEqC tone-burst, a 22-nF capacitor, C_{DiSEqC}, connects between the coax cable and the negative comparator input. When a tone-burst is sent on the coax cable, the IN+ voltage exceeds the IN- voltage, which generates a pulse on the output of the comparator. A protection Schottky diode, D_{DiSEqC} connects between IN+ and the comparator power supply input, V+, to avoid an over-voltage fault on the IN+ pin.

To avoid false output pulse triggering due to the radio signal travelling on the cable, a 1-nF bypass capacitor, C_{BP}, is placed between the comparator inputs, from IN+ to IN-. The radio signal received from the remote antenna and amplified with the LNA is injected onto the cable with a capacitor C_{TX}.

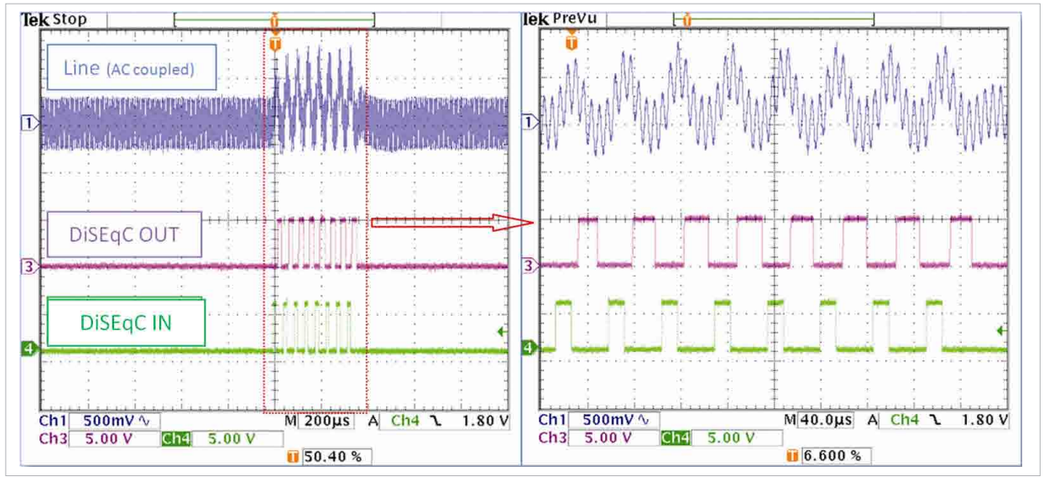

Figure 3: RF signal at 148 kHz and 500-mV amplitude

Bench Tests

Bench tests were performed, which generated eight 5-V-amplitude 22-kHz tone-bursts with a waveform generator connected to the gate of the external NMOS. A sinusoidal 500-mV amplitude RF signal was obtained with another waveform generator and injected with the CTX capacitor, emulating the LNA output radio signal. An oscilloscope monitored the output of the comparator to confirm whether the sent tone-bursts were received and to ensure that the injected radio signal did not influence the DiSEqC communication.

Figure 3 and **Figure 4** illustrate the scope traces corresponding to the performed tests. Figure 3 shows the results with an injected radio signal at 148 kHz and 500-mV amplitude, which coincides with the lower AM band frequency.

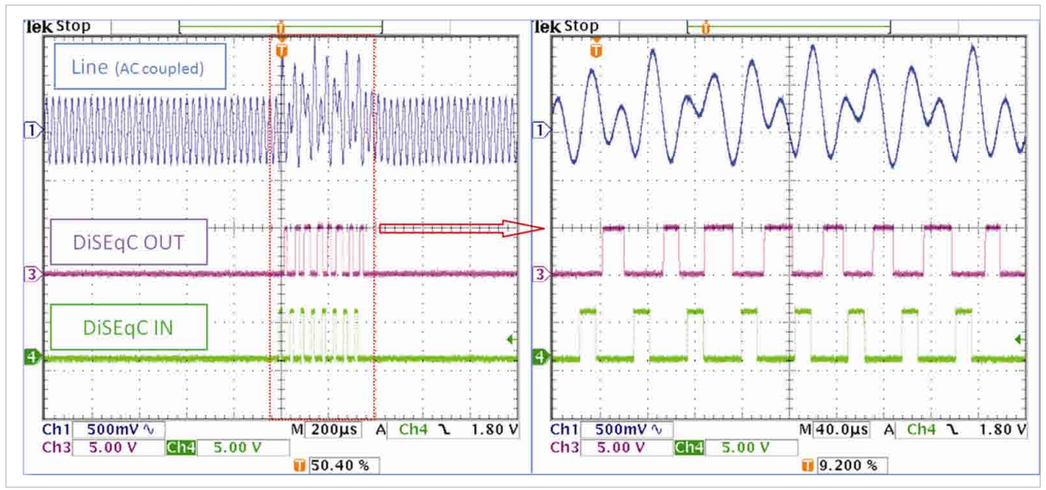

Figure 4: RF signal at 37 kHz and 500-mV amplitude

Figure 4 shows the results with an injected radio signal at 37 kHz and 500-mV amplitude, which is the second subharmonic of the lower AM band frequency, 148 kHz,.

The DiSEqC application circuit is a low-cost and flexible design for an antenna phantom power supply that is compatible with the DiSEqC communication standard. Additional bench test results have confirmed that DiSEqC communication still operates when selecting the DiSEqC tone-burst frequency in the range between 100 Hz and 30 kHz. This provides the flexibility to tune the most suitable frequency for DiSEqC communication, thus minimizing interference with other RF signals on the coax cable. Product designers can also regulate the tone-burst duty cycle and add hysteresis to the comparator to attain the best DiSEqC communication performance.

This application circuit enables one-way DiSEqC communication. If a receive-acknowledge signal is needed from the remote antenna, it can be generated by modulating the load current of the LDO/switch. A simple way to do this would be to connect an extra load in parallel to the LNA supply inside the remote antenna once the DiSEqC message is received. The microcontroller in the radio head unit can receive the acknowledgment by sampling the load current variation on the MAX16948's current-sense output, SENSE. The switched extra load could be easily built up with an NMOS switch in series with a pullup resistor connected to the LNA supply inside the remote antenna.

www.maximintegrated.com

References

1. MAX16948 data sheet: http://www.maximintegrated.com/MAX16948
2. Maxim application note 5271, *Selecting External Components for an Automotive Dual Remote Antenna Current-Sense LDO/Switch*, http://www.maximintegrated.com/AN5271.
3. Maxim application note 5271, *Selecting External Components for an Automotive Dual Remote Antenna Current-Sense LDO/Switch*, http://www.maximintegrated.com/AN5271.

Thermal design considerations for thin power MOSFETs

Applications with high pulse currents benefit from thin, low RDS(on) MOSFETs but demand care in thermal design

By: Ralf Walter, Application Engineer, Infineon Technologies AG

Thermal challenges accompany almost all advances in power-electronic designs. Load cases not covered under a datasheet's parameters are often particularly problematic.

Each new MOSFET generation generally marks a significant reduction in the area-specific on-state resistance—the product of $R_{DS(on)}$ and area. This means that a MOSFET with the same chip size offers dramatically reduced $R_{DS(on)}$ figures.

Many optimization steps are required to achieve this outcome including altering the cell structures and scaling individual areas. The thickness of the silicon is usually also reduced, which can bring about shifts in thermal behavior. The reduction in silicon thickness lowers the thermal resistance, also reducing the amount of silicon available for temporary high loads.

The challenge to the developer is to calculate these changes in the thermal behavior for each specific application and to optimize the design as appropriate. Datasheet values and the PSPICE library may not provide sufficient indicators for special cases, such as temporary loads, and analyses carried out using a thermal camera for such fast processes do not provide reliable values for the temperatures within the power device's package.

Thin MOSFETs, such as Infineon's 40- and 60-V OptiMOS power MOSFET generation, feature on-state resistances to about half of the previous generation's values without changing chip sizes. This was only possible by reducing the thickness of the silicon. Chip thickness values in this voltage range are typically still in tenths of millimeters.

Besides conventional switched-mode power supplies (such as synchronous rectifiers on the secondary side), typical applications in these voltage classes include motor controls in cordless power tools like screwdrivers and similar products. These designs must withstand high loads for short periods as they are often subject to such peaks under real operating conditions. It is therefore completely normal that a multiple of the rated current is permissible for a second, for example, when fully and securely driving a screw into a material.

The resulting losses in the MOSFETs are briefly much higher than in rated operation. One possible design choice would be to reduce the losses by massive parallelization. However, this is not practical (or even possible in some cases) for cost and space reasons. On the other hand, such overload cases are *one-time* events from a thermal point of view since a repetition is not likely until many seconds later (for the next screw).

Now, however, it is possible to buffer the power loss occurring during this short peak load and to dissipate it slowly over the lengthy period before the next power impulse. Silicon and the copper lead frame

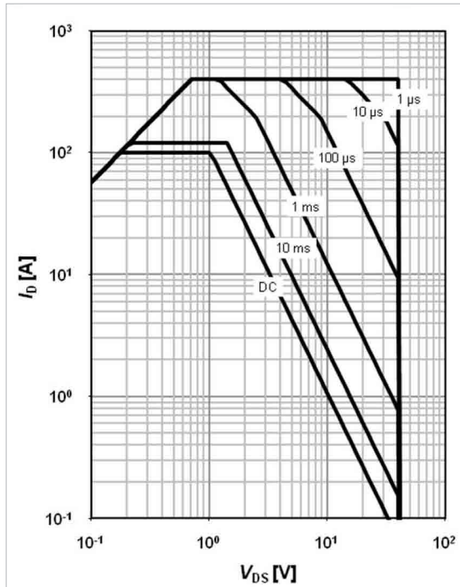

Figure a

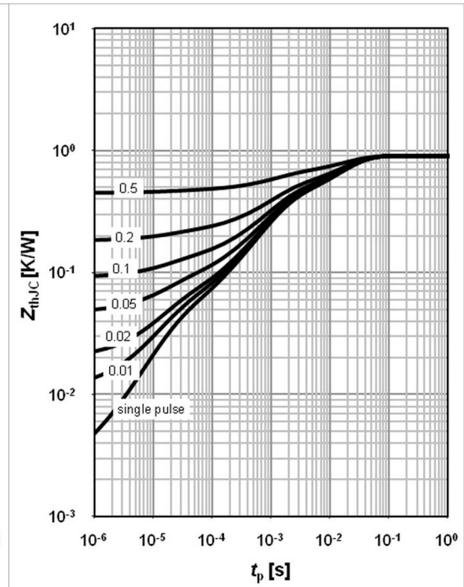

Figure b

Figure 1: SOA (a) and Z_{thJC} (b) charts for the BSC014N04LS

are ideal for this purpose. However, increasing levels of miniaturization have greatly reduced both structures in size already. In particular, high-performance packages, such as the SuperSO8, have fewer options for storing thermal energy due to their small volume.

Manufacturers can use special IMS (insulated metal substrate) PCB designs for such SMD components. These consist, in principle, of an aluminum or copper substrate—usually between 0.5 mm and 3 mm thick—with a thin laminated insulation layer of enriched epoxy. The copper conductors and the SMD components are located on that layer. The metal substrate can operate as a *thermal tank* for storing thermal energy.

Thermal performance

A useful starting point is an analysis of the datasheet areas that describe thermal behavior in relation to time and power dissipation. For example, the SOA and Z_{thJC} charts (**figure 1**) depict the performance of a BSC014N04LS 40-V MOSFET that offers an RDS(on) of 1.4-mΩ in a 5- x 6-mm SuperSO8 package.

As the figure depicts, the MOSFET reaches a quasi-static state within 10 ms. An examination of thermal equivalent circuit diagram for the different silicon generations indicates that the chip volume plays a role only for short events below 100 µs. Even for very short load impulses, thermal energy can no longer be stored in the chip, but must instead dissipate first to the solder and lead frame and, after several milliseconds, to the environment. This means that chip volume does not play any measurable role in the case of impulses lasting at least one second. This is also clear when regarding the thermal properties of the materials—silicon, solder, and copper—in combination with their volume. This thermal behavior also reflects in the equivalent circuit diagram or PSPICE parameters.

A thermal-equivalent circuit can depict the thermal system including the substrate (**Figure 2**). In this example, the BSC014N04LS in its SuperSO8 package combines with the properties of an acceptably priced IMS. Here the specific thermal conductivity of the IMS's

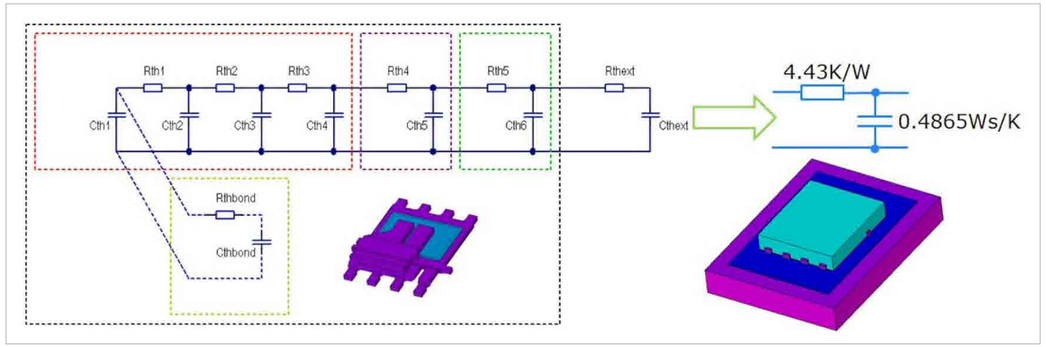

Figure 2: BSC014N04LS on IMS with thermal equivalent circuit diagram of the complete structure

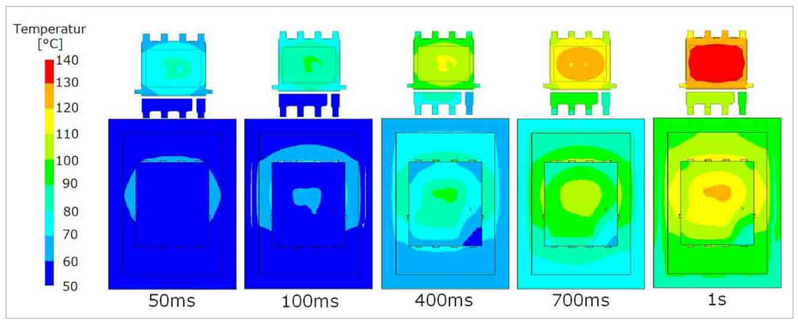

Figure 3: Bottom: View of the entire system including housed MOSFET. Top: View of the underside of the BSC014N04LS

insulation layer is 1.3 W/(m x K) with a layer thickness of 76 µm.

Calculating the system's temperature rise results in a value for theoretically homogeneous energy distribution, for which a ΔT of about 51 K is generally acceptable. In a real system, however, the time constants play a major role. Purely mathematically,

$$R_{th}C_{th} = 4.43\frac{\mathrm{K}}{\mathrm{W}} \cdot 0.4865\frac{\mathrm{Ws}}{\mathrm{K}} = 2.155\,\mathrm{s},$$

and non-homogeneous distribution of the temperature is to be expected. The real world behavior—the heating of the individual components in the system—is very much removed from the ideal model.

The finite-element method offers a way of obtaining realistic values for each individual point of such a system, and of visualizing the values. **Figure 3** shows how the temperature changes within one second for an assumed 25-W load. This corresponds to considerably more than 100-A drain current with 100% turn-on as a typical excess load that occurs when driving in a screw.

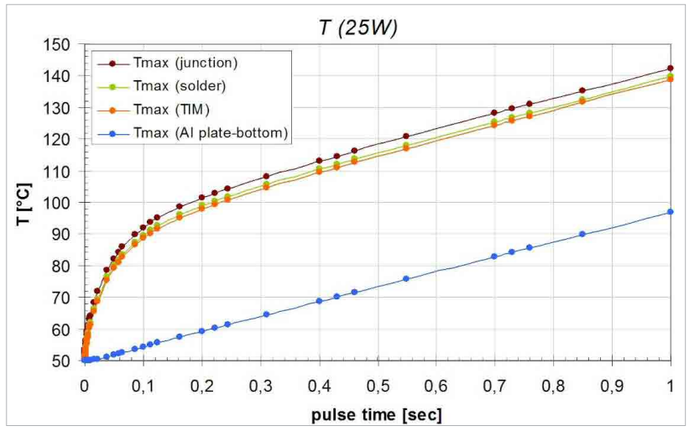

Figure 4: Temperature behavior over time at selected points

After several hundred milliseconds, the MOSFET's plastic housing becomes hot while the aluminum profile remains noticeably cooler due to non-homogeneous. After the second expires, the outer housing temperature of 120 °C has risen more than considerably and has reached the vicinity of the maximum permissible value. The underside of the component (top row of images) reaches a temperature near 140 °C. The temperature behavior can be seen in greater detail in the timing diagram (**Figure 4**).

As a result of the MOSFET cells' good thermal connection to the copper lead frame, the maximum temperature difference between the chip and lead frame is only about 4 K (dark red and orange lines). Also clear is a short nonlinear interval—up to about 100 ms—after which all temperature behaviors are more or less linear. This is due to the previously mentioned short time constants within the component, as Figure 1 indicated earlier. Figure 4 also shows that these findings can easily transfer to another situation (power loss, turn-on relationship, starting temperature, impulse length, etc.) with sufficient precision. In the case of lower power losses or turn-on ratios, the line moves down in the according ratio. A shift upwards corresponds to greater power losses. Given the linear characteristics beyond about 100 ms, extending the lines displays longer impulse lengths.

The 143-°C results for the example shown are close to the BSC014N04LS's permissible thermal limits—$T_{J(max)}$ = 150 °C. The fact that the simulation assumes worst-case values offers additional security. It would be problematic, of course, if, for example, the impulse lengths were greater or if the power losses were more than 25 W because the reliability of the application could suffer under certain conditions without a change in the setup. An unacceptable reduction in the lifetime of the chip would be the consequence.

Regarding the task in detail, there are a number of possibilities for improving the situation. Due to the non-homogeneous temperature distribution on the aluminum profile, it does not make much sense to increase it in size. The mounting volume may then no longer be acceptable and the costs would increase.

It is much more advantageous to increase the thickness of the profile, which can greatly increase the effectiveness of the thermal tank for this brief impulse. However, this can result in cost and space disadvantages. Yet another design alternative would be to use an IMS material with superior thermal conductivity. The sharp rise in costs for this special material represents a drawback though.

A further possibility is careful selection of power devices from the extensive range of available components. The developer can select an optimum combination, balancing cost and performance.

In the setup shown above, a MOSFET with a low $R_{DS(on)}$, such as the 1-mΩ BSC010N04LS, proved to be a good choice with the added bonus of enhanced switching performance. Also available are 40-V chips with monolithic integrated Schottky-like diodes such as the BSC010N04LSI, which offer advantages, in particular, for fast switching.

When assessing modern silicon technologies, consider whether the improvements in certain parameters, such as $R_{DS(on)}$, will result in disadvantages in other areas including thermal management. As the example given here shows, a reduction in the thickness of the silicon layer does not present a disadvantage in this or similar applications. Instead, the reduced thickness even lowers the thermal resistance between the MOSFET cells responsible for power loss and the copper lead frame, and the thermal capacity only suffers to an insignificant degree. The latter property only has an effect in the microsecond range in any case.

www.infineon.com

New technology and acquisitions benefit wireless sensing

Maturing WSN technologies from established sources ready for mainstream applications

By: Harry Forbes, Senior Analyst, Automation, ARC Advisory Group

Wireless sensing is the most important new technology in process measurement to appear in decades. For this reason alone, it has attracted the sustained attention of ARC Advisory Group and other industry analyst organizations. While growth of wireless measurements in industrial applications has proven to be slower than most forecasts, wireless continues to grow much faster than the overall automation market.

Looking at the changes in WSN (wireless sensor networking) during the past year, ARC is most impressed by the greater maturity of the industry along with significant new technical achievements. We believe that both will spur market growth because they enable WSN to serve a much larger set of customers. The events of the last year at WSN leader Dust Networks provide an excellent example of how these developments have changed—and will continue to change—the wireless-sensing market.

Commercial Maturity

Greater WSN industry maturity has come through a steady series of acquisitions. At present, much larger companies have acquired almost all the pioneering sensor-networking ventures. Most of the acquirers are well-established semiconductor suppliers, which have now snapped up virtually all the fabless WSN firms. Dust Networks has become one of the more recent, acquired in late 2011 by specialist Linear Technology.

Fortune 1000 companies must carefully manage the technologies they employ if they are to meet their customer expectations. A technology sourced from a group of venture-stage fabless semiconductor firms represents a major supply-chain risk to global firms whose customers expect product support for many years going forward. Today the list of WSN suppliers no longer represents an exceptional risk to the supply chain of global giants like Emerson, GE, Honeywell, and Siemens (**figure 1**).

Figure 1: Emerson Rosemount's 3051S series wireless instrumentation exemplifies the benefits WSN brings to industrial applications: It can reduce total installation costs by 45% and total deployment time by as much as 75%, eliminating design and installation time and cost associated with wiring, conduit, cable trays, and junction boxes.

For Dust Networks, besides joining a $1.5 billion chip company with global presence, the acquisition has enabled Dust to adopt a mature testing and commercialization process. Linear uses this internal rigor to ensure the high quality it uses as a differentiator. Roughly, 40% of Linear's business is with industrial customers, so Linear is familiar with their demands. Similarly, Dust Networks chose to focus on industrial WSN applications from its earliest days (**figure 2**). Through dialog with industrial customers, Dust developed its product mantra of *low power and reliable delivery*. These two attributes have governed Dust's design decisions for years and skillfully pursuing both goals accounts for Dust's market-share leadership in industrial WSN applications.

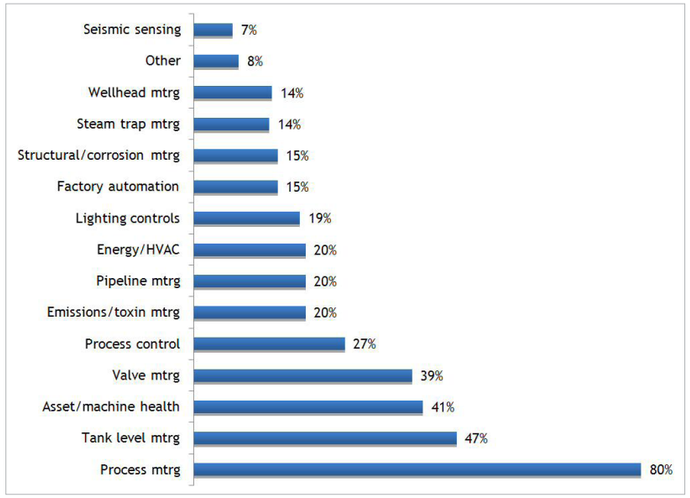

Figure 2: Industrial applications for WSN products. (source: OnWorld, 2012)

Dust has adopted Linear's processes for its latest product release. The Dust acquisition, then, is not merely a plan to grow volume through wider distribution. Rather, it also adds value through improvements to Dust's internal development and commercialization processes.

Technical Growth

WSN has also grown technically during the past year. The most important development is the completion of the *e* revision to the IEEE 802.15.4 standard. Though the 15.4 standard published years ago, most commercial WSN applications have used non-standard MAC (medium access) rules to optimize the performance of their own networks. WirelessHART, ISA100.11a, ZigBee, and IPv6 sensor networks all depended on a customized MAC layer to achieve low power consumption and reliable end-to-end message delivery. This limits the value of the IEEE standard, because applications needed their own specialized MACs.

IEEE 802.15.4e makes important changes to the defined MAC layer. It creates a standard and fully defined MAC that can support diverse types of networks. This includes 6LoWPAN-compressed IPv6 networks. It also supports the synchronized TDMA (time division multiple access) network properties used in industrial low-power applications such as WirelessHART and ISA100.11a. Finally, 15.4e accommodates extensions, so that these diverse networks—as well as future ones—can extend the standard MAC without violating the standard itself. This decouples WSN development from the three- to four-year long IEEE standards-development process. It will give the IEEE standard greater value and a much longer life, enable greater interoperability among WSN silicon and network stacks, and enable future WSN technologies to leverage an existing and fully standard MAC.

Dust Networks and 15.4e

The new standard incorporates many of the technology enhancements that Dust Networks discovered and championed from its earliest days, and has continued to use right through the newest SmartMesh WH and SmartMesh IP products (**figure 3**). But, the new standard was developed not only from Dust's technology but also with extensive and enthusiastic contributions from many major suppliers, notably including Siemens.

Dust now employs its most recent chip design, which conforms

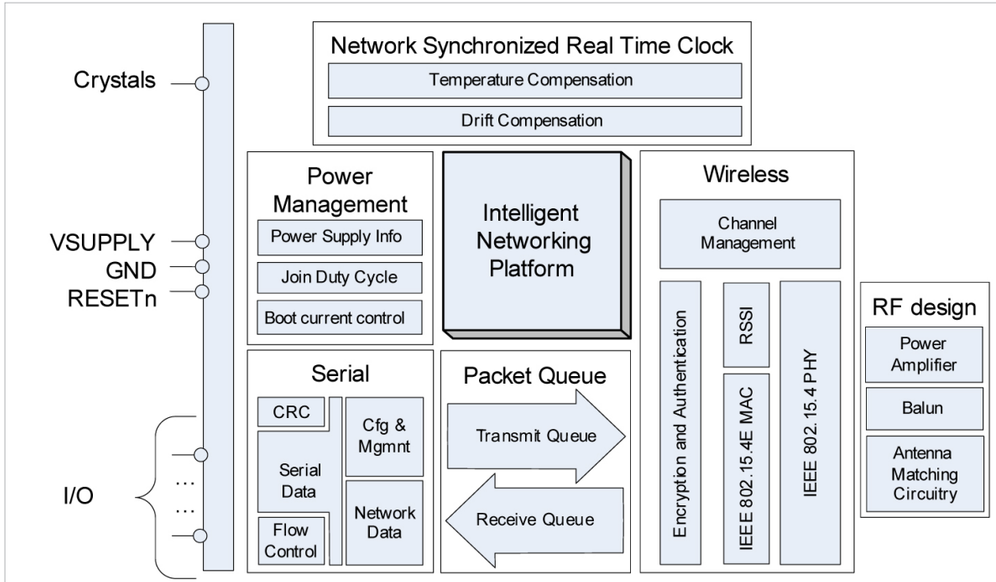

Figure 3: Block diagram of the LTC5800-WHM SmartMesh WirelessHART MoC (mote on chip).

to 15.4e, in both WirelessHART and IPv6 products. Using the latest chipset is a big benefit for WirelessHART industrial customers, since Dust has driven down the power consumption of its chipsets by roughly 50% with each new generation. In battery-powered industrial devices, this provides an *energy surplus* enabling devices to have longer life and to serve in applications that previously consumed too much battery power.

The ability to serve both WirelessHART and IPv6 applications from a single chip design also positions Dust for growth in both markets. IPv6 applications can use a wider variety of system architectures. In particular, they can route packets from field sensors to analytics and applications that are truly location-independent, including cloud-based applications. It remains to be seen what impact the Cloud will have on the wireless sensing market, but the combination of high scalability and low cost offered by cloud services may well enable new applications that are not feasible when built using traditional WSN gateways.

Dust Networks continues to focus on products that lead the market in terms of low energy consumption. This lengthens the life of devices in industrial applications, which is particularly valuable to both industrial OEMs and their customers.

Dust Networks also uses its huge experience base to optimize its network manager. In most WSN applications, the network manager is both extremely active and mission-critical. Active network management permits a sensor network to achieve highly reliable end-to-end service combined with low energy use. This is possible only because the network manager continuously makes decisions that dynamically reconfigure the network so that it operates reliably over point-to-point links that are inherently unreliable.

Another area to watch may be energy harvesting. Though energy harvesting represents an ongoing technical challenge, it is a very active development area. Low power products will be the first to benefit from any new harvesting components that emerge.

www.arcweb.com

IEDM, where the device is king

but the material base suggests new approaches to manufacture

By: Gail Purvis, Europe Editor, Power Systems Design

With a 60-year pedigree the IEDM (IEEE International Electron Devices Meeting) is that pre-eminent forum where the world reports technological breakthrough in the areas of semiconductors, electronic-devices technology, design, manufacturing, physics, and modeling. It is the realized R&D dream, the first steps to producing eventually the electronic components, devices, systems, and applications of the biennial Electronica, now in its 25th year, and alternating with Productronica manufacturing slant.

IEDM has also become the flagship conference for nanometer-scale CMOS transistor technology, advanced memory, displays, sensors, MEMS devices, novel quantum, phenomenology, optoelectronics, power, energy harvesting, and high-speed devices, not to mention process technology, device modeling, and simulation.

This year in December at San Francisco, those who have the thrill of attending will see an increased emphasis on circuit and device interaction. With the ever-increasing transistor count, nanometer design rules, and layout restrictions, circuit-device interaction is becoming highly critical to provide viable technology.

This is seen as one new emphasis that includes technology and circuit co-optimization, power to performance to area analyses design for manufacturing and process control, as well as the better known, more traditional CMOS platform technology with its continuous scaling.

Among some 220 presentations, which of course includes Intel unveiling its *hot* trigate manufacturing technology and Europe's imec presenting a plenary talk, two invited

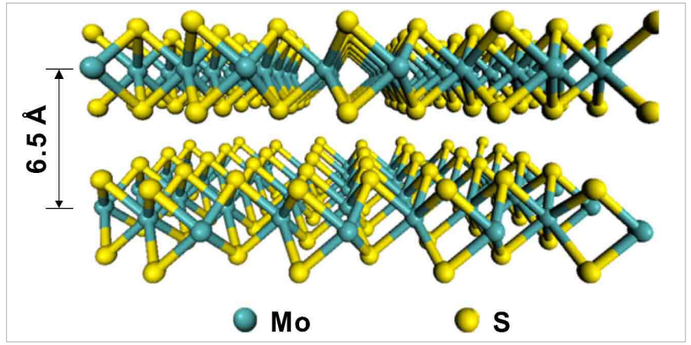

Figure a

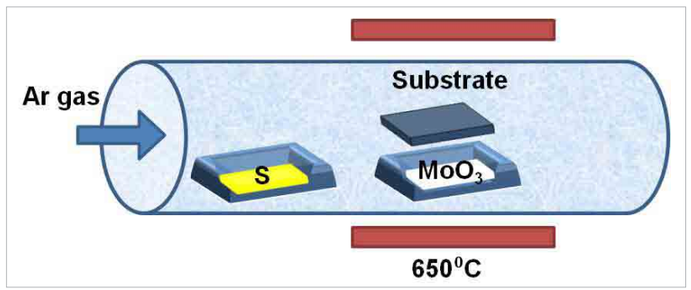

Figure b

Figure 1: The MoS lattice structure (a) and a schematic of the CVD process for growing single layer MoS, which shows great uniformity and coverage (b).

presentations, one tutorial, and nine papers with imec researchers as first authors and co-authored, there are several new material straws revealed for future devices.

MoS (molybdenum sulphide) is one of the new entrants cheerfully determined to nudge grapheme slightly off its pedestal. A 2D material, MoS has grapheme characteristics, but unlike grapheme, it also has a wide energy bandgap, which enable fabricators to build transistors and circuits directly from it.

MIT researchers will be describing the use of that traditional workhorse CVD (chemical vapor deposition) process to grow the uniform, flexible, single molecular layers of MoS, comprising a layer of Mo atoms sandwiched between two layers of S atoms (**figure 1**).

They then exploited the material's 1.8-eV bandgap to build MoS transistors and simple digital and analog circuits, namely a NAND logic gate and a 1-bit ADC converter. The transistors have demonstrated record MoS mobility > 190 cm^2/Vs, an ultra-high on-off current ratio of 108, record current density of ~20 μA/μm, and saturation, as well as putting on the first ever GHz RF performance from MoS. These results would seem to show MoS as suitable for mixed-signal applications, and useful for those that require both high performance and mechanical flexibility (**reference 1**).

Project	Pros	Cons
Darpa SyNAPSE (IBM)	· 256 neuron chip. · 256k CMOS SRAM synapses. · Received substantial SyNAPSE funding.	· Axon–neuron crossbar does not scale well. · Interconnection problem
LETI	· PRAM synapse. · Pattern recognition.	· Crossbar may not scale well.
SpiNNaker (Manchester)	· Connection info on chip, used for routing. · Uses and expands on silicon, architecture, networking technology. · Easy to experiment.	· Crossbar may not scale well.
FACETS (Heidelberg+)	· 100k times faster than real neurons. · 100k CMOS synapses. · HICANN *analog network chip.*	· Interconnection problem unsolved.
Blue Brain (EPFL Switzerland)	· 100k neurons with BlueGene. · STDP implemented. · Most bio-realistic. No hardware required. Software easy to change.	· Expensive simulation.
Neurogrid (Stanford)	· Programmable neurocore chip (1k neurons). · Synaptic addresses stored in RAM.	· Limited scalability to date. · Degree of neuron is unclear.
INC IFAT chips (UCSD)	· HiAER IFAT chips (65 neurons). · FPGA router for AER	· AER router tree bandwidth
This work	· 1k RRAM synapses · $4F^2$ is possible · Connection CMOS neuron to RRAM synapses	

Table 1: Comparison of Artificial Brain Projects

PCM (phase change memory) has frequently been a star performer at IEDM over the last handful of years. For 2012 however a neuromorphic, or brain-like, electronic system that mimics cognitive functions is the current showpiece and focus of research through its potential for complex tasks such as pattern-recognition.

IEDM papers in 2011 described studies using programmable PCM synapses in neuromorphic systems to carry out a function STDP (spike-timing-dependent plasticity). STDP is an electronic analog of a brain mechanism for learning and memory.

An electronic system that accurately performs STDP can be said to be *learning*. This year, a team led by Korea's Gwangju Institute of Science and Technology details a high-speed pattern recognition system; comprising CMOS *neurons* and an array of RRAM (resistive-RAM) based *synapses*, which also demonstrate STDP (**reference 2**).

The 1-Kb RRAM array has a simple cross-point structure and may be scalable to 4F—the theoretical minimum size for a cross-point array. The work shows the feasibility of using neuromorphic architecture for high-speed pattern recognition and charts the work in a comparison of artificial-brain projects (**table 1**).

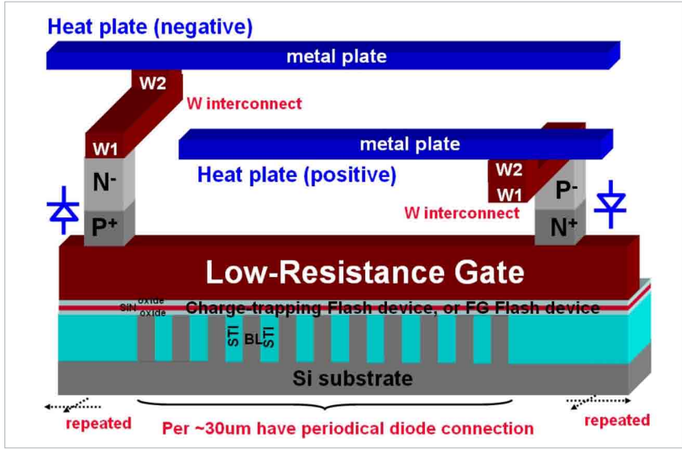

Figure 2: A schematic showing the structure of the diode-strapped wordline. A PN diode can be formed directly on top of the wordline, and local interconnect can be used to connect to the metal heat plates.

Without memory, much of the electronics business would be missing, but flash memory lifetimes are currently limited in use, because repeated P-E (program-erase) cycles degrade the tunnel oxide, which insulates the flash memory cells. In principle, while heating the oxide repairs the damage, this thermal annealing has been impractical, as flash memories do not tolerate high temperatures and long baking times.

Now it appears that Macronix researchers have managed to build flash memories that could self-heal, by using tiny on-board heaters that provide thermal annealing, just at the spots where it is needed.

They have modified the wordline from a single-ended to a double-ended structure, which in turn enables current to be passed through the gate to generate Joule heating (**figure 2**). High temperatures, > 800 °C, accordingly were generated, but only in the immediate proximity to the gate. The devices demonstrated record-setting endurance of > 100 million P-E cycles and with excellent data retention. Interestingly, is thought to be temperature-independent (**reference 3**).

That other electronics Holy Grail, the quest for flexible circuitry with its promise innovative biomedical, security, wearable, and other uses to date finds plastic substrates not compatible with the high temperatures and harsh processes needed for high-performance CMOS devices. It now that looks to be in IBM's remit, as researchers have developed logic on plastic with high performance CMOS circuits, including SRAM memory and ring oscillators, on a flexible plastic substrate.

Known as ETSOI (extremely thin silicon-on-insulator) devices, these have a thickness of just 60 angstroms. IBM built them on silicon, then used a simple, low-cost room-temperature process called *controlled spalling*, which essentially flakes off the Si substrate (**reference 4**).

Engineers then transferred them onto flexible plastic tape. The devices had gate lengths of < 30 nm and gate pitch of 100 nm. Ring oscillators had a stage delay of just 16 ps at 0.9 V, and this is believed to be the best-reported performance for a flexible circuit.

A slight degradation of delay for the flexible sample, after the layer transfer, comes from degradation of p-FET performance due to strain effects. The final 100-mm-diameter flexible ETSOI circuit is on plastic (**figure 3**).

The unreleased MEMS (micro-electro-mechanical systems) resonator, another M.I.T. device, might in passing seem to be just a curiosity, but it looks to be very helpful player in the communications and timing sectors. A common task in circuit design is to generate reference frequencies for timing and communications purposes, and one way to do this is via vibrating

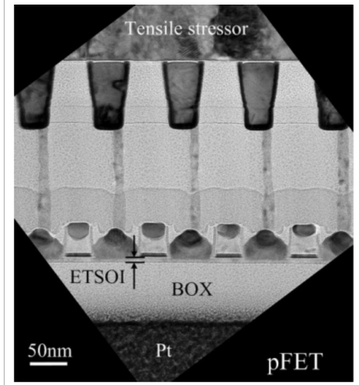

Figure a

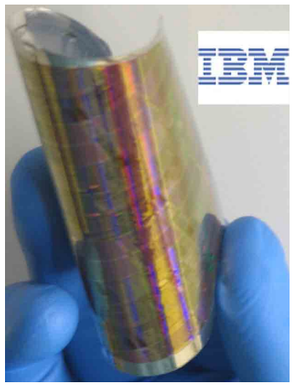

Figure b

Figure 3: A cross-sectional view taken by a TEM electron microscope after selective removal of the residual silicon, confirming the structural integrity of the device (a). The final 100-mm-diameter flexible ETSOI circuit on plastic (b).

micromechanical resonators.

A challenge with MEMS technology is integration of these MEMS devices with CMOS circuitry. An aspect of the problem is the difficulty of successfully releasing, or removing the MEMS devices from a substrate after fabrication, so that they can be integrated with CMOS.

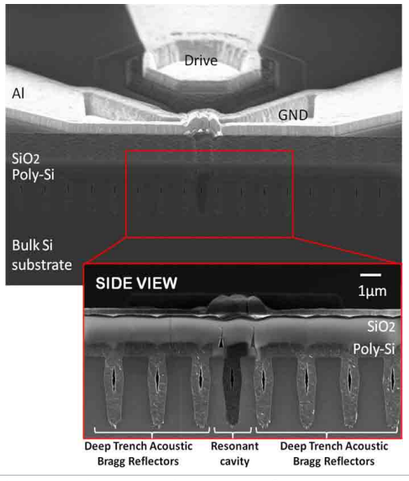

Figure 4: The SEM photos are of an unreleased deep trench resonator. Acoustic Bragg Reflectors formed from periodically spaced deep trenches define a high-Q resonant cavity in the center of the device. Deep-trench capacitors inside the resonant cavity form electrostatic drive and sense transducers.

Now an M.I.T team avoid that need altogether, by driving the MEMS resonator electrostatically, using deep trenches, that function as capacitors (**figure 4**). The resonator frequency could then be tailored easily, via the lithography used to build the trenches, and the trenches also serve as acoustic Bragg reflectors to confine and localize the resonance vibrations. Built using a 32-nm SOI process, the 3.3-GHz MEMS resonator had a Q of 2057—the highest reported to date for an unreleased MEMS resonator. The work paves the way for high-Q, multi-frequency sources to be built and intimately integrated in CMOS with no need for additional processing or packaging (**reference 5**).

www.powersystemsdesign.com

References:

1. Paper 4.6: Wang, H., et al, *Large-Scale 2D Electronics Based on Single-Layer MoS2 Grown by Chemical Vapor Deposition*, Massachusetts Institute of Technology
2. Paper 10.2: Park, S., et al, *RRAM-Based Synapse for Neuromorphic System with Pattern-Recognition Function*, Gwangju Institute of Science and Technology
3. Paper 9.1: Lue, H.T., et al, *Radically Extending The Cycling Endurance of Flash Memory (to >100M Cycles) by Using Built-In Thermal Annealing to Self-Heal the Stress-Induced Damage*, Macronix International
4. Paper 5.1: Shahrjerdi, D., et al, *Advanced Flexible CMOS Integrated Circuits on Plastic Enabled by Controlled Spalling Technology*, IBM
5. Paper 15.1: Wang, W., et al, *Deep Trench Capacitor Drive of a 3.3-GHz Unreleased Si MEMS Resonator*, Massachusetts Institute of Technology

Direct 48 V to Intel VR12.0 DC-DC conversion

Saving *Big Data* $500,000 per datacenter, per year

By: Stephen Oliver, Vice President, VI Chip Product Line, Vicor

The exponential rise in *Big Data* generation, processing, and storage from sources such as intensive industrial simulations, medical research, and social-media sites highlights the growing demand for datacenter and cloud computing power worldwide. The subsequent task is to maximize energy efficiency thus saving money and natural energy resources, minimizing pollution, and meeting the US Department of Energy's *Exascale* challenge: perform an ExaFLOP (10^{18} floating-point operations per second) of compute workload using only 20 MW of power.

For the power semiconductor industry, the challenge is to provide efficient and *high quality* power conversion from 480-V 3-phase AC power entering the datacenter, all the way to 1-V, > 100-A processors. Here, *high quality* means providing the optimum voltage and current as the processor dynamically demands to operate at peak performance.

High-volume data processing and data communication tasks require large-scale dedicated, optimized datacenters designed and operated by companies such as IBM, Amazon, Cisco, Hewlett-Packard, Google, Cray and others. Within the datacenter infrastructure, classical physics principles dictate that power should channel to data-processing locations at a high voltage to reduce current and so minimize distribution I^2R losses.

Compute density (related to the number of processors, the amount of memory, and input and output functionality) and subsequent power draw are also key issues. As power per rack increases above 10 kW, loss in traditional 12-V rack distribution becomes excessive, with additional financial- and size-related costs such as larger, more expensive copper bus bars and connectors. Applications approaching 20 kW and above per rack require efficient 48-V distribution. An example is the POWER7-based IBM Blue Gene/Q, which uses 48-V distribution, 80 kW per rack, and is amongst the world's highest performance and most efficient supercomputers, achieving 20,132 TFLOPs and 2,026 MFLOPs/W. In addition, high-reliability or *high-uptime* servers require an energy storage system (typically lead-acid batteries) to provide back-up power in the event of a main AC-feed interruption.

Historically, small 30- to 40-W ASICs performed processing tasks such as switching and routing in the telecommunications industry. With the rise of *triple-play* (voice, video, and internet) usage on mobile phones and tablets, voice-only *telecom* equipment has become *datacom*, with the use of standard computing processors. These systems typically use a 48-V rail but in *central office* equipment, the battery back-up voltage may be much wider, requiring wider input-range converters, further reducing the system efficiency. Energy storage is proportional to voltage squared, so 48 V is, again, the superior choice.

Voltage regulation

Compute workload varies with time. As demand increases, each processor requires more power to maintain performance (MFLOPs). As workload reduces, the processor may throttle back, moving to idle or sleep states to conserve power. In anticipation of a change in power requirement, the processor sends a serial *VID* (voltage identification) code to the power delivery system. During all steady-state and transient periods, the voltage delivered to the processor

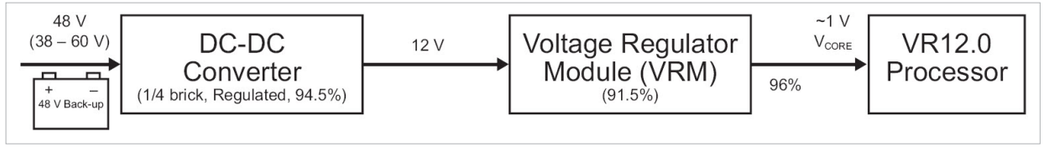

Figure 1: Traditional high-power server architecture showing the two-stage (48- to 12 to 1-V) approach, the wide 48-V input range (38 to 60 V), and back-up energy storage at 48 V.

must remain within tight, pre-defined limits to maximize performance and minimize the chance of a system crash.

VR (voltage regulator) specifications identify power requirements for Intel processors. The Sandy Bridge and Ivy Bridge processor chip-sets require compliance to the VR12.0 specification.

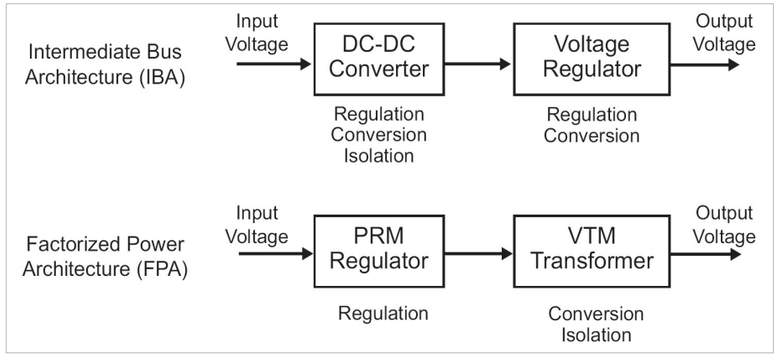

Figure 2: In traditional IBA, functions are duplicated thus increasing losses. FPA eliminates the duplication, increasing efficiency.

Traditional 12-V limitations

As power per rack increases, imposing a 48-V rack-distribution requirement, traditional 12-V to VR12.0 power converters require a separate 48- to 12-V conversion stage (**figure 1**).

It is critical to consider the power scheme in its entirety. *Headline* specification claims for DC-DC converters and VRMs ignore losses related to distribution loss on the motherboard and any connector losses. An accurate approach requires a measurement from the 48-V rail all the way to the processor socket, thus taking into account all potential loss elements. For example, the 96% in the figure accounts for the loss in the motherboard, due to the VRM being large and unable to be located adjacent to the processor socket.

An optimized 48-V alternative

Alternative architectures, such as the FPA (factorized-power architecture), employ different approaches to power conversion. The FPA architecture takes the regulation, isolation, and voltage transformation functions of a typical DC-DC converter and separates or *factorizes* them into individual elements implemented as VI Chips. A power-subsystem design then arranges these individual components (small, high-efficiency regulators and transformer / isolators) in the optimal power architecture (**figure 2**).

For example, VI Chip PRM regulators use a non-isolated buck-boost topology. The PRM accepts a varying DC input and creates a tightly regulated, adjustable DC output, V_F (the *factorized* bus voltage), which feeds into a downstream VTM transformer. The VTM is a fixed-ratio DC-DC transformer using an *SAC* (sine amplitude converter), which down-converts V_F directly to the processor's core voltage, V_{CORE}. ZVS (zero-voltage) and ZCS (zero-current) MHz-switching techniques achieve high efficiency and high power densities, with the PRM up to 97% peak and more than 1,000 W/in³, and the VTM more than 94% peak and over 100 A/in² .

The FPA power-system architecture maintains high-efficiency 48-V distribution along the entire path to the motherboard and uses the PRM and VTM VI Chips adjacent to the processor socket. The result is a highly efficient, small power system with a proven record in high-performance systems including the 48- to 1-V Blue Gene/Q system referenced earlier.

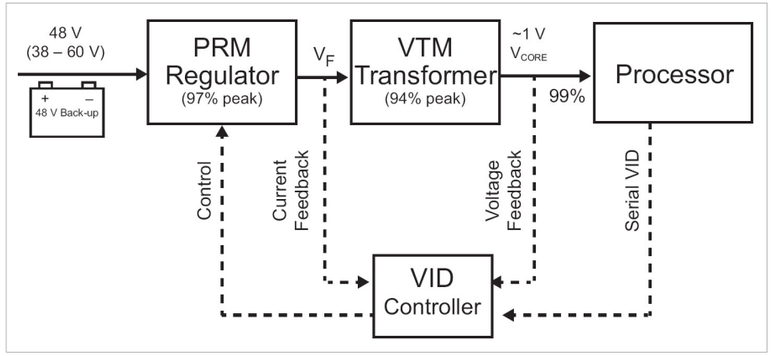

Figure 3: Information from the processor (VID instructions) and feedback on current and voltage input to the VID control IC. The IC sends a control signal to the PRM regulator, which controls the powertrain.

Direct 48 V to VR12.0 implementation

For Intel processor systems, the VI Chips form a *pure* powertrain accompanied by a separate VID controller (**figure 3**). VID controller acts as a translator between the processor's digital VID and the FPA powertrain, which, in turn, uses the optimal fast analog-control loop to provide an accurate processor core voltage, V_{CORE}.

To demonstrate VR12.0 compliance, Vicor engineers created a voltage-regulator-test board configured to support a 145-W, *socket-R* processor. A VTT (voltage-test tool), inserted into the processor socket, emulates processor behavior to characterize the performance of the powertrain. Monitoring instrumentation records the performance and compares it to the VR12.0 specifications using an automated spreadsheet.

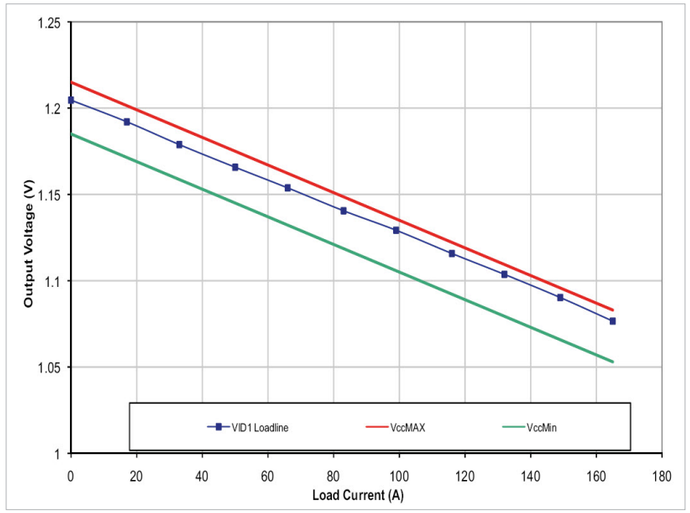

Figure 4: Test results show that the FPA maintains the processor core voltage, V_{CORE}, (blue) well within VR12.0 limits as the load increases from zero to 165 A.

As processor-current demand increases, the power subsystem cannot respond instantly, so the VR12.0 standard allows V_{CORE} to droop in relation to the load.

A plot of V_{CORE} versus processor current is known as a *load line*. The standard requires that the system design maintains the V_{CORE} load line within tight limits to ensure processor stability and performance (**figure 4**).

Under load transients (16 to 147 A) the FPA system has a clean, stable response within 5 µs with only SMT ceramic caps.

The FPA system does not use traditional large, unreliable electrolytic capacitors.

The core voltage response to a VID instruction—Dvid_pso_121A_up (1.05 V to 1.07 V) —shows that the FTP reacts with a stable output in only 2 µs. Additional, more stringent tests, using sweeps of

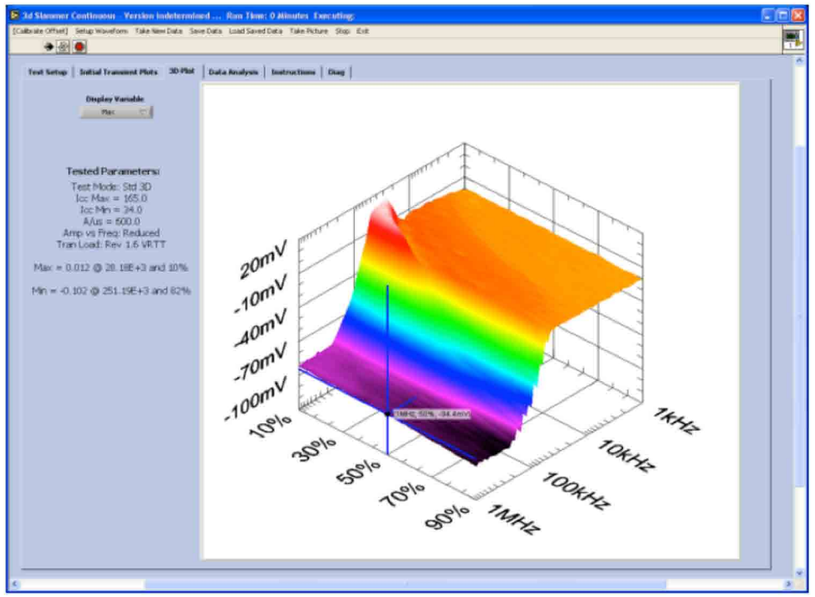

Figure 5: 3D-matrix result at Intel Dupont. X axis: load (processor) switching frequency; Y axis: load-switching duty cycle; Z axis: maximum excursion above set point

processor frequency, load power, and VID commands provide *3D-matrix* plots (**figure 5**).

Savings

A complete evaluation of the powertrains showed significant size and efficiency savings for FPA versus the traditional IBA system. Server motherboard real estate (PCB area) is expensive. A reduction in powertrain size means an increase in space available for more compute functionality (processors, memory, and input and output functions). The FPA architecture uses 50% less motherboard than IBA, while eliminating the off -board DC-DC converter—in total, a 2/3 reduction in size.

In power terms, the FPA architecture is > 5% more efficient from 60% to 100% processor load (**figure 6**). This is a significant improvement in VR12.0 systems, representing a 10 W or 30% reduction in power loss per processor.

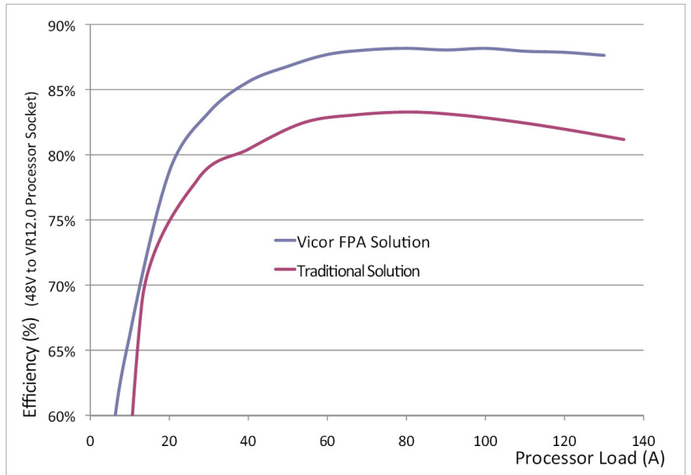

Figure 6: 48-V to VR12.0 socket efficiency (accounts for conversion and distribution losses).

Adjusting for 85% usage rate and air-conditioning costs (+70%), the final value per processor is 14.5 W saved. A new-build datacenter typically uses 30,000 processors. Using a $0.13 per kWh price for electricity, 14.5 W quickly becomes $500,000 reduction per year in datacenter operating costs (equivalent to 2,300 imported barrels of oil). The annual saving means that in less than three years, the VI Chip VR12.0 powertrain completely pays for itself.

www.vicorpower.com

Improve server efficiency with advanced control electronics

Power-management methods and three-phase BLDCs for cooling reduce data-center energy use

By: Mark Gaboriault, Strategic Marketing Director, Alexander Latham, Systems Engineer, and Thomas Rowan, Strategic Marketing Manager, Allegro MicroSystems

The worldwide growth in computer server farms and internet traffic has resulted in this infrastructure consuming global energy production at an accelerating rate. It is now estimated that the world's 500,000 data centers and 32 million individual servers consume 1.5 % of global electricity—about 300 TWhr of electricity per year (**reference 1**).

With significant efficiency improvements already attained, attention is now turning to power and heat management at the server-component level, specifically the on-board cooling fans themselves, which consume 10% to 15% of the total power used by the server (**reference 2**).

Recent advances in integrated control electronics provide local closed-loop control of both supply to the server and demand within components. These advances also make it practical to migrate from traditional single-phase BLDC (brushless DC) motors to highly efficient three-phase BLDC motors for the fans, typically realizing up to 25% improvement in efficiency.

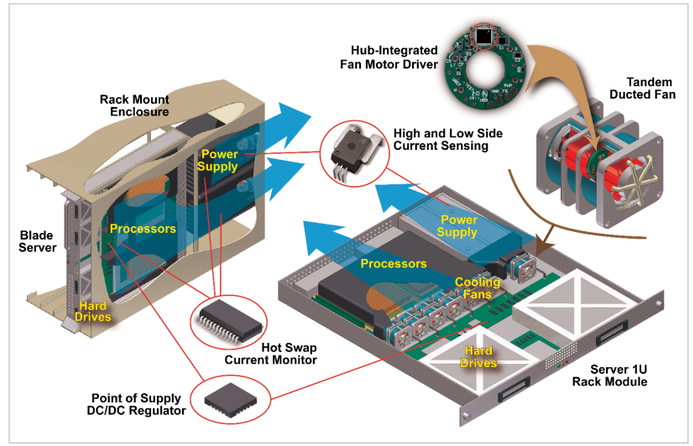

Figure 1: Fan-management, current-sensing, hot-swap-management, and PoS-regulation applications in standard rack-mount and blade servers.

The electronic devices allow inexpensive management of server components with minimal contributions to thermal signature, power draw, or physical size. Some, such as the Allegro MicroSystems A4942 three-phase sensorless BLDC-fan-motor driver chip, are small enough to fit onto the hub PCB of mini ducted fans. The hub PCB is a small ring-shaped board with as little as 5-mm effective width, to accommodate the rotor shaft (**figure 1**). Monitor ICs, such as the ACS761, provide current and power monitoring and control, enabling hot-swap management at the individual server blade level.

Energy reduction strategies

The latest generations of servers provide several new approaches to energy management, which allow rapid recovery of conversion costs—often within a year. For example,

microprocessors have been designed for higher throughput in smaller packages, requiring less power and generating less heat.

Studies of the individual thermal sources, principally power supplies and microprocessors and their enclosures, have led to optimized heatsink geometries and component layouts, with channeled shrouds to direct laminar airflow across these key areas. This complements the more recent high-efficiency ducted miniature (less than 40-mm) tandem fan-motor assemblies, arranged in series or parallel arrays within these flow paths.

To increase airflow efficiency and minimize footprint, the integrated fans assemble in tandem pairs that share the same ducting. The two fans are, however, completely independent in terms of mounting shaft and drive electronics. While this could gain an advantage from modular control, in fact it can introduce problems effecting reliable sensorless motor startup: Left to themselves, one of the motors will start first, causing airflow over the other fan and dragging the motor and interfering with the open-loop startup sequence.

A similar problem can occur when one fan has not yet stopped turning when the motors restart. In the past, this phenomenon made it necessary to allow both fans to come to a complete stop before restarting. The new motor-driver ICs contain an adaptive startup algorithm that can interpret when the motor is being driven by airflow over the fan blades from the tandem fan or when the motor and fan are already in motion from a previous power cycle. The advanced IC can modify the power-on sequence to adjust for this and allow both fans to operate synchronously at maximum efficiency throughout the power cycles.

Optimization of airflow is fixed, however, and improvements in PID control systems are required to optimize fan usage in terms of speed and idle time. Many servers are utilized only a small percentage of the time. Energy during the low demand periods can be saved by low-power or even power-down modes with automatic startup.

This can be accomplished by monitoring current consumption as the components operate, using current sensing ICs that can mount on the PCBs in the servers for lower-current onboard applications and on supply lines for high-side current sensing. These compact ICs measure current magnetically, using the Hall effect, eliminating the need for sense resistors, which dissipate heat. For example, an integrated conductor, such as in the Allegro ACS758, presents only 100 μΩ resistance, which is an order of magnitude lower than typical sense resistors, and results in significant power savings.

This technology also provides isolated current sensing in a compact package, providing a low voltage output signal for closed-loop feedback. Applied with advanced PWM motor drivers, these devices can control supply current surges and to ensure direct closed loop fan speed control to hold the airflow rate consistent and in proportion to the actual cooling requirement.

This also results in material savings because motors do not have to be overdesigned to compensate for large motor-to-motor torque and speed variations. Individual motors often have electrical characteristics that vary more than 10% between units. In addition, the local environments in which the motors mount vary substantially and inconsistently in terms of electrical supply and load, as well as thermal loading from coolant flow and adjacent heat sources.

Advanced PWM motor drivers and hot-swapping current-monitoring ICs can suppress current surges as the motors turn on. New device types apply soft-start PWM current-ramping techniques that allow the designer to optimize tradeoffs between surge current and power cycle times (**figure 2**).

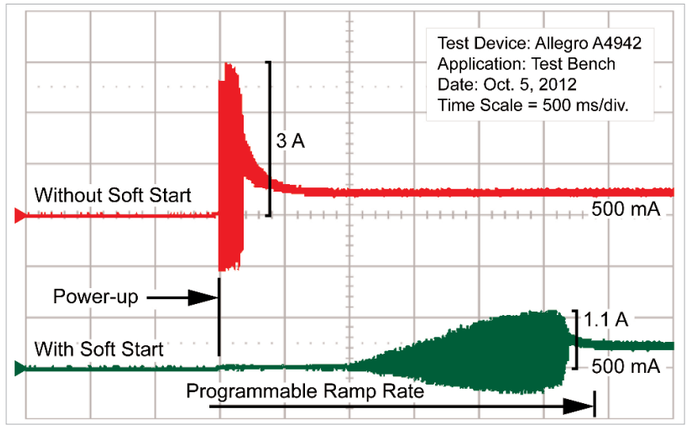

Figure 2: Effect of soft start in reducing surge current

Additional efficiency is gained by the test device—in this case, the A4942—which has advanced features that start to energize the motor phase windings in advance of the timing defined by the rotor position.

This phase advance technique ensures that the phase windings have reached the required current level at the point where the resulting forward torque on the rotor will be most effective, thereby improving motor efficiency. Note that the start and stop conditions are the same but with soft start, the maximum current is greatly reduced. The longer power-up may not be significant in start-stop fan application, and can be programmed to tradeoff with power surge.

Integrating Hot-Swap Management

Existing server-blade technology seeks to minimize these variances by the modular approach, placing power supplies and cooling

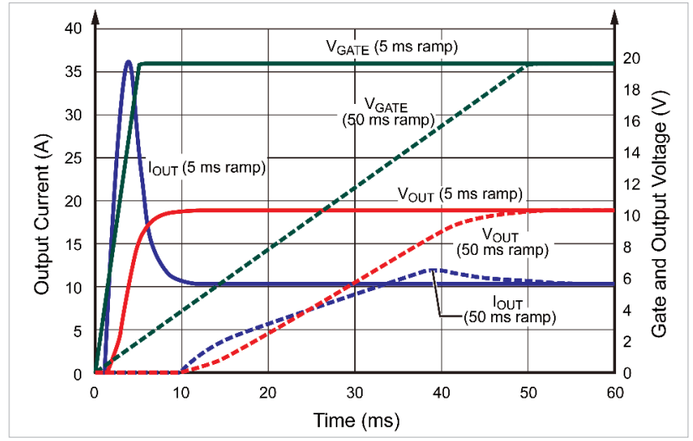

Figure 3: Hot-swapping current-surge suppression simulation

power is connected, the hot-swap current-sensor IC, in this example the ACS761, reduces the inrush current from 32 A to 12 A.

Hot-swap management affects the design of the other components in the server. This reduces the requirement for components to be rated for high inrush current levels. Additionally, by integrating current and power limiting, the hot-swap IC not only minimizes the board area that must be isolated from the operator for fans off-board from the memory storage and processor elements. However, this incurs significant risks in hot swapping. Current-sensor ICs with integrated hot-swapping controls manage power surges that occur due to the makes and breaks of electromechanical connections. The soft start of an external FET controls the hot-swap power surge and provides current limiting (**figure 3**). By controlling the FET's turn-on time when compliance with UL 60950, but also provides short-circuit protection.

Three-phase motor advantages

Although single-phase BLDC motors cost less than three-phase motors, increasing energy costs have made the higher efficiency of the 3-phase motor an economic offset. Typical efficiency improvements from single-phase BLDC motors to three-phase BLDCs are approximately 25%.

Designs achieve further cost reductions using techniques such as motor soft start to reduce the current surges from the power supply at startup. This reduction in surge current also allows smaller FETs and reduces costs for power supplies.

Along with optimized motor drivers, power-regulation techniques can optimize the operation of various components and systems within the server. QFN-sized DC-DC regulators provide point-of-supply management integrated with advanced features, such as synchronous rectification for high efficiency, short minimum-controllable on times, and optimized high- and low-side FET $R_{DS(on)}$ ratio for V_{IN}/V_{OUT} ratios commonly found in servers. These provide robust fault-tolerant power management to withstand variable operating conditions, and detect and report a wide variety of fault conditions.

3-phase BLDC motors used with advanced integrated circuit control and monitoring are providing significant efficiency gains now, and provide a path to future improvements. Because these technologies can apply at the subsystem level, they can scale to both DG (distributed generation) and CHP (combined heat and power) systems. With the improved electronic-evaluation techniques, these devices enhance server-system microgrid integration with Smart Grid systems.

www.allegromicro.com

References:

1. Biello, David, *Can Facebook Show How to Reduce the Growing Energy Use of the Internet?*, Scientific American, August 3, 2012.
2. Neudorfer, Julius, *How to Optimize the Energy Efficiency of Your Server*, eWeek, March 5, 2009.

SPECIAL REPORT : COMMUNICATION, DATA, AND COMPUTING

Voltage mode versus peak-current mode

Synchronous rectification provides rapid response to large load-current steps but demands care in control-loop operation

By: Milan Marjanovic, Power Design Engineer, Texas Instruments

High-speed bi-directional power supplies benefit from synchronous rectification, which plays a key role in backward power flow and greatly improved load-current transient response. This behavior, however, depends on never entering discontinuous conduction mode.

This means if huge load-current steps of 0% to 100% are expected, then the converter has to stay permanently in continuous-conduction mode. Huge load-current steps produce negative current flow for a much longer period of time than only one switching cycle. A careful selection of regulation techniques in isolated power supplies with bi-directional power flow is, therefore, necessary.

Voltage or peak-current mode?

Many modern PWM controllers, such as the UCC28950 from Texas Instruments, provide a burst mode implemented to keep the efficiency high at light loads. Activating burst mode leads to a very low primary current and narrow duty-cycle operation on the output at light or no load. This means that the signal on the current sense pin is also very small. Furthermore, this condition inflicts a hundred-fold switching noise on the current signal in the millivolt range. This can cause controller instability, because the controller tries to regulate a peak current signal that is not clean enough.

Because the peak-current mode needs voltage-slope compensation anyway, the only solution is to add even more slope compensation to keep the controller stable in this light-load condition. Then, however, it is moving away from a real peak-current-mode to voltage-mode regulation. Applying a large load-current step produces a large voltage drop at first on the output until the controller comes up from burst mode and the loop starts to react.

Disabling the burst mode is the right approach for this high-speed bi-directional application. In this case, the synchronous rectification works continuously

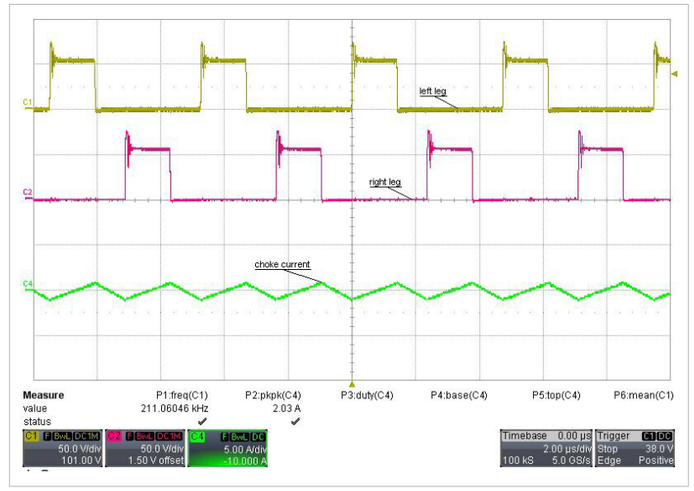

Figure 1: No load conditions: CH1-transformer secondary side, left leg; CH2-transformer secondary side, right leg; CH4 output choke current

and the duty cycle remains almost constant over the entire load range. Here again, however, at light- or no-load conditions the current-sense signal is still small and, additionally, has a negative component.

The explanation is simple: the current in the output inductor has two components: an average DC current and an alternating ripple current. At no load, the average current is zero but, due to the synchronous rectification, the current in the inductor remains continuous, therefore the alternating component is ever present (**figure 1**).

This behaviour forces the current sense signal on the primary side to zero for each period of cycle when the current is negative (flows back from the secondary side to the primary side). The higher the ripple current, the more negative this voltage. Because a diode blocks the negative voltage in the unipolar current sense, the voltage on the current sense pin of the controller is approaching zero (**Figure 2**). Now, if the dead time between each switching cycle is not large enough for demagnetization of the transformer, it will go into saturation.

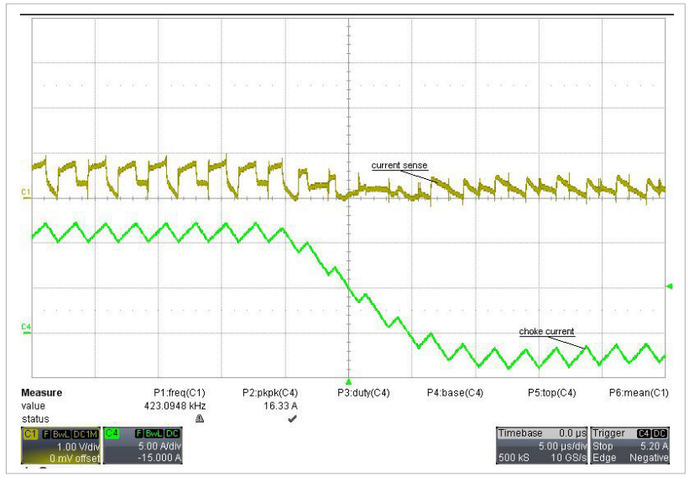

Figure 3: High load step condition—down slope: CH1-current-sense pin; CH4-output-choke current

Together this scenario is problematic and can confuse the controller, especially during large load-current steps. In this case, even high slope compensation doesn't help to keep the converter stable.

Figure 3 shows a high load-transient down slope of 1 A/µs. The average current crosses the zero level, becoming negative, and remains there until the overloaded energy returns to zero. In the meantime, the signal on the current sense pin has negative slope. A controller working in peak-current mode cannot regulate under these conditions.

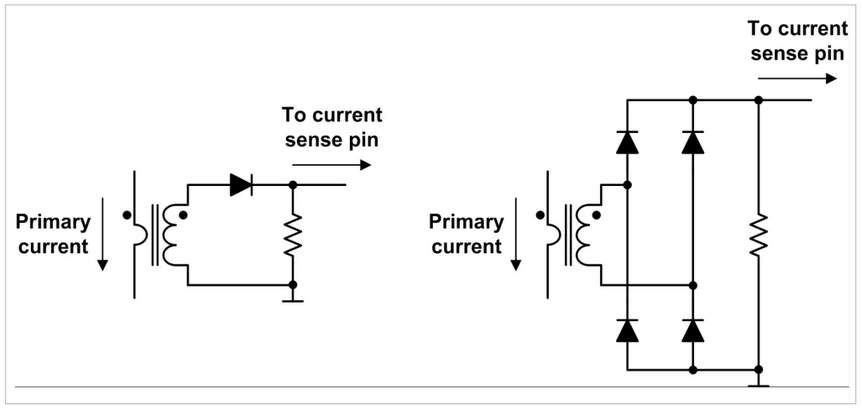

Figure 2: Unipolar (left) and bipolar (right) current sense

So, sensing unipolar current with a current-sense transformer is only suitable for converters where the current flows only in one direction, that is not in bi-directional mode. Unipolar current sensing requires a high-voltage diode with low parasitic capacitance and leakage current. Such sensing also requires a clamping circuit to protect against the energy overshoot coming from the leakage inductance caused by the current-sense transformer.

The experience gained during development showed that achieving a good transient response at load steps from 0% to 100% requires bi-directional working capabilities. Two requirements have to be fulfilled:

- The PWM controller has to work in voltage mode.
- Protection in the form of cycle-by-cycle over-current limitation is necessary.

Therefore, the current-sense transformer has to connect in series with the full-bridge transformer's primary winding. This ensures demagnetization, even when a negative current is measured (current flow from the secondary side to the primary side).

To protect the full-bridge transformer against saturation in voltage mode, the easiest way is to place a capacitor in series with the transformer. This blocks any DC component caused by any waveform asymmetry. The type and value of this capacitor is defined by the accepted ripple voltage and primary peak current. Usually in this case, 2% of the input voltage is allowed for the ripple voltage and a X7R capacitor is well suited for this task.

$$C = \frac{\Delta I_{op}\frac{1}{n}U_o}{2 f_{sws}\Delta U_{pp}U_i}$$

$$= \frac{(16.5\,A)\left(\frac{5}{4}\right)(32\,V)}{(2)(210\,kHz)\left(\frac{36\,V}{100}\right)(2)(36\,V)}$$

$$= 60\,\mu F$$

At the switching frequency of 210 kHz and minimum input voltage of 36 V, the capacitor needs to be 60 µF minimum. Considering the tolerances and allowable ripple current, we chose two 47-uF capacitors with the X7R dielectric. By combining the capacitors in parallel, PCB routing is easier and the parallel connection reduces the ciruict's overall RESR.

As shown, unidirectional current sensing works well only if the current remains in the same direction of flow. On the other hand, bi-directional current sensing works only with alternating current. As the current sense transformer connects in series with the primary winding of the full-bridge transformer and the system is working in voltage mode, this demand is fulfilled. Furthermore, the transformer is demagnetized inherently as only AC current flows due to the capacitor in series with the full-bridge transformer. The rectification of the current signal is not an issue any longer: the circuit can use low-voltage Schottky diodes. In voltage mode with bipolar current sensing, the current-sense signal is only used for protection. If the signal hits the threshold, the controller goes into cycle-by-cycle peak-current limit. For this reason, slope compensation is needed again.

Recommendations:

- Use a topology with synchronous rectification for a high-performance and high-speed power supply to achieve a good load-transient behavior.
- Disable any burst or power-safe mode and enable continuous conduction mode all the time. This reduces efficiency at low or no load conditions, but is the only way to react fast to large load-current changes.
- Avoid any components with low-pass behaviour like an optocoupler in the voltage feedback loop. The highest bandwidth is achieved, if the PWM controller is placed on the secondary side.
- Use voltage-mode control and bi-directional current sensing to avoid instabilities when current flows backwards (from the secondary side to the primary side).

www.ti.com

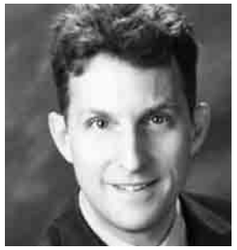

Control method maximizes efficiency of single-phase PFC stages

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

The drive to improve power supply efficiency for computing and other applications no longer focuses on simply on achieving maximum efficiency at the power supply's rated output, but rather on maximizing efficiency across the full load range. This concern is reflected by energy efficiency standards such as 80Plus, which aim to raise power supply efficiency under the conditions reflecting actual product usage, and therefore establish minimum requirements for either average efficiency or the efficiency at different load levels. Since PFC (power factor correction) stages, typically boost converters, have a noticeable impact on overall power supply efficiency, there are efforts underway to optimize the efficiency of these stages across the load range.

Among those working to improve PFC stage efficiency are the suppliers of PFC controller ICs. These companies are developing new control techniques to enable better PFC stage performance in line with the power-supply efficiency standards. An article by Joel Turchi in the October issue of How2Power Today explains the principles of operation for a technique called CCFF (current-controlled frequency foldback). Developed by ON Semi and implemented in some of their newer PFC controller ICs, CCFF enables improvement in efficiency at both light and medium loads in single-phase PFC stages where CrM (critical conduction mode) control is typically used. Like some other methods, CCFF works to reduce switching losses by reducing the switching frequency as the load level decreases.

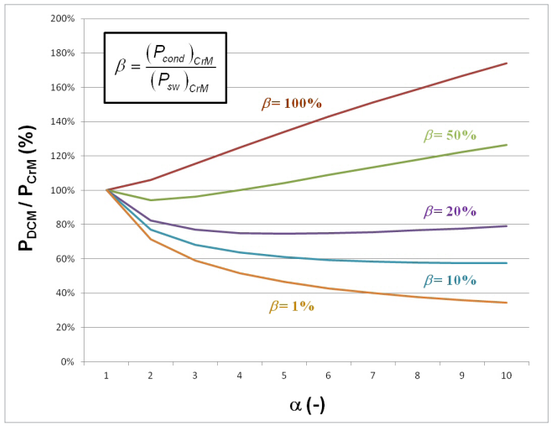

DCM losses as a percentage of the CrM losses with respect to the α = f_{CrM} : f_{DCM} ratio. (Courtesy of ON Semiconductor).

Reducing the switching frequency (frequency foldback) results in high efficiency at light loads. However, it's more difficult to apply this technique at intermediate load levels because frequency foldback also increases conduction losses. It is challenging to predict the optimum switching frequency that minimizes the PFC boost converter's total losses at each operating point (each set of line and load conditions). Without being able to predict the switching losses, it is hard to know under which operating conditions frequency reduction should apply.

But this article shows that using the line current to control frequency reduction is an efficient approach. It is demonstrated that, rather than trying to compute switching losses directly, an easier method is to predict the trends in losses under DCM (discontinuous conduction mode) operation versus losses under CrM operation as a function of frequency reduction. From this analysis, it is seen that efficiency is optimized when the frequency is reduced in accordance with the conduction losses relative weight on the total CrM losses—in other words, as a function of the line current. This conclusion is confirmed by experimental data.

As the figure below suggests, as the line current falls, so does the optimal switching frequency. The CCFF technique exploits this relationship to optimize and flatten the efficiency curve of the PFC stage over the load range.

For the targeted applications, this approach is shown to be more effective than implementing frequency foldback as a function of the output power level. For more on the analysis behind this control method, experimental results, and comparisons with other PFC control techniques, see *Frequency-Foldback Technique Optimizes PFC Efficiency Over The Full Load Range*, by Joel Turchi, ON Semiconductor, Toulouse, France, in the October issue of How2Power Today, which is available online at www.how2power.com/newsletters.

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 http://www.how2power.com/newsletters/.

www.how2power.com

Software security urged for medical devices and networks

By: Gail Purvis, Europe Editor, Power Systems Design

Edinburgh University scientists working on host-pathogen biology and advanced biochip technologies for infectious diseases have developed a new test using a strip with electrical sensors that show if wounds or lesions are infected with bacteria, including MRSA (methicillin-resistant staphylococcus aureus). The head of biochip research at the Division of Pathway Medicine, Dr Till Bachmann, notes that hand-held tests provide rapid results, allowing almost immediate detection of bacteria, so patients get more-effective drugs, quicker to speed up recovery. In UK hospitals, conventional laboratory tests to confirm MRSA presence in a wound can take a full day.

While infection is common, a less publicized aspect of malware and bugs is found in both the computing software and data embedded in hospital equipment and implanted medical devices, increasingly vulnerable to software infection. Europe has as its standard the IEC 62304 for developing safety-critical and high-reliability software for medical devices. But, while medical-device-software developers comply with standards, medical equipment becomes increasingly interconnected. With many systems run on Windows or variants, these unwittingly become a common target for hackers.

Devices are usually connected to an internal network, that is itself connected to the internet, and accordingly are vulnerable to infections from laptops or other device that are brought into hospitals. In the US, the issue seems exacerbated by the fact that manufacturers often will not allow their equipment to be modified, even to add security features, because such modifications could run afoul of the US Food and Drug Administration regulatory reviews. Instead, infected computers must be removed for cleaning. Devices can become sufficiently compromised that they can't record and track data. Vulnerable equipment ranges from intravenous-drug and nutrition compounders to picture-archiving systems associated with diagnostic equipment, and includes MRI devices.

The US Government Accountability Office recently issued a warning report, http://www.gao.gov/products/GAO-12-816, that implanted defibrillators and insulin pumps could be vulnerable to hacking, though no attacks on these devices have been reported. Compounding the issue is hospital devices are rarely reported to regulators, and such reporting is not required unless a patient is harmed, although researchers report that the FDA is now reviewing its regulatory stance on software.

As the internet of things gathers pace, news is Canadian LionsGate Technologies is to use the audio jack of the smart phone/tablet/PC to run low-power diagnostic equipment, taking advantage of the host computer's superior processing. It's a really neat move to power-saving, but still is left open to the vulnerability to the host device's software security.

www.powersystemsdesign.com

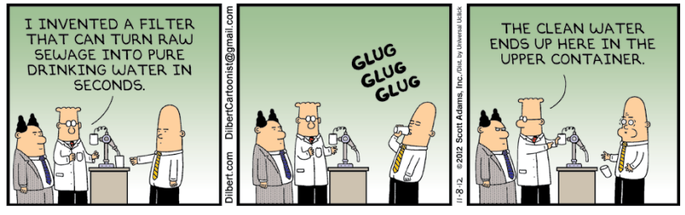

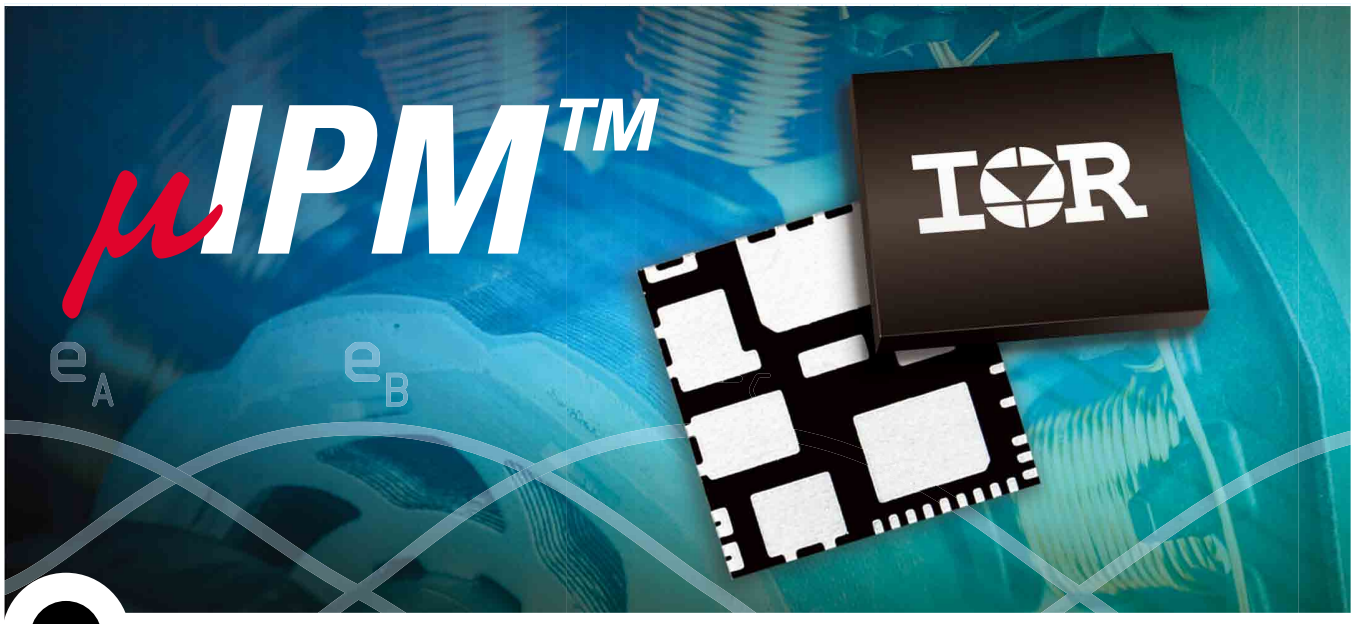

Innovative Power Module Reduces System Size

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

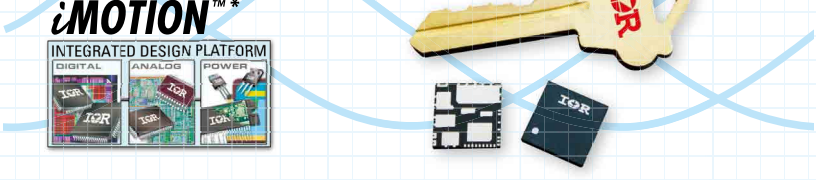

For more information call +49 (0) 6102 884 311
or visit us at www.irf.com

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