LINEAR TECHNOLOGY

Journal of Analog Innovation

April 2017 Volume 27 Number 1

IN THIS ISSUE

Robust Industrial Sensing with the Temperature-to-Bits Family

Logan Cummings

Precision temperature measurements provide important data points for industrial control systems to maximize efficiency, safety and reliability. Of course, industrial environments are challenging for performing precision measurements. High voltages, large sources of EMI and long cable runs between sensor and measurement device all conspire to deter accuracy. This article shows how to overcome these barriers, and improve the noise immunity and overvoltage tolerance of the LTC®2983, LTC2984 and LTC2986 temp-to-bits devices, without degrading measurement accuracy.

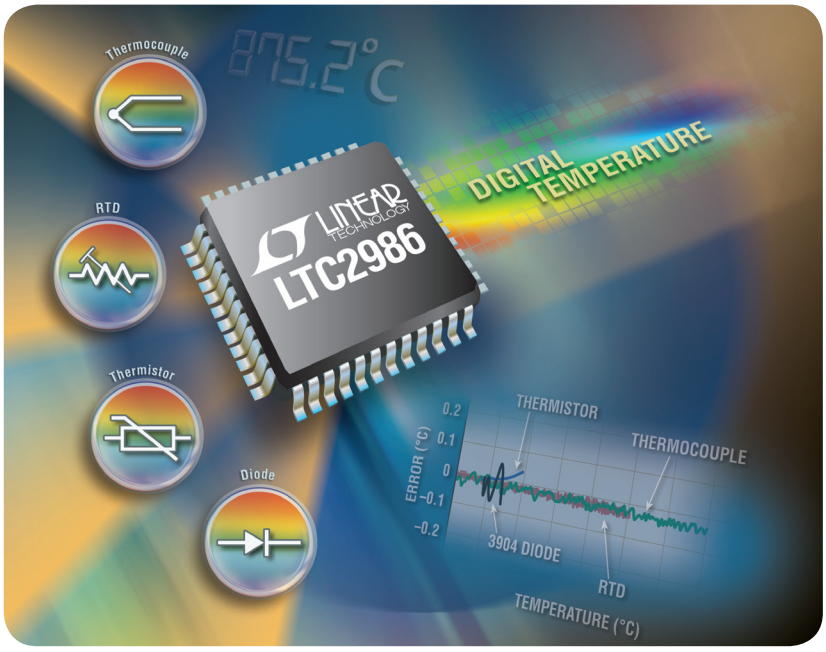

The temp-to-bits family simplifies industrial temperature measurement.

With up to 20 input channels, the temp-to-bits family is well suited for industrial temperature monitoring. Multiple sensors can be used to evaluate temperature gradients and identify anomalous behavior in time to prevent equipment failures.

Chassis wiring often exposes high voltages to sensors—the sensor wires provide a low impedance path back to the measurement circuit, which makes possible a short circuit; some cause for concern. This is especially true with bare-tip thermocouples, essentially a wire probe.

EMI in the industrial environment includes 50Hz or 60Hz line noise radiated from synchronous motors and other high current devices, as well as higher frequency noise from radio links, unshielded communications cables, and broadband noise from arcing.

(continued on page 4)

In this issue...

Linear in the News

LINEAR TECHNOLOGY COMBINES WITH ANALOG DEVICES

As of March 10, Linear Technology is now part of Analog Devices. The combined company offers customers a broad array of capabilities to address their design challenges. The companies' expanded portfolio now includes among the industry's most complete analog, mixed signal and power solutions. With these vast tools, we look forward to continue partnering with customers across markets spanning automotive, communications, consumer, healthcare, industrial automation and process control, and aerospace and defense. We plan to keep you apprised of progress and developments as two innovative analog companies combine to provide designers with a broad range of solutions.

8-SWITCH MATRIX LED DIMMER WINS AWARD

Linear Technology's LT®3965 8-switch matrix LED dimmer was named as a winner of *LED Magazine's* Sapphire Awards in the ICs and Electronic Components for Solid State Lighting category. The award was announced at the publication's award event in early March. The LT3965 is a LED bypass switching device that enables independent dimming and diagnostics of eight individual LEDs or LED segments. The LT3965 works in conjunction with an LED driver circuit configured as a current source. It features eight individually controlled floating 17V/330mΩ NMOS switches. Using these internal power switches and an integrated I²C serial interface, individual LEDs within the string can be turned on or off, or they can be PWM dimmed to offer unique patterns within the LED matrix.

Typical applications include automotive matrix LED headlights, industrial lighting and large LED display lighting. The I²C serial interface enables digital programming with 256:1 dimming ratios with or without the 11-bit resolution fade transition between the dimming states. Each switch can control and monitor a single LED or a segment of up to 16V of series-connected LEDs. The LT3965's 8V to 60V input voltage range can accommodate a wide range of LED drivers commonly used in automotive and industrial applications. For larger LED matrix applications, multiple LT3965s (up to 16) can use a common I²C bus.

ANNUAL PRODUCTS AWARD PRESENTED TO LINEAR FOR LTM4631

Linear Technology was presented with a 2016 Annual Products Award by *Electronic Products China* magazine for the LTM®4631. The LTM4631 dual 10A or single 20A µModule® step-down regulator is in an ultrathin 1.91mm LGA package with a 16mm × 16mm footprint. It can be placed on a board close to an FPGA, with a shared heat sink. It can be mounted on the backside of circuit boards, freeing space on the topside for other components. Applications include data storage systems, embedded systems, gateway controllers and 40Gbps to 100Gbps networks.

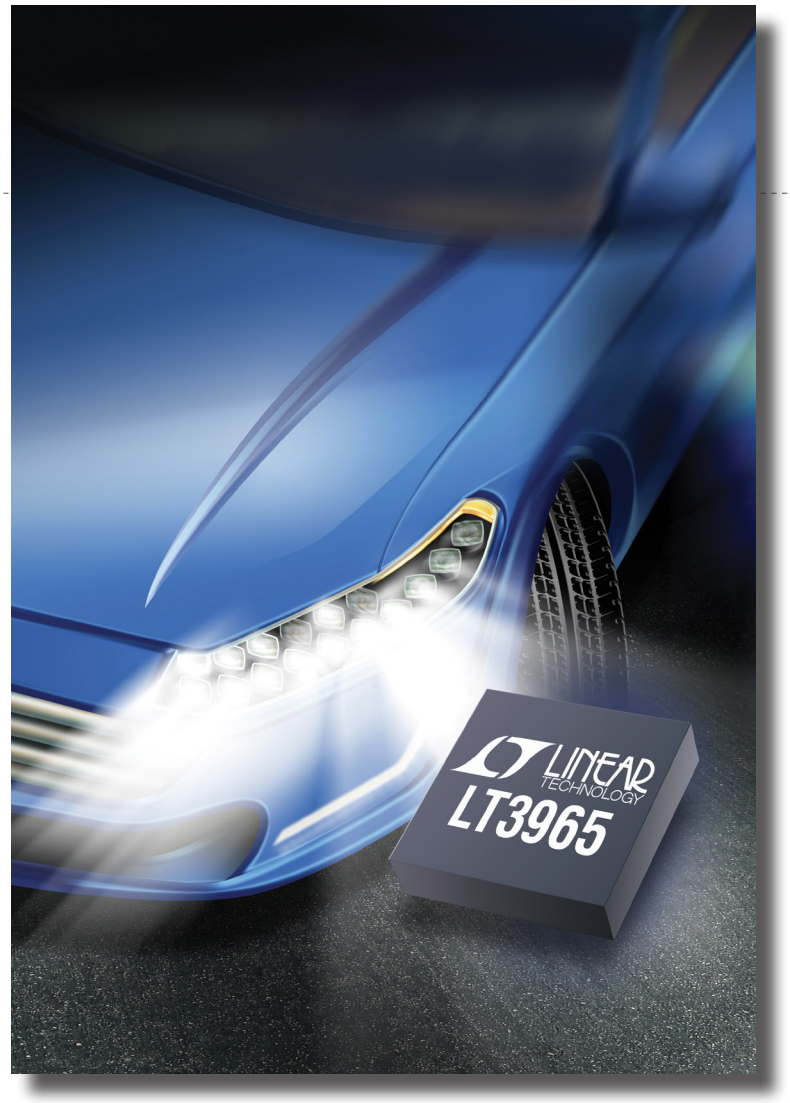

Linear Technology's LT3965 8-switch matrix LED dimmer was named as a winner of *LED Magazine's* Sapphire Awards in the ICs and Electronic Components for Solid State Lighting category.

CONFERENCES & EVENTS

Tech Taipei – EV & Power Management Seminar, Taipei World Trade Center, April 20, 10:10 am, Exhibition Hall 1, F2—"Advanced High Voltage Battery Stack Monitor and Active Balancer" by Galant Chen

GPU Technology Conference 2017, San Jose Convention Center, San Jose, CA, May 8–11, Booth 504—Featuring power products for graphics engines, including Linear's µModule regulators and Silent Switcher® switching regulators. Dave Dwelley is presenting "Powering the Brain & Sensors in Autonomous Vehicles," see conference website for time/location: www.gputechconf.com

One Mega Event ~ 2nd Transport India 2017 Expo, Pragati Maidan, New Delhi, May 10–12, Hall 12, Booth D12—Featuring Linear's automotive solutions, including battery management systems for hybrid/electric vehicles and a broad array of power solutions. http://10times.com/smart-transport

Internet of Things World, Santa Clara Convention Center, Santa Clara, CA, May 16–18, Booth 520—Showcasing Linear Technology/Analog Devices solutions for applications ranging from Smart City to wearables, industrial and medical. https://tmt.knect365.com/iot-world/

CommunicAsia 2017, Marina Bay Sands, Singapore, May 23-25, Hall L5, Booth 5G4-07—Showcasing the broad range of Linear Technology/Analog Devices solutions for wireless and wired communication systems. www.communicasia.com

CTI ISO 26262 Conference, Royal Park Hotel Rochester, Detroit, MI, June 26–28—Linear Technology will showcase its breakthrough wireless BMS concept car. Kyle Fabris is presenting "SEooC and wireless systems" at 2:40 pm, June 27. Info at www.euroforum.de/iso-conference/usa/

The LTC2983-6 temp-to-bits family provides numerous features that enable robust temperature measurement in industrial environments. Low input leakage permits large series resistance for overvoltage protection and RC filtering.

(LTC298x, continued from page 1)

PROTECTION

As discussed above, the industrial environment is full of hazards that must be mitigated for robust temperature measurements. Fortunately, the simple resistor can aid against the two biggest enemies: overvoltage and EMI.

The internal ESD structures in the LTC2983/4/6 family can sustain 15mA maximum, so we must ensure that an overvoltage condition from a thermocouple tip contacting the rail will not cause more than 15mA to flow into or out of the IC. A protection resistor (Figure 2) can be used to limit the current seen at the inputs in case of contact with the supply. For a 48V supply:

$$R_P = \frac{48V}{15mA}$$

$$R_P = 3.2k$$

Figure 3 shows the minimum protection resistance as a function of max fault voltage.

Figure 1. The LTC2986's Kelvin excitation modes support universal protected inputs for remote 2- and 3-wire sensors

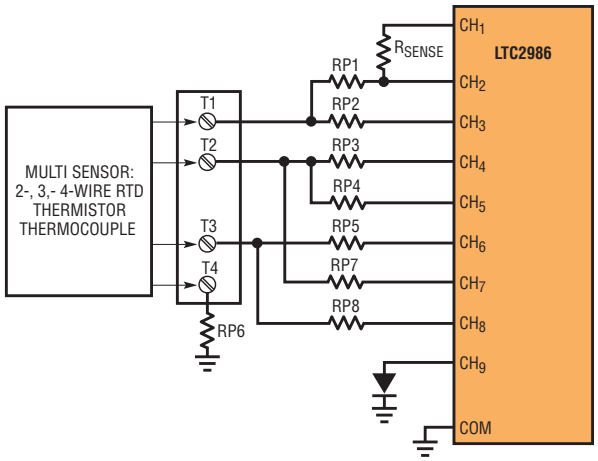

The protection resistor must also be able to handle its power dissipation:

$$P_{R(P)} = 3.2k \bullet (15mA)^2$$

$$P_{R(P)} = 720mW$$

Figure 4 shows the minimum protection resistor power rating as a function of max fault voltage.

This power handling requirement can be reduced if we use a larger resistance at the cost of increased offset due to leakage currents. The LTC2983-6 inputs have less than 1nA leakage current. The minimum 3.2k protection resistors add a worst-case temperature offset of 0.16°C; 10k protection resistors could contribute up to 0.5°C error but would only need to dissipate 230mW, affording the designer a smaller protection circuit.

Figure 2. Thermocouple with 48V protection resistors

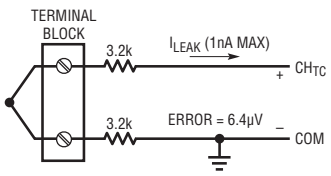

Figure 3. Maximum fault voltage vs minimum protection resistance

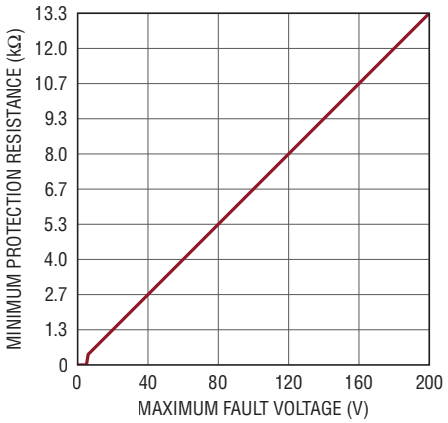

Figure 4. Maximum fault voltage vs minimum protection resistor power rating

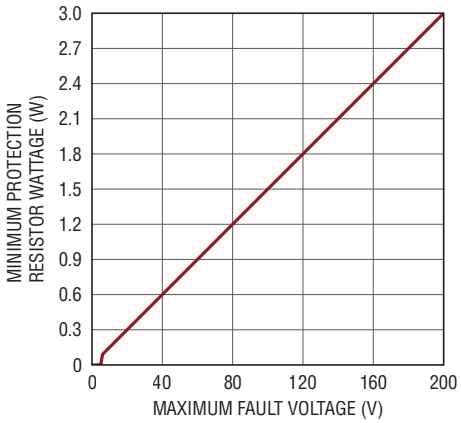

A major issue with remote sensors is knowing when they have failed. In many installations, the sensor has been placed remotely due to hazards encountered at the sensor location. In the case of remote thermocouples this is often a very high temperature and/or a corrosive atmosphere that can cause the thermocouple wiring to degrade over time. One solution to this problem is to apply a current through the thermocouple…

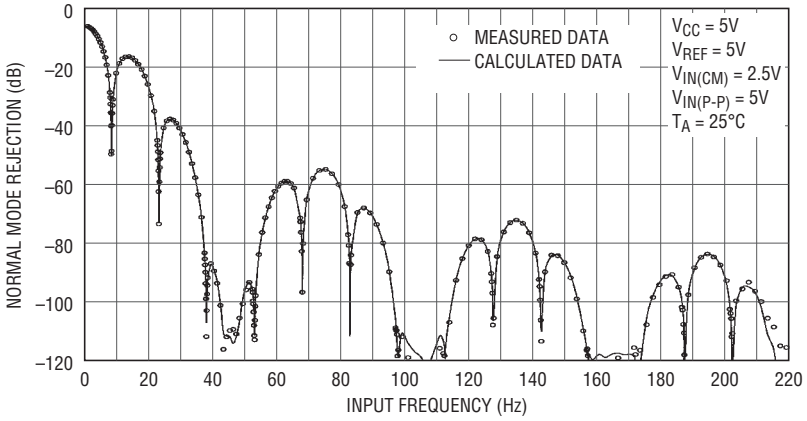

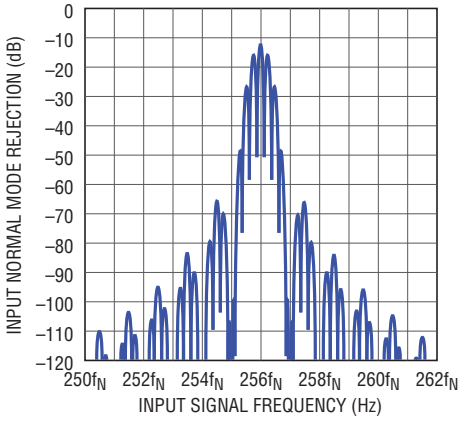

Figure 5. Input normal mode rejection at f_S = 15.36kHz

FILTERING

The LTC2983 family incorporates a sinc^4 filter that provides excellent simultaneous 50Hz/60Hz rejection and good wideband noise rejection at higher frequencies.

However, the ADC is a sampled system and there is a narrow passband at 256 times the notch frequency, as shown in Figure 5. This is where a simple RC lowpass or anti-aliasing filter on the inputs can improve the quality of the measurement in conjunction with the digital filtering.

Adding a capacitor to ground on the input side of the protection resistor forms a single pole RC filter. Since the frequencies present in the end user's environment are frequently unknown at design time, an input filter should strive to give as much attenuation at 15.36kHz as possible without causing an offset due to excessive settling time. Too large of an RC time constant will cause errors if the voltage being measured has not settled when the ADC begins sampling.

REMOTE THERMOCOUPLE BURN-OUT DETECTION

A major issue with remote sensors is knowing when they have failed. In many installations, the sensor has been placed remotely due to hazards encountered at the sensor location. Remote thermocouples are often placed in very high temperature or corrosive atmospheres that can cause the thermocouple wiring to degrade over time. As the voltages produced by the thermocouple are very small, a completely broken thermocouple could go unnoticed since the temperature reported may still be within expected limits.

One solution is to apply a current through the thermocouple, which will generate a large voltage if the thermocouple has become an open circuit as shown in Figure 6.

There is an offset voltage produced by any parasitic resistance in the leads. This offset voltage changes if either the current or the resistance changes, making it difficult to remove through calibration. Additionally, simple current sources and pull-up resistors transmit power supply noise to the thermocouple inputs, forcing stronger filtering in the measurement circuit.

Even a tiny constant current source produces a comparatively large voltage across a protection resistor. Here, two of them are in series producing an upper voltage of 64mV with only a 10µA burn-out detection current, equivalent

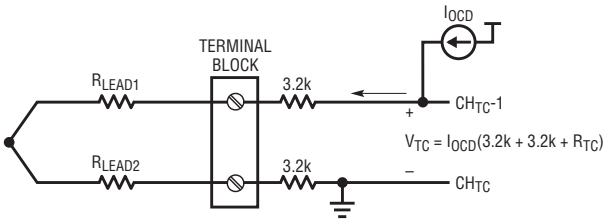

Figure 6. Pulsed OCD current produces detectable voltage which decays in normal use before measurement is made, eliminating offset error

A preferred approach to thermocouple burn-out detection is to supply a pulse of current that is of sufficient amplitude and duration to create a large detectable voltage if an open circuit is present, but short enough that this voltage decays below the noise floor within a determined time if the thermocouple is healthy. This ensures no measurable offset in the reported temperature.

to an offset of approximately 1600°C for an example type K thermocouple.

A preferred approach is to supply a pulse of current of sufficient amplitude and duration to create a large detectable voltage if an open circuit is present, but short enough that this voltage decays below the noise floor within a determined time if the thermocouple is healthy. This ensures no measurable offset in the reported temperature.

In the LTC2983/4/6 family, the user can select an internally sourced excitation current from 10µA to 1mA. When open-circuit detection is enabled, this current source is applied for 8ms at the beginning of one 81ms conversion cycle, allowed to decay, and then the standard 2-conversion thermocouple measurement is made.

With an overvoltage protection resistance of 3.2k, the effective series resistance seen by the current source is over 6.4k including the resistance of the thermocouple and cabling. Therefore, our maximum selectable excitation current for a 5V-powered device would be 500µA since a larger current would exceed V_{DD}.

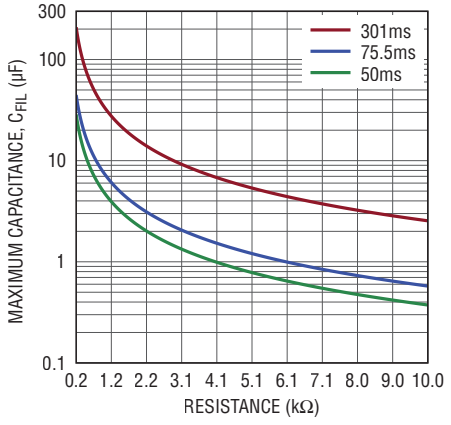

Figure 7. Thermocouple open-circuit detect settling time maximums (100µA)

FILTERING A REMOTE THERMOCOUPLE

A large RC filter reduces noise at the inputs at the cost of increased settling time from the open-circuit detection current pulse. The chart in Figure 7 shows the maximum RC combinations that settle to 1µV with various settling times.

The LTC298x provides an adjustable MUX delay of up to 25.5ms that can be used to extend the time available for settling—this is shown as the 75.5ms series in Figure 7. Very large RC filters can be supported by performing a dummy conversion after an open-circuit detection enabled conversion, extending the settling time by 160ms. Open-circuit detect disabled conversions can be used for the majority of data points, enabling open-circuit detect as needed or with a low duty cycle.

The traces in Figure 7 are maximum RC combinations for the different settling times. The graphs are based on an equation derived from the pulse response:

$$t_{SETTLE} = -\ln\left\{\frac{1\mu V}{R_{EFF} \bullet I_{OCD}\left[1-e^{\left(-8ms/RC\right)}\right]}\right\} \bullet RC$$

$$\text{where } RC = R_{EFF} \bullet C_{FIL1}$$

where R_{EFF} is the sum of the series resistance through the protection resistors and any thermocouple lead and tip resistance from CH_{TC-1} to CH_{TC} or to COM.

Using the graphs and following the line for the 50ms delay and 100µA excitation we find that for a 6.4kΩ total series resistance we can put approximately 560nF of capacitance on each input of a differential connection.

3.2k and 560nF per leg gives the RC filter an f_{3dB} of just under 90Hz. This yields over 40dB attenuation in the narrow band around our worst-case 256 • f_N frequency. Increasing the resistance can increase attenuation, but at the expense of increased thermal noise and offset voltage.

COLD JUNCTION MEASUREMENTS FOR REMOTE THERMOCOUPLES

Thermocouples are formed when two wires made of materials with different Seebeck coefficients meet. A temperature gradient along each wire generates a potential, where the potential difference between the wires is related to the temperature difference from the tip (mated end) to the (unmated) ends.

In order to calculate the absolute temperature at the thermocouple tip, the temperature at the ends must be known. The temperature at this junction of thermocouple wire and connector or circuit board is known as the cold junction. Since the absolute temperature at the tip is a function of this cold junction temperature, accurate measurement is a must.

A commonly encountered source of error in thermocouple measurement systems arises from the assumption that the temperature at the thermocouple connector is the same as that of the rest

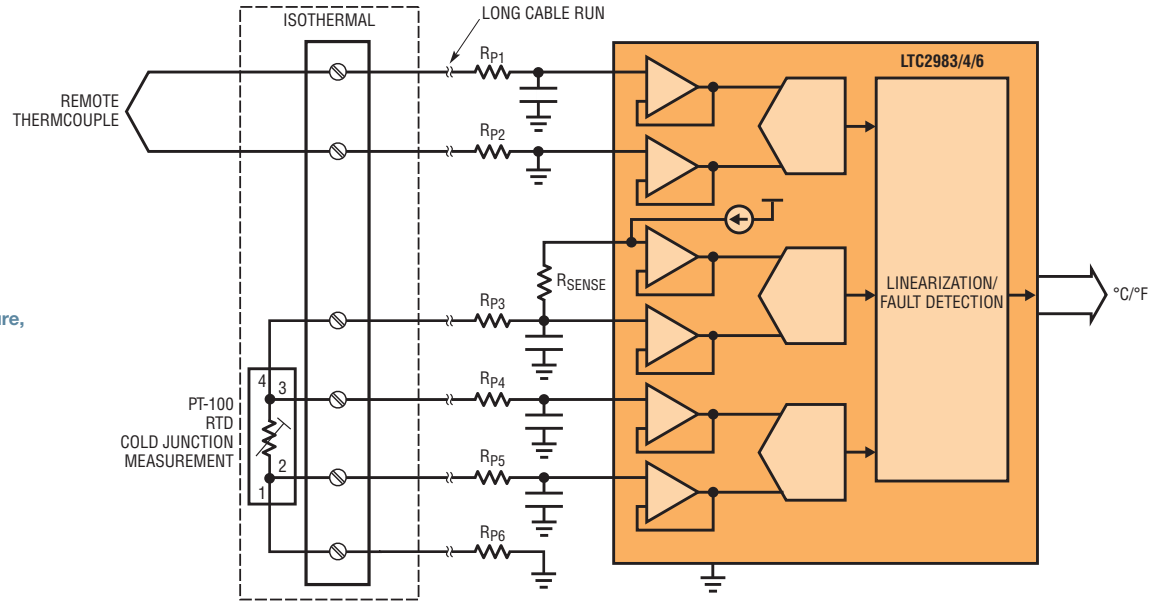

Figure 8. Remote RTD measures thermocouple cold junction temperature, improving accuracy and eliminating requirement for costly thermocouple extension wire.

of the circuit board—many systems simply take the temperature of the measurement device as the cold junction compensation value. Errors of several degrees arise from thermal gradients across the circuit board created by power supplies or other warm objects nearby.

Filtering Resistive Sensors

Resistive sensors provide a well defined method of measuring absolute temperature and are widely used for cold junction compensation. The temp-to-bits devices make it easy to add a cold junction measurement device, which is automatically used in the remote thermocouple temp calculation (Figure 8).

Like thermocouples, it is often necessary to protect and filter inputs measuring resistive devices like RTDs and thermistors. Since they are resistors themselves, it can be difficult to separate the temperature dependent resistance from leads and other series resistance in the measurement path.

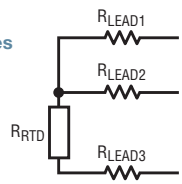

Figure 9. 3-wire RTD provides extra lead for correction

Three-wire RTDs (Figure 9) are commonly used to solve this problem. When all three leads are exactly the same length and resistance, the resistance between leads 1 and 2 can be measured to determine the lead resistance. In practice, it is difficult to ensure identical leads and the problem becomes much more difficult when additional series resistance is added to provide input protection or filtering. Any mismatch in the effective lead resistance results in a 1:1 measurement error due to mismatch.

If a fourth lead is added to the RTD, a 4-wire Kelvin measurement can be performed (Figure 10). The excitation current no longer flows through the sensed leads connected to the measurement inputs, eliminating any resistance and associated voltage drop in the leads. Additionally, the Kelvin connection allows us to rotate the direction of excitation current and remove any offsets created by parasitic thermocouples.

Figure 10. A 4-wire RTD is an ideal choice for a remote cold junction sensor as the 4-wire connection eliminates the effects of parasitic resistance from the long cable.

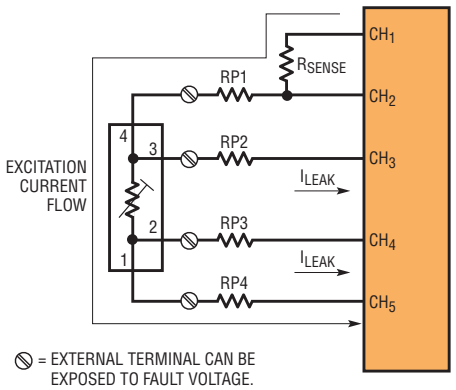

The LTC2986 and LTC2986-1 provide special modes for achieving low error measurements with 2- and 3-wire RTDs and thermistors. The LTC2986 achieves this accuracy by using additional connections to Kelvin sense the sensor element and provide excitation current from separate channels. In this way, the filtering and protection resistances are removed from the current path and their mismatch removed from the measured resistance.

Filtering the Cold Junction Sensor

Unlike a thermocouple, which may have an exposed tip, a cold junction sensor can be physically protected within a housing and does not require overvoltage protection.

Settling time for a resistive measurement is an important consideration, as too much RC delay leads to an effectively reduced resistance value. Unlike the

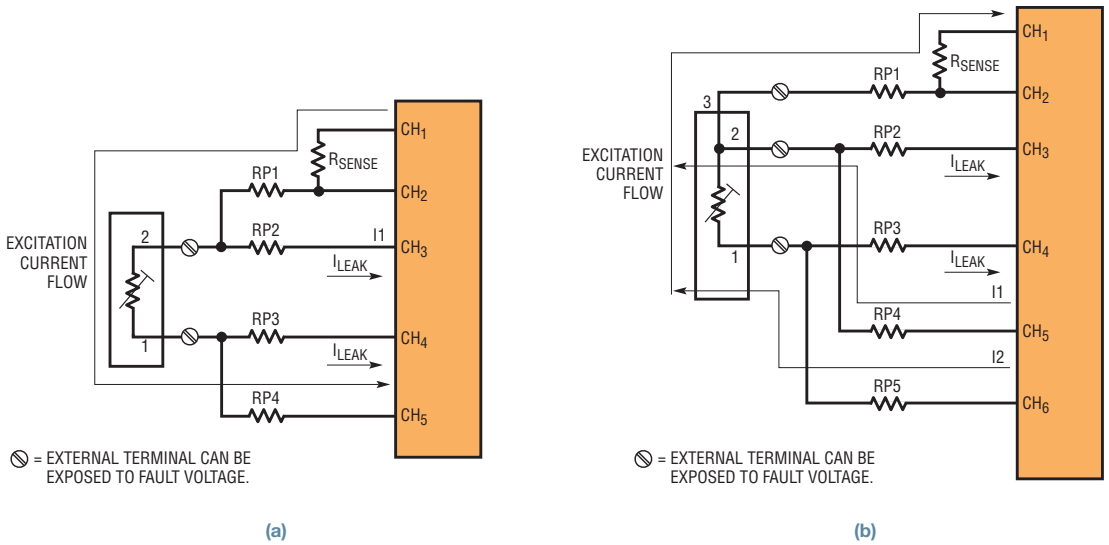

Figure 11. The LTC2986 enables the use of remote sensing with 2- or 3-wire RTDs. Shown are Kelvin current sourcing modes for (a) 2-wire and (b) 3-wire RTDs.

open-circuit detect pulse for a thermocouple, the excitation current for RTDs and thermistors is applied just prior to the start of a conversion cycle. The RC network should settle to required accuracy within a few milliseconds.

A 4-wire RTD is an ideal choice for a remote cold junction sensor as the 4-wire connection eliminates the effects of parasitic resistance from the long cable. With the LTC2986 we also have the option of using a remote thermistor or 2- or 3-wire RTDs, as the Kelvin current sourcing modes allow us to eliminate the cable resistance from the measurement in the same way (Figure 11).

A PT-100 RTD excited by a 250µA current has a temperature coefficient of about 96µV/°C. To keep offset error to less than 0.1°C the RC circuit should settle to within 9.6µV in approximately 3ms.

To ensure we do not exceed our current source compliance limits, calculate the steady state voltage at the sourcing channel with a 250µA excitation:

$$V_{CS} = I_{EXC}(R_{SENSE} + R_{FIL1} + R_{PT\text{-}100} + R_{FIL4})$$

Since we cannot exceed 4V, R_{FIL1} and R_{FIL4} must be less than 1.5k each. To maximize noise performance with the 250µA excitation source, use a smaller series resistor here, ideally less than 500Ω. Note that the resistances on the sensed leads (R_{FIL1} and R_{FIL3}) are not in series with the excitation current and do not factor in the steady state calculation. This is important because they do not contribute an offset in the 4-wire case. As a result, precision matching of these resistors is not necessary, as in the 3-wire RTD case.

A combination of R_{FIL} = 330Ω and C_{FIL} = 68nF (Figure 12) was found to settle to within 10µV in 3ms when all four leads from the RTD are filtered. The first and fourth connections do not factor in the measurement and are directly connected to either GND or the current source during measurement, giving us the possibility of reducing or eliminating the filter on these leads. It is good practice to provide some filtering, as noise on these leads could couple into the signal inputs.

Figure 12. RC filters for RTD should settle to within 9.6µs in approximately 3ms

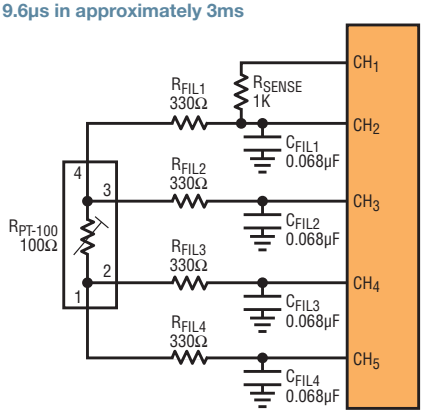

INCREASING THE CHANNEL COUNT WITH KELVIN CONNECTIONS

One limitation of the Kelvin sensing scheme is the number of channels in use by each sensor. While the LTC2986 trades channels for current sourcing flexibility, an external multiplexer can be used to increase the number of Kelvin-connected sensors, as shown in Figure 13.

Here we show how to Kelvin-connect nine thermistors on an LTC2983 or LTC2984 20-channel part. Each thermistor requires only two channels and the input protection resistors are removed from the current path.

The system controller connects each sensor's ground through the multiplexer just prior to initiating a conversion. In this way, the excitation current flows from the sense resistor through R_{P1}, across the thermistor and to ground through the analog MUX. Assuming the MUX has low channel-to-channel leakage, this technique should isolate the protection/filtering resistance from the measurement result.

Kelvin modes of the LTC2986 remove offsets due to extra series resistance when protecting and filtering 2- and 3-wire RTDs and thermistors. The Kelvin sensing technique implemented internally in the LTC2986 can be extended to larger numbers of sensors by using an external switch.

It is important to note that settling time will be a factor here, as the chain of potentially large valued thermistors with parallel protection resistors and filter caps creates a structure not unlike a lumped element transmission line.

CONCLUSION

The LTC2983-6 temp-to-bits family provides numerous features that enable robust temperature measurement in industrial environments. Low input leakage permits large series resistance for overvoltage protection and RC filtering. Kelvin modes of the LTC2986 remove offsets due to extra series resistance when protecting and filtering 2- and 3-wire RTDs and thermistors. The Kelvin sensing technique implemented internally in the LTC2986 can be extended to larger numbers of sensors by using an external switch. ■

Figure 13. Using an analog MUX to Kelvin-connect nine thermistors to an LTC2983 or LTC2984 20-channel part.

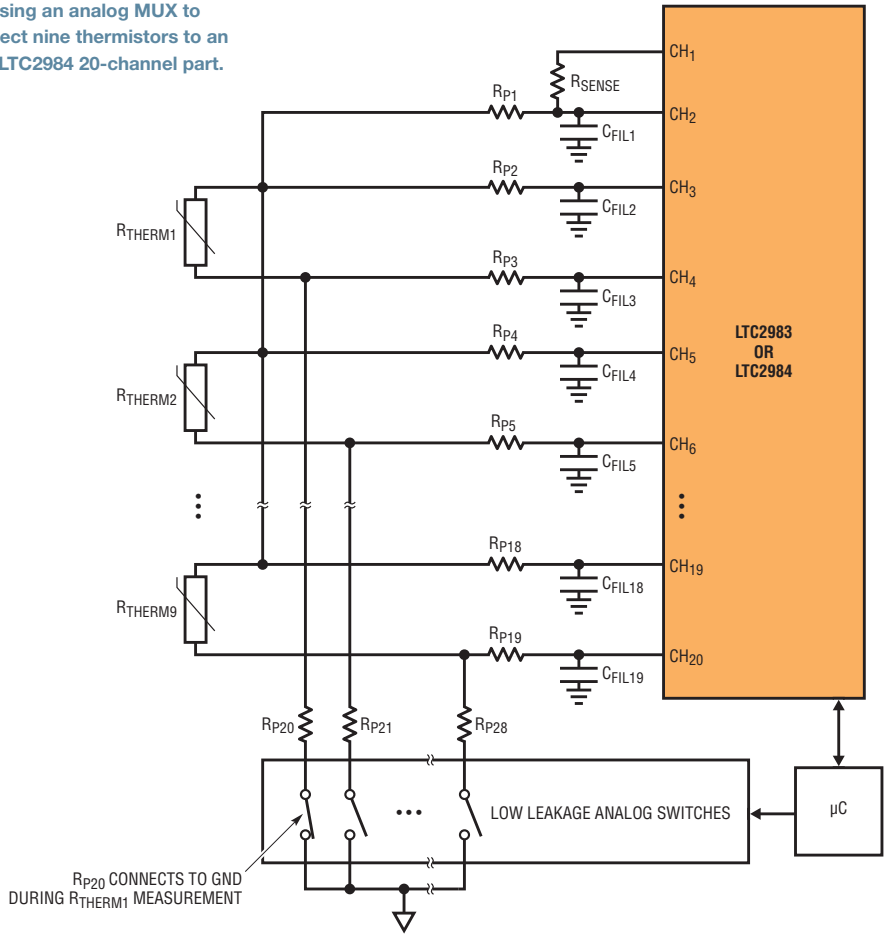

Table 1. Temp-to-bits converters for industrial applications

	LTC2983	LTC2984	LTC2986	LTC2986-1
Channel count	20	20	10	10
Any sensor any input	✓	✓	✓	✓
Automatic burnout and fault detection	✓	✓	✓	✓
Built-in and custom sensor coefficients	✓	✓	✓	✓
Automatic cold junction compensation—uses any sensor	✓	✓	✓	✓
Universal protected inputs (2- 3- and 4-wire)			✓	✓
Custom table lookup for analog voltage sensors			✓	✓
EEPROM		✓		✓
Package	7mm × 7mm LQFP-48			

How to Choose Cool Running, High Power, Scalable POL Regulators and Save Board Space

Afshin Odabaee

The art of designing efficient and compact DC/DC converters is practiced by a select group of engineers with a deep understanding of the physics and supporting mathematics involved in conversion design, combined with a healthy dose of bench experience. A deep understanding of Bode plots, Maxwell's equations and concerns for poles and zeros figure into elegant DC/DC converter design. Nevertheless, IC designers often escape dealing with the dreaded topic of heat—a job that usually falls to the package engineer.

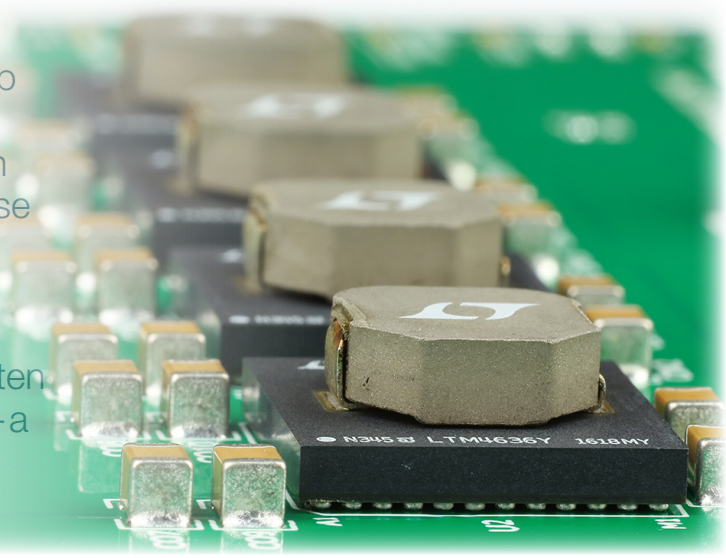

Heat is a significant concern for point-of-load (POL) converters, where space is tight among delicate ICs. A POL regulator generates heat because no voltage conversion is 100% efficient (yet). How hot does the package become due to its construction, layout and thermal impedance? Thermal impedance of the package not only raises temperature of the POL regulator, it also increases the temperature of the PCB and surrounding components, contributing to the complexity, size and cost of the system's heat removal arrangements.

Heat mitigation for a DC/DC converter package on a PCB is achieved through two major strategies:

- **Distribute it through the PCB:** If the converter IC is surface mountable, the heat conductive copper vias and layers in the PCB disperse the heat from the bottom of the package. If the thermal impedance of the package to the PCB is low enough, this is sufficient.
- **Add airflow:** Cool airflow removes heat from the package (or more precisely, the heat is transferred to the cooler fast air molecules in contact with the surface of the package).

Of course, there are methods of passive and active heat sinking, which, for simplicity of this discussion, are considered subsets of the second category.

When faced with rising component temperatures, the PCB designer can reach into the standard heat-mitigation toolbox for commonly used tools: additional copper, heat sinks or bigger and faster fans, or simply more space—use more PCB real estate and increase the distance between components on the PCB, or thicken the PCB layers.

Any of these tools can be used on the PCB to maintain the system within safe temperature limits, but applying these remedies can diminish the end product's competitive edge in the market. The product, say a router, might require a larger case to accommodate necessary component separation on the PCB, or it may become relatively noisy as faster fans are added to increase airflow. This can render the end product inferior in a market where companies compete on the merits of compactness, computational power, data rates, efficiency and cost.

Successful thermal management around high power POL regulators requires choosing the right regulator, which demands careful research. This article shows how a regulator choice can simplify the board designer's job.

DON'T JUDGE POL REGULATORS BY POWER DENSITY ALONE

A number of market factors drive the need to improve thermal performance in electronic equipment. Most obvious: performance continuously improves even as products shrink in size. For instance, 28nm–20nm and sub-20nm digital devices burn power to deliver performance, as innovative equipment designers use these smaller processes for faster, tinier, quieter

Stacked-inductor construction rewards system designers with a compact POL regulator, with the additional benefit of superior thermal performance.

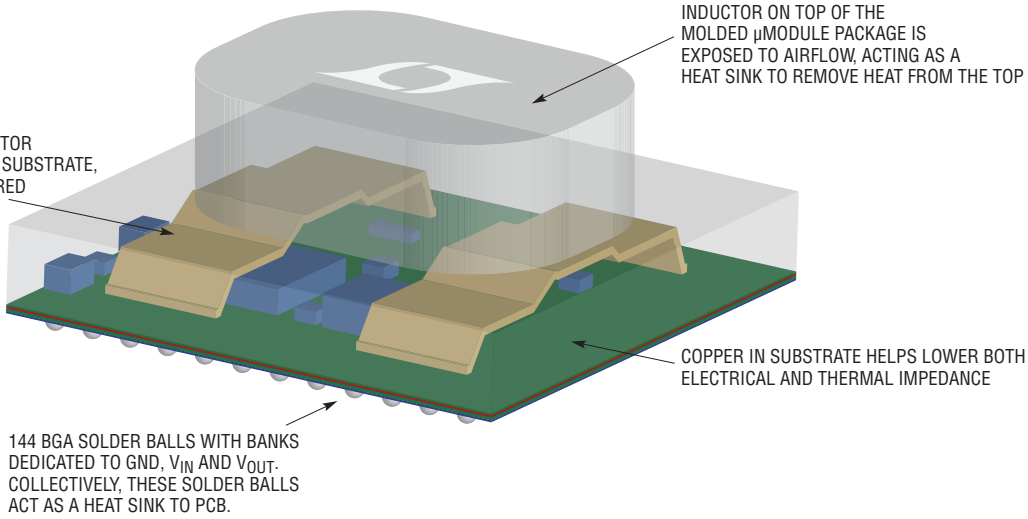

Figure 1. High power POL regulator module uses 3D (vertical) packaging technology to elevate the inductor and expose it to airflow as a heat sink. Remaining DC/DC circuitry is assembled on the substrate under the inductor, minimizing required PCB area while improving thermal performance.

and more efficient devices. The obvious conclusion from this trend is that POL regulators must increase in power density: (power)/(volume) or (power)/(area).

It is no surprise that power density is often cited in regulator literature as the headline specification. Impressive power densities make a regulator stand out from the pack—giving designers quotable specs when choosing from the vast array of available regulators.

Figure 2. The LTM4636's stacked inductor doubles as a heat sink to achieve impressive thermal performance in a complete POL solution with a small footprint.

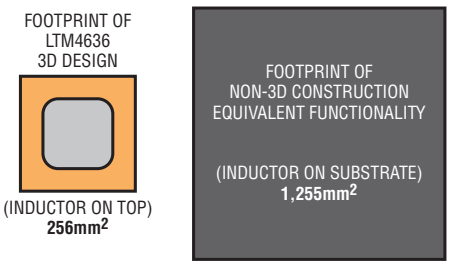

A 40W/cm2 POL regulator must be better than a 30W/cm2 regulator.

Product designers want to squeeze higher power into tighter spaces—superlative power density numbers appear at first blush to be the clear path to the fastest, smallest, quietest and most efficient products, akin to comparing automobile performance using horsepower. But, how significant is power density in achieving a successful final design? Less than you might think.

A POL regulator must meet the requirements of its application. In choosing a POL regulator, one must assure its ability to do the job on the PCB, where the treatment of heat can make or break the application. The following recommended step-by-step selection process for a POL regulator makes the case for prioritizing thermal performance:

1. **Ignore power density numbers:** Power density specifications ignore thermal derating, which has a significantly greater effect on the effective, real world "power density."
2. **Check the regulator's thermal derating curves:** A well documented and characterized POL regulator should have graphs specifying output current at various input voltages, output voltages and airflow speeds. The data sheet should show the output current capability of the POL regulator under real world operating conditions, so you can judge the regulator by its thermal and load current abilities. Does it meet the requirement of your system's typical and maximum ambient temperature and airflow speed? Remember, output current derating relates to the thermal performance of the device. The two are closely related, and equally important.

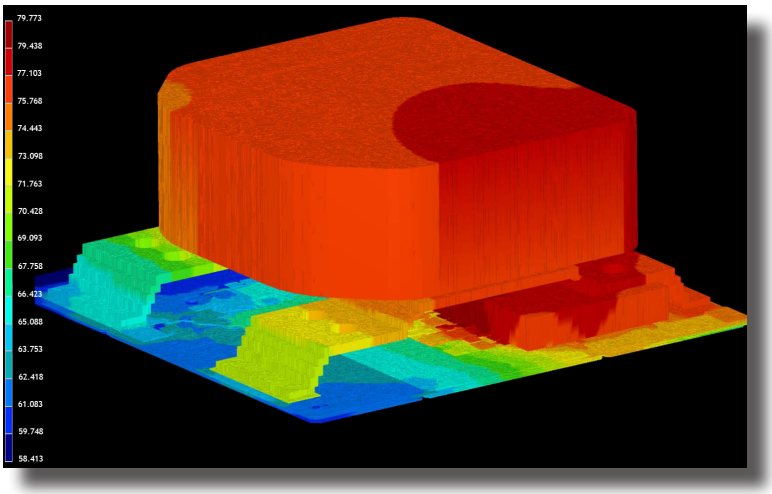

Figure 3. Modeled thermal behavior of LTM4636 shows heat is readily moved to the inductor package, which is exposed to airflow.

3. Look at efficiency: Yes, efficiency is *not the first* consideration. Efficiency results, when used exclusively, can present an inaccurate picture of the thermal characteristics of a DC/DC regulator. Of course, efficiency numbers are required to calculate input current and load current, input power consumption, power dissipation and junction temperature. But, efficiency values must be combined with output current derating and other thermal data related to the device and its package.

For example, a 98% efficient DC/DC step-down converter is impressive; even better when it boasts a superior power density number. Do you purchase it over a less efficient, less power dense regulator? A savvy engineer should ask about the effect of that seemingly insignificant 2% efficiency loss. How does that power translate into the package temperature rise during operation? What is the junction temperature of a high power density, efficient regulator at 60°C ambient with 200LFM (linear feet per minute, airflow)? Look beyond the typical numbers that are listed at room temperature of 25°C. What are the maximum and minimum values that are measured at the extremes: –40°C, 85°C or 125°C? At a high power density, does the package thermal impedance rise so high that the junction temperature shoots over the safe operating temperature? How much derating does an impressively efficient, but expensive, regulator require? Do derated output current values curtail output power capability to the point that the additional cost of the device is no longer justified?

4. Consider the ease of cooling the POL regulator: The package thermal impedance values provided in the data sheet are key to simulate and calculate rise in junction, ambient and case temperatures of the device. Because much of the heat in surface mount packages flow from the bottom of the package to the PCB, layout guidance and discussions about thermal measurements must be articulated in the data sheet to minimize surprises during system prototyping.

A well designed package should efficiently dissipate heat evenly throughout its surfaces, eliminating hot spots, which degrade the reliability of a POL regulator.

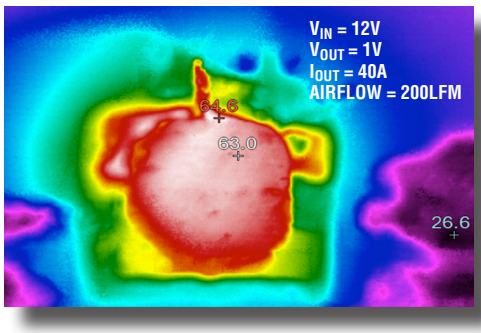

Figure 4. Thermal results of regulator at 40W shows temperature rise of only 40°C

As described above, the PCB is responsible for absorbing and routing much of the heat from surface mountable POL regulators. But with the prevalence of forced airflow in today's dense and complex systems, a cleverly designed POL regulator should also tap this "free" cooling opportunity to remove heat from heat generating components such as MOSFETs and inductors.

GUIDING HEAT TO THE TOP OF THE PACKAGE, AND INTO THE AIR

A high power switching POL regulator depends on an inductor or transformer to converter the input supply voltage to a regulated output voltage. In a nonisolated step-down POL regulator, the device uses an inductor. The inductor and accompanying switching elements, such as MOSFETs, produce heat during DC/DC conversion.

About a decade ago, a new packaging advances allowed an entire DC/DC regulator circuit including the magnetics to be designed and fitted inside molded plastic, called modules or SiP, where much of the heat generated inside the molded plastic is routed to the PCB via the bottom of the package. Any conventional attempt to improve heat removal capability of the package contributes to a larger package such as attaching a heat sink to the top of the surface mount package.

A few years ago, an innovative module packaging technique was developed to take advantage of available airflow to aid in cooling. In this package design, a heat sink is integrated into the module package and over molded. Inside the package, the bottom of the heat sink is directly connected to the MOSFETs and inductors, the heat generators, while the topside of the heat sink is a flat surface exposed at top of the package. This new intra-packaging heat sinking technique allows a device to be cooled quickly with airflow (for an example, see the LTM4620 TechClip videos, www.linear.com/solutions/4936).

GO VERTICAL: POL MODULE REGULATOR WITH STACKED INDUCTOR AS HEAT SINK

The size of an inductor in a POL regulator depends on voltage, switching frequency, current handling and its construction. In a module approach, where the DC/DC circuit including the inductor is overmolded and encapsulated in a plastic package and resembles an IC, the inductor dictates the thickness, volume and weight of the package more than any other component. The inductor is also a significant source of heat.

Integrating the heat sink into the package, the strategy discussed above, helps to conducts heat from the MOSFETs and inductor to the top of the package, where it can be dissipated to air, a cold plate or a passive heat sink. This technique is effective when relatively small, low current inductors easily fit inside the plastic mold compound of the package, but not so effective when POL regulators depend on larger and higher current inductors, where placement of the magnetics inside the package forces other circuit components to be farther apart, significantly expanding the PCB footprint of the package. To keep the footprint small while improving heat dissipation, the package engineers have developed another trick: vertical, stack or 3D (Figure 1).

3D PACKAGING WITH EXPOSED STACKED INDUCTOR: KEEP FOOTPRINT SMALL, INCREASE POWER, IMPROVE HEAT DISSIPATION

Small PCB footprint, more power and better thermal performance—all three are simultaneously possible with 3D packaging, a new method in construction of POL regulators (Figure 1). The LTM4636 is a µModule regulator with onboard DC/DC regulator IC, MOSFETs, supporting circuitry and a large inductor to decrease output ripple and deliver load currents up to 40A from 12V input to precisely regulated output voltages ranging from 3.3V to 0.6V. Four LTM4636 devices running in parallel can current share to provide 160A of load current. The footprint of the package is only 16mm × 16mm. Another regulator in the family, the LTM4636-1, detects overtemperature and input/output overvoltage conditions and can trip an upstream power supply or circuit breaker to protect itself and its load.

Horsepower advocates can calculate the power density of the LTM4636, and safely tout its numbers as impressive, but as discussed above, power density numbers tell an incomplete story. There are other significant benefits that this µModule regulator brings to the system designer's toolbox: superior thermal performance resulting from impressive DC/DC conversion efficiency and an unparalleled ability to disperse heat.

To minimize the regulator's footprint (16mm × 16mm BGA), the inductor is elevated and secured on two copper lead frame structures so that other circuit components (diodes, resistors, MOSFETs, capacitors, DC/DC IC) can be soldered under it on the substrate. If the inductor is placed on the substrate, the µModule regulator can easily occupy more than 1,225mm² of PCB, instead of small 256mm² footprint (Figure 2).

Stacked-inductor construction rewards system designers with a compact POL regulator, with the additional benefit of superior thermal performance. The stacked inductor in the LTM4636 is not overmolded (encapsulated) with plastic, as are the rest of the components. Instead, it is exposed directly to airflow. The shape of the inductor casing incorporates rounded corners for improved aerodynamics (minimal flow blockage).

THERMAL PERFORMANCE AND EFFICIENCY

The LTM4636 is a 40A-capable µModule regulator benefiting from 3D packaging technology, or component-on-package (CoP), as shown in Figure 1. The body of the package is an overmolded 16mm × 16mm × 1.91mm BGA package. With the inductor stacked on top of the molded section, the LTM4636's total package height, from the bottom of BGA solder balls (144 of them) to the top of the inductor, is 7.16mm.

In addition to dissipating heat from the top, the LTM4636 is designed to efficiently disperse heat from the bottom of the package to the PCB. It has 144 BGA solder balls with banks dedicated to GND, V_{IN} and V_{OUT} where high current flows. Collectively, these solder balls act as a heat sink to the PCB. The LTM4636

Figure 5. Thermal derating shows full current of 40A delivered up to 83°C ambient, 200LFM

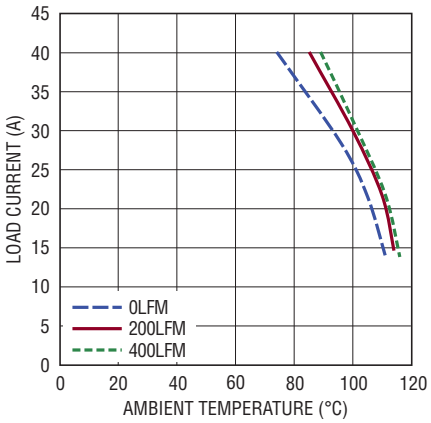

Figure 6. High DC/DC conversion efficiency over a variety of output voltages

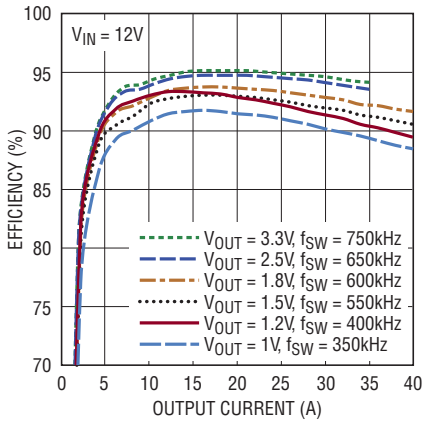

Figure 7. It is easy to lay out parallel LTM4636s. Simply duplicate the layout of one channel.

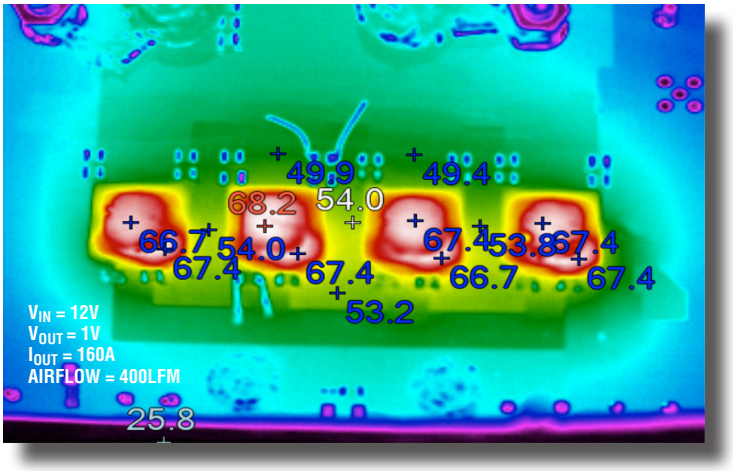

Figure 8. Precision current sharing among four LTM4636s running in parallel, resulting in only 40°C rise in temperature for 160A application.

is optimized to dissipate heat from both the top and bottom of the package.

Even operating with a significant conversion ratio, 12V input/1V output, and at a full load current of 40A (40W) and standard 200LFM airflow, the LTM4636 package temperature rises only 40°C over ambient temperature (25°C–26.5°C). Figure 4 shows the thermal image of the LTM4636 under these conditions.

Figure 5 shows the output current thermal derating results. At 200LFM, the LTM4636 delivers an impressive full current of 40A up to an 83°C ambient temperature. Half-current, 20A, derating only occurs at an excessively high ambient temperature of 110°C. This allows the LTM4636 to perform at high capacity as long as some airflow is available.

Figure 9. Efficiency of 140W 4-µModule regulator

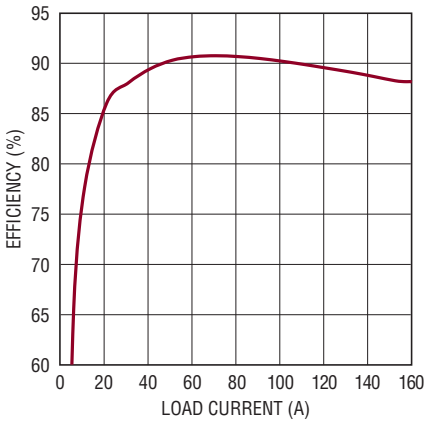

The high conversion efficiency shown in Figure 6 is mainly a result of top performing MOSFETs and strong drivers of the LTM4636. For example, a 12V input supply step-down DC/DC controller achieves:

- 95% for 12V input to 3.3V, 25A
- 93% for 12V input to 1.8V, 40A
- 88% for 12V input to 1V, 40A

140W, SCALABLE 4 × 40A µMODULE POL REGULATOR WITH THERMAL BALANCE

One LTM4636 is rated for 40A load current delivery. Two LTM4636s in current sharing mode (or parallel) can support 80A, while four will support 160A. Upscaling a power supply with parallel LTM4636s is easy: simply copy and paste the single-regulator footprint, as shown in Figure 7 (symbols and footprints are available).

The current mode architecture of the LTM4636 enables precision current sharing among the 40A blocks. Precise current sharing, in turn, produces a power supply that spreads the heat evenly between devices. Figure 8 shows that all devices in the 4-µModule 160A regulator operate within a degree C of each other, ensuring that no individual device is overloaded or overheated. This greatly simplifies heat mitigation.

Figure 10 shows the complete 160A design. Note that no clock device is required for the LTM4636s to operate out-of-phase respective of each other—clocking and phase control is included. Multiphase operation reduces output and input ripple current, reducing the number of required input and output capacitors. Here, the four LTM4636s in Figure 10 are running 90° out-of-phase.

CONCLUSION

Choosing a POL regulator for a densely populated system requires scrutiny beyond voltage and amperage ratings of the device. Evaluation of its package's thermal characteristics is essential, as it determines cost of cooling, cost of PCB and final product size. Advances in 3D, also referred to as stacked, vertical, CoP, allow high power POL module regulators to fit a small PCB footprint, but more importantly, enable efficient cooling. The LTM4636 is the first series of µModule regulators to benefit from this stacked packaging technology. As a 40A POL µModule regulator with stacked inductor as heat sink, it boasts 95% to 88% efficiency, with only a 40°C rise at full load, occupying only 16mm × 16mm of PCB area. A video description of the LTM4636 is available at www.linear.com/LTM4636. ■

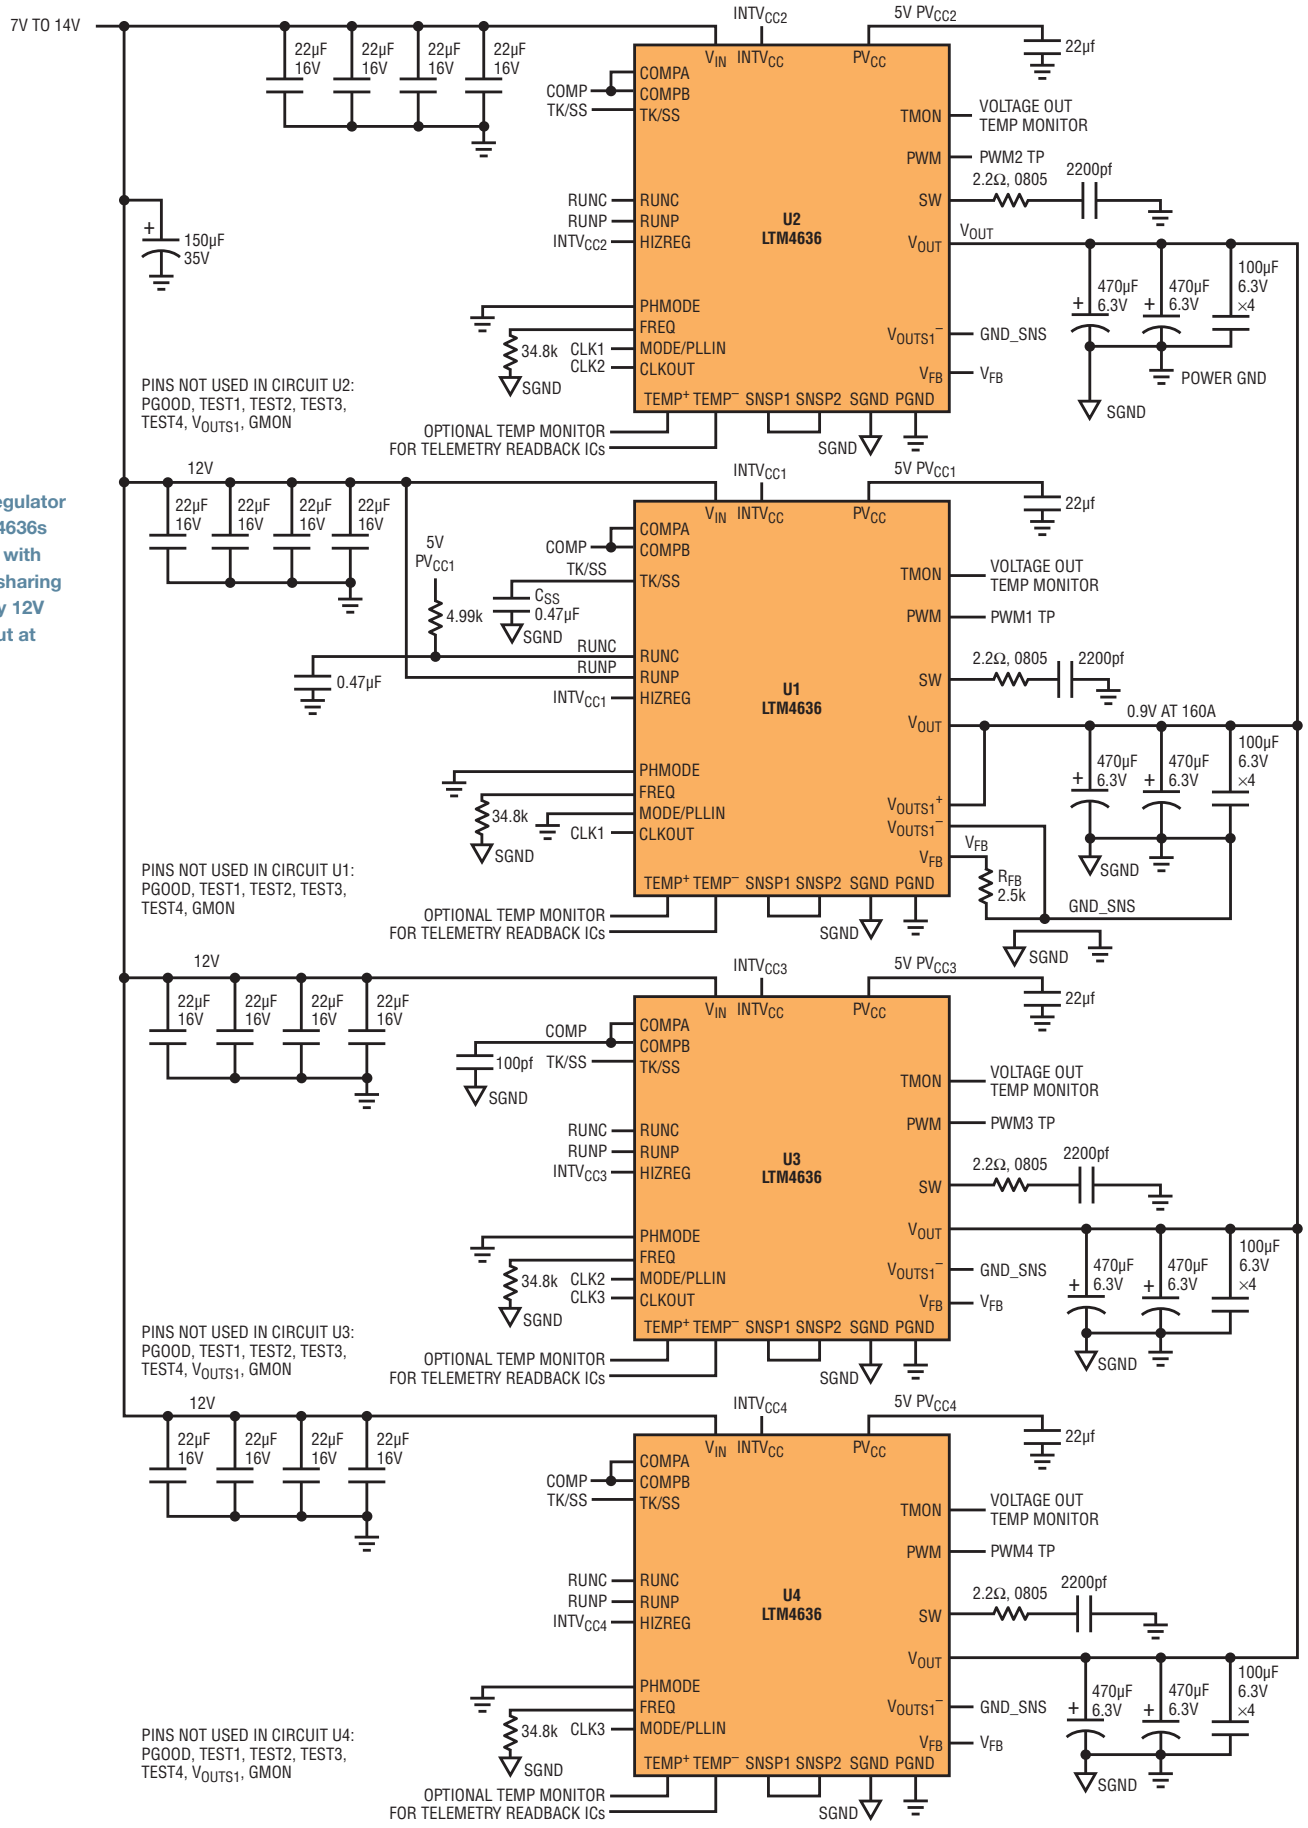

Figure 10. 140W regulator features four LTM4636s running in parallel with precision current sharing and high efficiency 12V input to 0.9V output at 160A

What's New with LTspice?

Gabino Alonso

LTspice Blog
www.linear.com/solutions/LTspice

twitter—Follow @LTspice at *www.twitter.com/LTspice*
—Like us at *www.facebook.com/LTspice*

NEW VIDEO: Creating and Working with Symbols

Markus Holtkamp walks you through the process of creating a symbol automatically using either a subcircuit definition in ASCII fomat or creating a schematic hierarchy. Hierarchical schematic drafting allows large circuits to be drafted in a single schematic, with component circuits presented in abstract, i.e. black boxes. www.linear.com/solutions/7723

SELECTED DEMO CIRCUITS

For a complete list of example simulations, please visit www.linear.com/democircuits.

Linear Regulators

- **LT3045:** Low noise, high PSRR linear regulator (3.8V–20V to 3.3V at 500mA) www.linear.com/solutions/7413

Buck Regulators

- **LT8608:** 2MHz low EMI high voltage synchronous buck regulator (5.5V–42V to 5V at 1.5A) www.linear.com/solutions/7313
- **LT8630:** High efficiency µPower buck regulator (13V–100V to 12V at 600mA) www.linear.com/solutions/7678
- **LTC3864:** 60W PMOS step-down converter with 100% duty cycle capability (12V–60V to 12V at 5A) www.linear.com/solutions/7776
- **LTM4632:** Triple output ultrathin buck regulator for DDR-QDR4 (3.6V–15V to 1.5V at 3A, 0.75V at ±3A, 0.75V at 10mA) www.linear.com/solutions/7240
- **LTM4642:** Wide input voltage, high efficiency, dual buck regulator (4.5V–20V to 1.8V at 4A, 1.2V at 4A) www.linear.com/solutions/7210
- **LTM4650:** High efficiency dual 25A step-down regulator with output tracking (4.5V–15V to 1.5V at 25A & 1V at 25A) www.linear.com/solutions/7379
- **LTM4650-1:** High efficiency 8-phase 200A step-down regulator (4.5V–15V to 1V at 200A) www.linear.com/solutions/7381
- **LTM4677 & LTM4650:** High current, parallel µModule buck regulators with power system management (4.5V–16V to 1V at186A) www.linear.com/solutions/7353
- **LTM8064:** CVCC source/sink step-down regulator (7.5V–58V to 5V at ±6A) www.linear.com/solutions/7225

Buck-Boost, Boost, SEPIC & Inverting Regulators

- **LT8390:** High efficiency 250W buck-boost regulator (9V–36V to 12V at 25A) www.linear.com/solutions/7645
- **LTM8049:** ±12V SEPIC & inverting regulator (2.8V–18V to 12V at 1A & -12V at 1A) www.linear.com/solutions/7610

Isolated Converters

- **LT3752-1/LT8311:** 200W active clamp forward converter with synchronous rectification (150V–400V to 12V at 16.7A) www.linear.com/solutions/4956
- **LT3752/LT8311:** Active clamp forward converter with synchronous rectification (36V–72V to 12V at 12A) www.linear.com/solutions/4698
- **LT3753/LT1431:** 80W active clamp non-synchronous forward converter for PoE (10V–54V to 54V at 1.5A) www.linear.com/solutions/5885
- **LT8303:** µPower no-opto isolated flyback converter (36V–72V to 5V at 0.65A–0.84A) www.linear.com/solutions/7500
- **LT8304:** µPower no-opto isolated flyback converter (18V–72V to 5V at 2A) www.linear.com/solutions/7297

LED Drivers

- **LT3922:** Low EMI, high efficiency boost LED driver (4V–28V to 34V LED at330mA) www.linear.com/solutions/7425
- **LT8391:** High efficiency 50W buck-boost LED driver (4V–60V to 25V LED at 2A) www.linear.com/solutions/7277

Op Amps

- **LT6018/LT1678:** Low impedance source, high common mode range amplifier www.linear.com/LT6018
- **LTC6362:** Baseband design example for a low power IQ modulator www.linear.com/solutions/7116

SELECT MODELS

To search the LTspice® library for a particular device model, press F2. To update to the current version, choose Sync Release from the Tools menu.

Buck Regulators

- **LT8609B:** 42V, 2A/3A peak synchronous step-down regulator with 2.5µA quiescent current www.linear.com/LT8609

- **LTC3884:** Dual output PolyPhase® step-down controller with sub-milliohm DCR sensing and digital power system management www.linear.com/LTC3884

Buck-Boost, Boost, SEPIC & Inverting Regulators

- **LTC3130:** 25V, 600mA buck-boost DC/DC converter with 1.6µA quiescent current www.linear.com/LTC3130
- **LTC3896:** 150V low IQ, synchronous inverting DC/DC controller www.linear.com/LTC3896
- **LTC3897:** PolyPhase synchronous boost controller with input/ output protection www.linear.com/LTC3897
- **LTC7813:** Low IQ, 60V synchronous boost + buck controller www.linear.com/LTC7813

Energy Harvesting

- **LTC3107:** Ultralow voltage energy harvester and primary battery life extender www.linear.com/LTC3107

Hot Swap, Surge Stopper & Protection Controllers

- **LTC4368:** 100V UV/OV and reverse protection controller with bidirectional circuit breaker www.linear.com/LTC4368
- **LTC4420:** 18V dual input µPower PowerPath™ prioritizer with backup supply monitoring www.linear.com/LTC4420
- **LTM9100:** Anyside™ high voltage isolated switch controller with I²C command and telemetry www.linear.com/LTM9100

Op Amp

- **LTC6261:** 30MHz, 240µA power efficient rail-to-rail I/O op amps www.linear.com/LTC6261 ■

Power User Tip

USING .MEAS AND .STEP TO CALCULATE EFFICIENCY

Predicting the efficiency of an application is vital to evaluating design trade-offs of a switching mode power supply. Two useful tools, the .step and .meas commands, can be used to calculate and plot efficiency over a range of load currents.

To evaluate efficiency, clearly label your input and output voltage net as IN and OUT, respectively. Press **F4** to place net names. Replace your resistive load with a independent current source as shown below and define the value using a global variable {Iload}. Press **F2** and type load2 in the search box to select and place the component. Edit the component value by right-clicking on the symbol and enter the variable {Iload}. Note the names of the input voltage source (V1) and the load current source (I1).

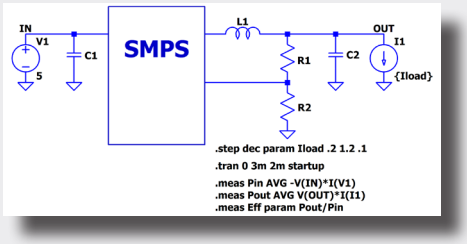

The .step command is useful for sweeping a variable across a range of values in a single simulation run. The variable can be temperature, a model parameter, a global parameter or in our case an independent source. These steps can be defined as linear, logarithmic or as a list of specific values.

Insert a .step command into your schematic as a SPICE directive and step the independent current source from a light load to maximum current load and define the step increments. (You can use the **S** hotkey to add and place a SPICE directive.)

Here, we step the independent current source, I1, from 0.2A to 1.2A in 0.1A increments using a global variable defined as {Iload}. The param directive is used here in the .step command to allow for the creation of this user-defined variable. Please see the help file (F1) for more details on .step command and param directive.

```
.step param Iload .2 1.2 .1
```

It is important to calculate efficiency when the circuit is operating in a steady state. To ensure this, simulate your circuit, and note when steady state is achieved for all conditions in the .step command, extending simulation stop time if needed. Use this observation to set the the "time to start saving data" and "stop time" to encompass a short duration within the observed steady state period. In the .tran statement below, we start to save data at 2ms and stop at 2.1ms.

```
.tran 0 2.1m 2m startup
```

The .meas command is useful for measuring a range over the abscissa (as well as one point on the abscissa). Add the following expressions as a SPICE directive to calculate the average input power (Pin), average output power (Pout) and the efficiency (Eff). Please note the current direction convention for the input voltage source, V1, is *into* the device, hence the negative sign in the Pin calculation. The final expression calculates the efficiency using the param directive for clarity. Run your simulation. Please see the help file (F1) for more details on .meas command.

```
.meas Pin AVG -V(IN)*I(V1)
.meas Pout AVG V(OUT)*I(I1)
.meas Eff param Pout/Pin
```

Once the simulation completes, right-click one of the windows, select view and select Spice Error Log (or use the **Ctrl**+**L** hotkey). The Spice Error Log contains data points for the .meas statements to include the efficiency calculations.

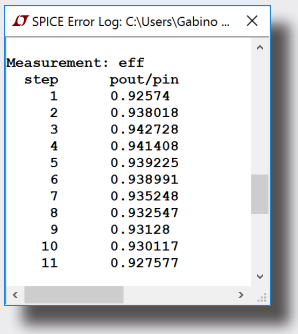

One neat feature of LTspice is the ability to plot the stepped .meas data over the abscissa (Iload). To plot the data, right-click the error log and select the Plot step'ed .meas data, right-click on the blank screen to select Add Trace (or use **Ctrl**+**A**) and select Eff. This will display the efficiency calculation over stepped load current.

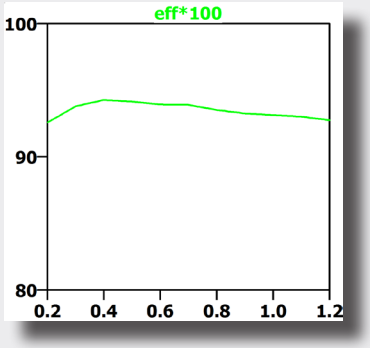

Of course, calculating efficiency is only one example. The .meas and .step commands can be combined in countless other ways to characterize your analog circuit designs.

Happy simulations!

Making Accurate Temperature Measurements with Devices with Different Ideality Factors

Hamza Salman Afzal and Shuilin Tian

The ability to accurately measure temperature is critical to the reliable operation of power systems. Accurate temperature data can reveal the instantaneous health of a system, and enable real-time protection functions such as thermal shutdown. For instance, significant rates of change in temperature can indicate increased power consumption. If a proper temperature feedback loop is implemented, this information can be used to reduce power consumption before temperatures escalate to dangerous levels.

Perhaps the least expensive and most ubiquitous temperature sensor is a diode. A silicon p-n junction has a forward drop of approximately 700mV and a temperature coefficient of –2.2mV/C at room temperature. This sensitivity to temperature is exploited in many systems as a means of measuring temperature by simply biasing the diode with a constant current and measuring the resulting voltage (see Figure 1)

This method has achieved wide acceptance and popularity as many microprocessors, FPGAs, DC/DC converters and other high power devices include diode junctions (often in the form of bipolar transistors) for the purpose of monitoring temperature within the device itself.

THEORY OF OPERATION

The forward voltage drop of a diode junction as a function of temperature is:

$$V_D = \eta \bullet V_T \bullet \ln\frac{I_D}{I_S} \quad (1)$$

where I_D is the diode current, V_D is the diode voltage, I_S is the reverse saturation current (process dependent parameter) and η is the ideality factor (typically close to 1.0). V_T is defined as:

$$V_T = \frac{k \bullet T}{q}$$

where T is the diode junction temperature in Kelvin, k is the Boltzmann constant and q is the electron charge. V_T is approximately 26mV at room temperature and scales linearly with Kelvin temperature.

I_S increases with temperature, causing the $\ln(I_D/I_S)$ term in equation 1 to drop with temperature, yielding an approximately –2.2mV/°K composite diode voltage slope.

Unfortunately, the accuracy of these measurements depends heavily on I_S, which varies from production lot to lot. Additionally, the initial offset must be calibrated out. This variability makes an accurate absolute value of temperature nearly impossible using a single point, forward voltage measurement.

As η is a constant, to obtain a linear voltage proportional to temperature the I_S term is the only term that must be cancelled out. This can be achieved by

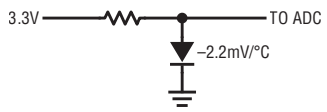

Figure 1. A resistor from 3.3V, or more, provides sufficiently constant current

measuring the change in diode voltage at two different current levels where

$$\Delta V_D = \left(\eta \bullet V_T \bullet \ln\frac{I_{C1}}{I_S}\right) - \left(\eta \bullet V_T \bullet \ln\frac{I_{C2}}{I_S}\right)$$

$$\Delta V_D = \eta \bullet V_T \bullet \ln\frac{I_{C1}}{I_{C2}}$$

If

$$\frac{I_{C1}}{I_{C2}} = 10$$

then

$$\Delta V_D = \eta \bullet V_T \bullet \ln 10$$

As

$$V_T = \frac{k \bullet T}{q}$$

$$T = \frac{\Delta V_D}{\frac{\eta \bullet k}{q} \bullet \ln 10}$$

If we take a diode measurement at two different currents with a ratio of 10, the resulting voltage is 198µV per Kelvin. The absolute value of the test currents is not important, rather it is the ratio of the two currents that dictates ΔV_D. This technique eliminates dependence on I_S, leaving the ideality factor as the lone variable in producing a temperature from a measured voltage, where the ideality factor is essentially a constant for a

Perhaps the least expensive and most ubiquitous temperature sensor is a diode. A silicon p-n junction has a forward drop of approximately 700mV and a temperature coefficient of –2.2mV/C at room temperature. This sensitivity to temperature is exploited in many systems as a means of measuring temperature by simply biasing the diode with a constant current and measuring the resulting voltage.

given junction type, with minor variation among different junction types.

SIMPLIFY TEMPERATURE MEASUREMENTS

Although the ΔV_D method of measuring temperature appears simple, its use is complicated by the fact that a number of accurate data points must be collected and used in real-time calculations, namely:

- the diode is driven by two, precisely ratioed, currents
- two precise measurements of the corresponding voltages must be made
- the known current values and voltage results are recorded and the resulting temperature calculated.

Temperature monitor ICs, such as the LTC2990, LTC2991 and LTC2997 take care of these complications, while using the differential measurement technique above to eliminate errors and provide temperature measurements within 1°C of accuracy. They also compensate for series resistance errors by taking V_D measurements at multiple operating currents. They measure remote and ambient temperature as well as voltage and current.

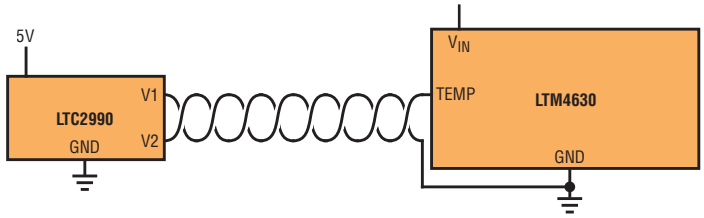

Figure 2. Block diagram of the LTC2990 interfacing with LTM4630

The LTC2997 provides an analog output and is ideal for overtemperature alarms and measurement. No calibration is necessary. Ratioed currents and delta voltage measurements are all made inside the part, and a voltage result is presented at the output, with a sensitivity of 4mV/°C.

Like the LTC2997, the LTC2990 and LTC2991 make all necessary measurement and temperature calculations internally. They feature four and eight measurement channels, respectively. Both devices feature an onboard ADC, which can be used to make single-ended or differential voltage measurements. They also measure temperature using an internal diode, or external, remote diodes using the measurement channels. They communicate all measurement results through an I²C interface.

Dealing with Ideality Factor

All three devices are designed to measure temperature with an accuracy of better than 1°C, using the base emitter junction of devices with an ideality factor of 1.004 such as the widely used MBBT3904. If the ideality factor of the target sensor differs from 1.004, it can be compensated in software using the following equation

$$T_{ACTUAL} = \frac{T_{MEASURED} \bullet \eta_{SENSOR}}{1.004} \quad (2)$$

It is important to consider this parameter when selecting a temperature sensor, as it could produce errors in the absolute temperature measurement. An ideality factor with a ±1% deviation from 1.004 would result in an error of 2.7°C at 0°C and an error of 4°C at 100°C. However, for most junctions, η_{SENSOR} error is <1% and contributes less error than device unit to unit variation.

For example, Figure 2 shows an application where the LTC2990 is interfacing with an LTM4630 µModule regulator, which has a TEMP pin with an internal PNP diode.

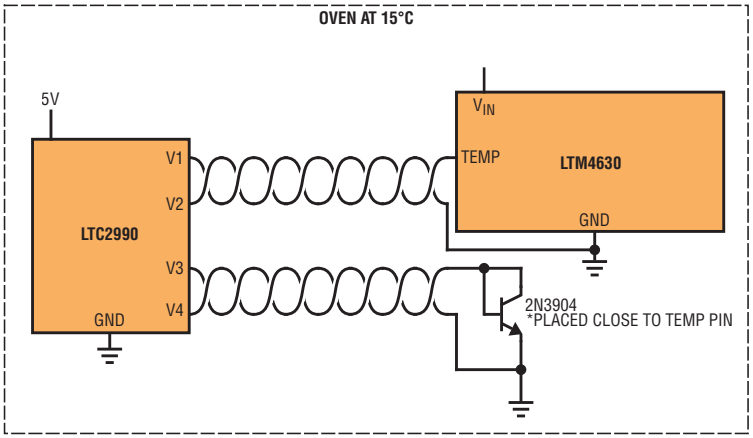

Figure 3. Test setup with a 2N3904 diode

Like the LTC2997, the LTC2990 and LTC2991 make all necessary measurement and temperature calculations internally. They feature four and eight measurement channels, respectively. Both devices feature an onboard ADC, which can be used to make single-ended or differential voltage measurements. They communicate all measurement results through an I^2C interface.

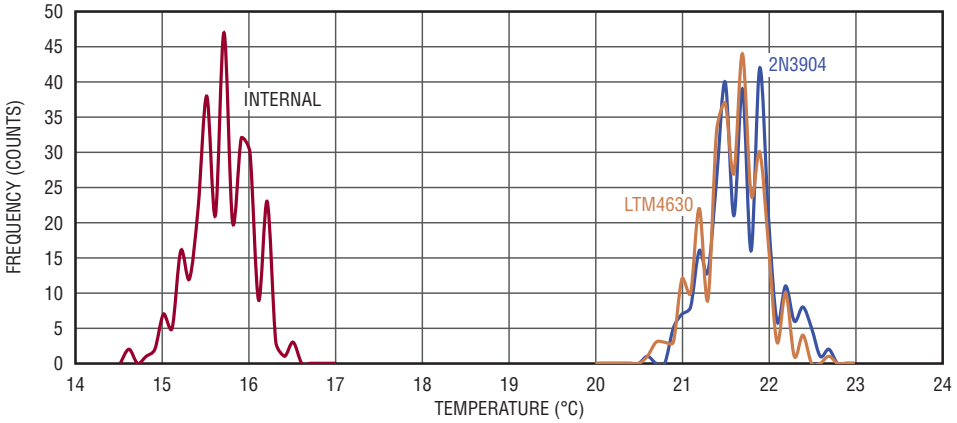

Figure 4. Temperature measurement comparison of the LTC2990 internal sensor, a 2N3904 diode and the LTM4630 TEMP pin (connected to internal PNP diode)

The ideality factor of this diode is 1.008, which is only a deviation of 0.4% from 1.004. As a result, absolute temperature readings from the LTM4630 while interfacing with the LTC2990 are well within 1°C of accuracy until the temperature exceeds 250°C, which is well beyond the operating range of both devices.

To demonstrate this, an experiment was conducted using the setup shown in Figure 3. A 2N3904 NPN transistor was placed very close to the TEMP pin so that it would measure the same local temperature as the junction present inside the LTM4630. Both the LTC2990 and the LTM4630 unit were placed inside an oven and heated to 15°C, 25°C, 40°C, 60°C and 80°C. The LTM4630 module was powered, but no load was present so that the internal temperature of the LTM4630 does not rise above the local board temperature. Temperature measurements are made using the LTC2990's internal sensor, the 2N3904 and the LTM4630 TEMP pin. Figure 4 shows a histogram of the measured data at 15°C.

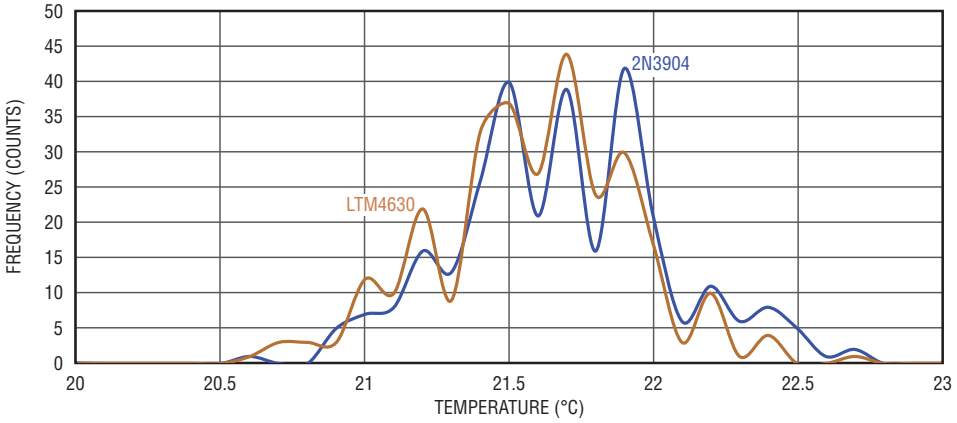

Figure 5. Close-up of LTC2990 temperature measurements of the 2N3904 and the LTM4630

The internal sensor of the LTC2990 shows the oven's ambient temperature of approximate 15.5°C. The LTM4630 module is running a few degrees higher than the ambient temperature due to it being powered. As a result, both the 2N3904 and LTM4630 read a temperature of approximately 21.6°C. A zoomed in view of the 2N3904 and LTM4630 readings is shown in Figure 5.

The correlation between the two sensors is easily seen as the distribution centers around 21.8°C for both sensors and is within 1°C of each other, which validates our previous claim of the error being within 1°C of each other due to the ideality factor deviation of only 0.4%. No correction for η is necessary.

Oversampling for Better Resolution

The widths of the curves represent the ADC quantization noise, which can easily be filtered out through software. While filtering can be done using a running average, it requires the processor to store data over the number of samples being averaged. For processors with memory constraints this might not be desirable. Figure 6 shows the pseudocode used to implement a lowpass filter.

Filtering is done by adding new temperature data and then subtracting the mean each time. The advantage of this technique is that the history of data need not be stored. If N is a power of 2, the division

Our temperature monitors make temperature measurement simple and convenient. Measurements are independent of reverse saturation current and series resistance. Errors introduced by interfacing with devices of different ideality factors, while often negligible, can be precisely calibrated, allowing use of temperature monitors with a wide variety of diode temperature sensors.

can be accomplished by a simple shift. Readings were taken again, with Figure 7 showing the filtered data distribution. A code example using the DC2026 Linduino® along with the LTC2990 is available at www.linear.com/docs/57871.

Notice how the temperature peaks are offset by 0.1°C. This error is due to the ideality factor error of 0.4%. If we compensate for the ideality factor using Equation 2, the two peaks line up nicely, as shown by Figure 8. Compensation ensures that the sensors with different η values can produce the same absolute temperature under the same conditions.

SUMMARY

Our temperature monitors make temperature measurements simple and convenient. Measurements are independent of reverse saturation current and series resistance. Errors introduced by interfacing with devices of different ideality factors, while often negligible, can be precisely calibrated, allowing use of temperature monitors with a wide variety of diode temperature sensors. ■

```
void low_pass_measurement(){
    uint8_t N = 5; //Number of Samples Per Average
    float filteredReading; //Final value
    float runningAverage = 0; //Initial Seed Value For Filter
    do{
    float temperature;
        temperatureReading = readFromSensor();
        runningAverage = runningAverage + temperature - runningAverage/N;
        filteredReading = runningAverage/N; //Filtered Reading
        Serial.print("Low Pass Reading: ");
        Serial.println(filteredReading);
    }
}
```

Figure 6. C psuedocode implementation of lowpass filtering. Sample code available at www.linear.com/docs/57871.

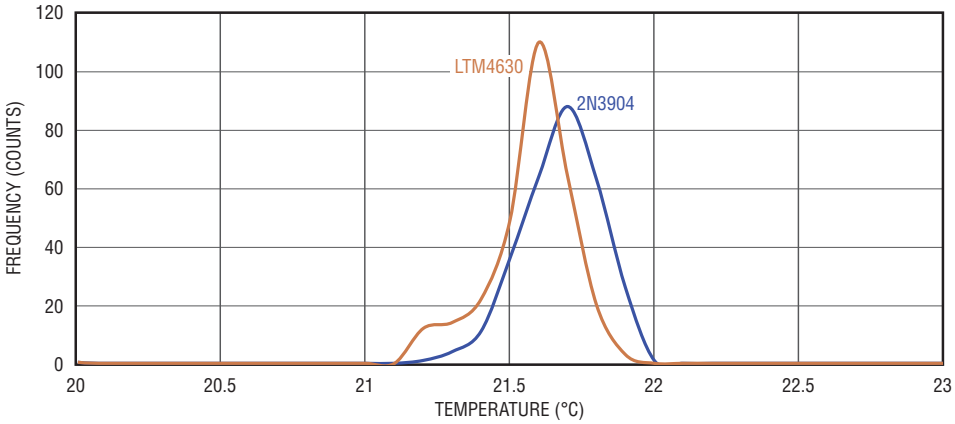

Figure 7. Filtered LTC2990 temperature measurement comparison between 2N3904 And LT4630

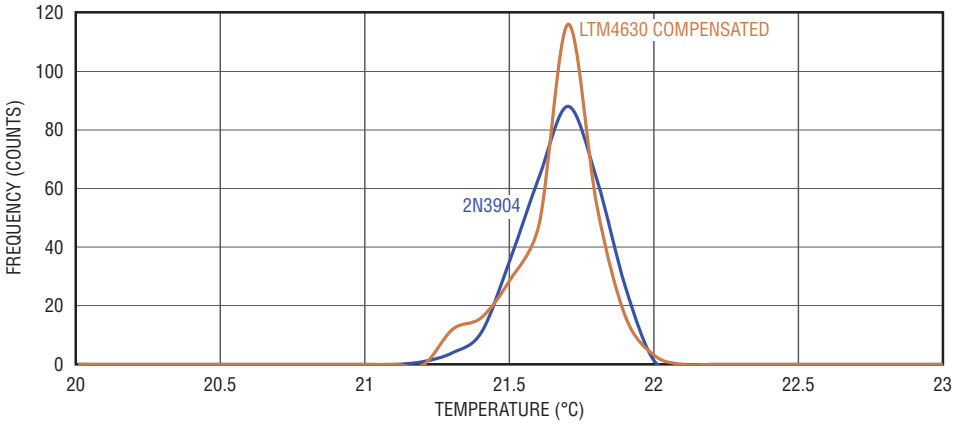

Figure 8. LTC2990 temperature measurement comparison between a 2N3904 and a software-calibrated LTM4630

Multi-Chemistry Battery Charger Supports Maximum Power Point Tracking for Solar Panels

Trevor Barcelo

The LTC4015 is a versatile synchronous step-down charger capable of supporting a variety of battery chemistries including lead-acid, Li-ion and $LiFePO_4$. The LTC4015 features an extensive list of battery charging functions, including coulomb counting and an array of battery and system monitor capabilities. However, this article focuses on its input control loop, which enables solar panel maximum power point tracking (MPPT) functionality.

For readers who have not yet encountered the general concept of MPPT, or could use a knowledge refresh, visit www.linear.com/solutions/4545. Regardless of your general knowledge of MPPT, to understand the LTC4015's implementation, it is important to understand the LTC4015's multi-control-loop operation.

BASIC DEVICE OPERATION

The LTC4015 charges batteries with a peak current mode synchronous step-down controller driving MN2 and MN3 (see Figure 1). The controller can regulate four parameters: input voltage (using the UVCLFB pin), input current (CLP and CLN), battery charge voltage (BATSENS) and battery charge current (CSP and CSN). Both peak inductor current control and battery charge current regulation are accomplished with sense resistor R_{SNSB}. In addition to these two functions, R_{SNSB} allows the LTC4015 to monitor battery charge and discharge current, battery ESR and battery coulomb count. The input voltage regulation is an integral part of MPPT operation and is discussed in detail in the next section.

The LTC4015 uses an ideal diode-OR PowerPath architecture to seamlessly interface both the input supply and the battery to the system load. Ideal diode MN1 connects V_{IN} to V_{SYS} if V_{IN} is larger than V_{CSP} (battery voltage) while MP1 connects the battery to V_{SYS} if V_{CSP} is larger than V_{IN}. In addition to powering the system from V_{IN}, the two diode controllers work with the charger to provide power from the battery to the system without back driving V_{IN} and guarantee that power is available to the system even if there is insufficient or absent power from V_{IN}.

When limited power is available to the switching charger because either the programmed input current limit (input current regulation) or input undervoltage limit (input voltage regulation) is active, charge current is automatically reduced to prioritize power delivery to the system load. However, it is important to note that the LTC4015 only limits charge current, but does not limit current from the input to the system load—if the system load alone requires more power than is available from the input after charge current has been reduced to zero, V_{SYS} must fall to the battery voltage in order for the battery to provide supplemental power.

Figure 1. Simplified LTC4015 application topology (not necessarily optimized for solar panel input)

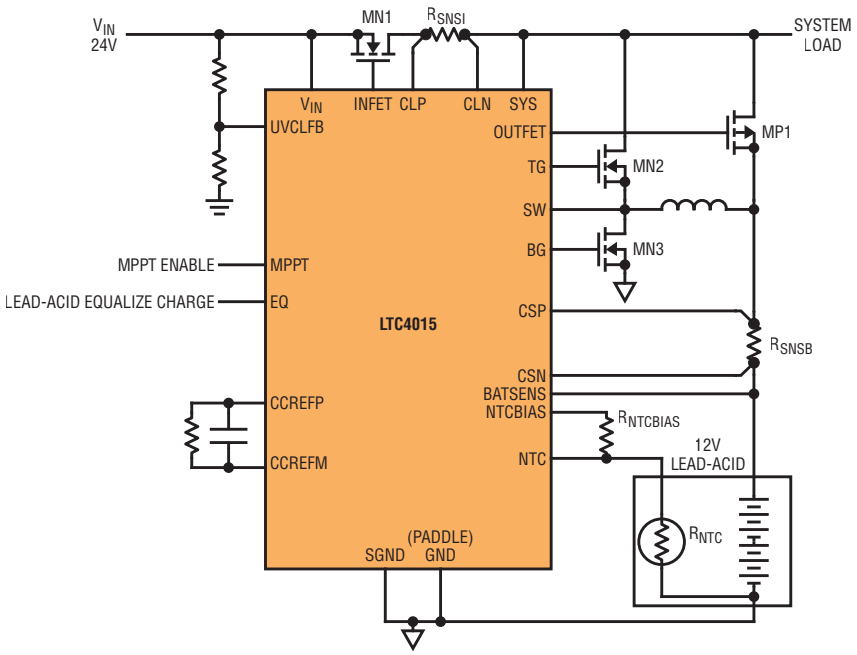

The LTC4015 maximum power point tracking algorithm performs a periodic global search as well as a continuous local dither to ensure that the solar panel powering the system remains at its peak power operating condition. The global search is necessary to ensure that the continuous dithering algorithm does not get stuck at a local maximum power point instead of the global maximum. Depending on the exact panel construction, this can occur during partial shading conditions.

This is important for MPPT operation. The LTC4015 effectively uses its ability to manipulate charge current to regulate both input current and input voltage. In other words, if the input voltage decreases enough such that the UVCLFB pin voltage falls below its DAC-programmed servo voltage, then charge current is reduced in an attempt to maintain that input voltage level. Likewise, if the input current starts to exceed the DAC-programmed input current limit, then charge current is reduced in an attempt to maintain that input current level. However, if charge current is reduced to zero, the LTC4015 loses its ability to further affect input current or input voltage. Consider LTC4015 MPPT operation in more detail to understand why these issues matter.

MPPT OPERATION

The LTC4015 maximum power point tracking algorithm performs a periodic global search as well as a continuous local dither to ensure that the solar panel powering the system remains at its peak power operating condition. The global search is necessary to ensure that the continuous dithering algorithm does not get stuck at a local maximum power point. Depending on the exact panel construction, this can occur during partial shading conditions.

Both the local dither and global search make use of the LTC4015 input voltage regulation function called UVCL, or undervoltage current limit. The UVCL control loop prevents a resistive or current limited input power source from falling too low (e.g., below the undervoltage lockout, UVLO threshold) by automatically reducing charge current as V_{IN} (observed at the UVCLFB pin using a V_{IN} voltage divider) drops to a programmable level (VIN_UVCL_SETTING).

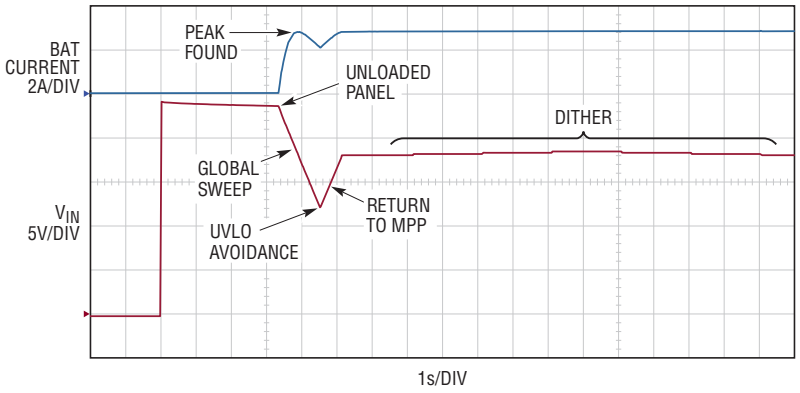

Figure 2. MPPT search algorithm

The global search steps VIN_UVCL_SETTING through its full range of values, being careful to avoid pulling V_{IN} below UVLO or V_{IN_DUVLO}, the differential undervoltage lockout. The differential UVLO condition is met if the input voltage falls to within about 100mV of the battery voltage. At each VIN_UVCL_SETTING, the charge current is measured. When the sweep is complete, the LTC4015 applies the VIN_UVCL_SETTING value corresponding to the maximum measured battery charge current.

Because the battery voltage is low impedance and relatively stable throughout the sweep, maximum battery charge current corresponds well to maximum output power. Following the global search, small changes in maximum power are tracked by slowly—approximately once per second—dithering the VIN_UVCL_SETTING. The LTC4015 periodically—approximately once per fifteen minutes—performs a new global search of VIN_UVCL_SETTING values, applies the new maximum power point, and resumes dithering at that point. Figure 2 shows a typical MPPT global search followed by local dithering.

The dithering algorithm begins by incrementing the VIN_UVCL_SETTING one step and measuring the new charge current. If the new charge current is greater than the previous measurement, then VIN_UVCL_SETTING continues to be incremented approximately once per second until the charge current decreases or VIN_UVCL_SETTING reaches full-scale, at which point the dither direction is reversed. Full-scale corresponds to V_{UVCLFB} = 1.2V and an input voltage of 36.5V with the required UVCLFB MPPT resistor divider values. In the reverse direction, VIN_UVCL_SETTING is decremented approximately once per second until either the charge current decreases or the input voltage falls too close to the UVLO thresholds, at which point the dither direction is reversed again.

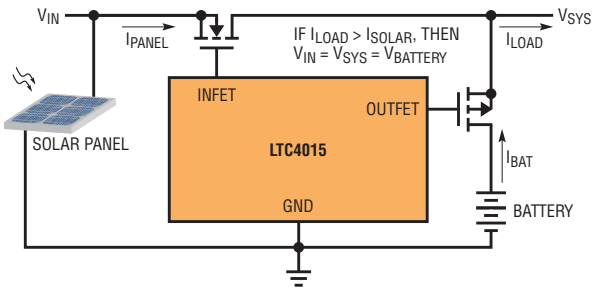

Figure 3. LTC4015 PowerPath architecture

MPPT SPECIAL CONSIDERATIONS

While MPPT operation is fairly straightforward under most conditions, there are a few cases outside the norm. The LTC4015 steps outside the basic algorithm in these cases in an effort to maximize the time the panel spends at its true maximum power point.

Significant Changes in Charge Current During Dither

When the LTC4015 is using the dithering algorithm, if the battery charge current falls by 1% or more in a single dithering step, the dither direction reverses after only 7ms, rather than the normal one second. This maximizes the time spent at the highest power setting. Similarly, if the step-to-step change in charge current is more than ±25%, the algorithm repeats a global search without waiting the standard 15 minutes. The maximum global search repetition rate is once per five minutes.

Input Current Limit Setting

As mentioned above, the LTC4015 monitors the input voltage during the MPPT algorithm to ensure that it does not fall below one of the UVLO thresholds. Another criterion under constant monitoring is whether or not the LTC4015 is actually in UVCL regulation using the vin_uvcl_active bit of the digital telemetry system. Remember that four parameters can be regulated: input voltage (VIN_UVCL_SETTING), input current (IIN_LIMIT_SETTING), charge voltage (VCHARGE_SETTING) and charge current (ICHARGE_TARGET). For MPPT applications, it is recommend that the input current limit (IIN_LIMIT_SETTING) is set greater than or equal to the maximum short-circuit current capability of the solar panel. This ensures that input current regulation does not interfere with MPPT operation. However, two other regulation loops can take control: charge voltage and charge current.

Available Current is Enough

During either the global search or dithering phase, if charge voltage or charge current regulation require less current than undervoltage current limit, UVCL, then it means that the solar panel can satisfy normal charging conditions at that particular VIN_UVCL_SETTING. At this point, the dither direction is reversed or the global search is stopped. During a global search, the VIN_UVCL_SETTING—which resulted in an exit from the UVCL regulation loop—likely corresponds to maximum charge current. If for some unusual reason it does not, the LTC4015 will ramp back to the VIN_UVCL_SETTING corresponding to maximum charge current.

Low Available Power

Another special case occurs when the maximum charge current, as measured by the completed global search, is below approximately 5% of full-scale, where full-scale corresponds to 32mV across R_{SNSB} (e.g., 200mA for a 4A charger). In this case, the LTC4015 returns to the VIN_UVCL_SETTING found during the global search, but does not attempt to dither. At this charge current level, noise in individual ADC readings becomes significant and dithering potentially leads to erratic operation.

Even Lower Available Power

If the maximum charge current, as measured by the completed global search, is even lower, less than approximately 1% of full-scale (e.g., 40mA for a 4A charger, or a mere 320µV across R_{SNSB}) then the LTC4015 has nearly lost its ability to control solar panel power. Nevertheless, one last attempt is made to maximize panel output power. The LTC4015 returns to a VIN_UVCL_SETTING that corresponds to 70% of the solar panel open-circuit voltage as measured when VIN_UVCL_SETTING was at full-scale. Because solar panels typically produce maximum power with a voltage of 70%–80% of their open-circuit voltage, with power rolling off gently at lower panel voltages, this is a best attempt at maximizing power with minimal available information.

POTENTIAL ISSUES WITH DIODE-OR TOPOLOGY

Depending on the specific application conditions, it is possible for the diode-OR topology (see Figure 1) to result in sub-optimal utilization of the solar panel power. Consider the simplified LTC4015 PowerPath architecture shown in Figure 3.

If the system load increases beyond the current capability of the solar panel, then both of the ideal diode controllers will turn on and MN1 and MP1 conduct to support the increased load. The solar panel output voltage collapses to the system load voltage, which collapses to the battery voltage.

The advantage of a battery fed topology is that the LTC4015 maximizes the combination of battery current and system load current. Nevertheless, a battery fed topology does come with trade-offs.

Operation with the solar panel voltage equal to the battery voltage is unlikely to result in maximum power, but in most applications, this should not be a serious concern.

The solar panel should be sized such that on average, its power capacity is greater than the average load power. If this condition is not met, then the battery will not charge. Therefore, the scenario depicted in Figure 3 should not be typical.

Furthermore, any solar panel paired with the LTC4015 must have an open-circuit voltage of less than 40V to avoid violating the LTC4015's absolute maximum ratings. Many commercially available panels that meet this requirement have a maximum power voltage of about 17V. When charging a 12V lead-acid battery, a 3S Li-ion stack (~11.7V), or a 4S Li-ion or $LiFePO_4$ stack (~15.6V and 14V, respectively), the panel will likely still be operating above 75% or 80% of its maximum power. In other words, even if the difference between the maximum power voltage of the panel and the battery voltage is relatively small, performance is not significantly impacted. For panels whose maximum power voltage is not 17V, the same logic applies. If the maximum power voltage is relatively close to the typical battery voltage, then the brief time periods when the system load exceeds the panel current will not significantly impact performance. However, if this scenario is still a concern, there is a solution.

BATTERY FED TOPOLOGY

In order to ensure that the LTC4015 can always maintain complete control of solar panel power it is necessary to move the connection for the system load. See Figure 4 for a simplified schematic of this topology, which can be referred to as a battery fed topology. This configuration forces the load to share programmed charge current with the battery. In other words, the system load current directly subtracts from the programmed charge current and reduces battery current. If the system load exceeds the programmed charge current, then the battery simply supplies the additional current required.

The advantage of this topology is that the LTC4015 maximizes the combination of battery current and system load current. In other words, the LTC4015 maximizes the total output power. Because the input PowerPath only feeds current to the switching regulator, the LTC4015 has complete control over the input current. Since the LTC4015 output charges the battery and powers the load in this configuration, it can reduce output power to zero. The load remains supported by the battery under this condition.

Nevertheless, a battery fed topology does come with trade-offs, namely:

- The coulomb counter functionality of the LTC4015 is critically impaired because the LTC4015 cannot differentiate battery current from system load current. This inability to distinguish between the two currents has other consequences. The programmed charge current is no longer a fixed battery charge current. Instead, the battery charge current varies with system load. While charging, the digital telemetry

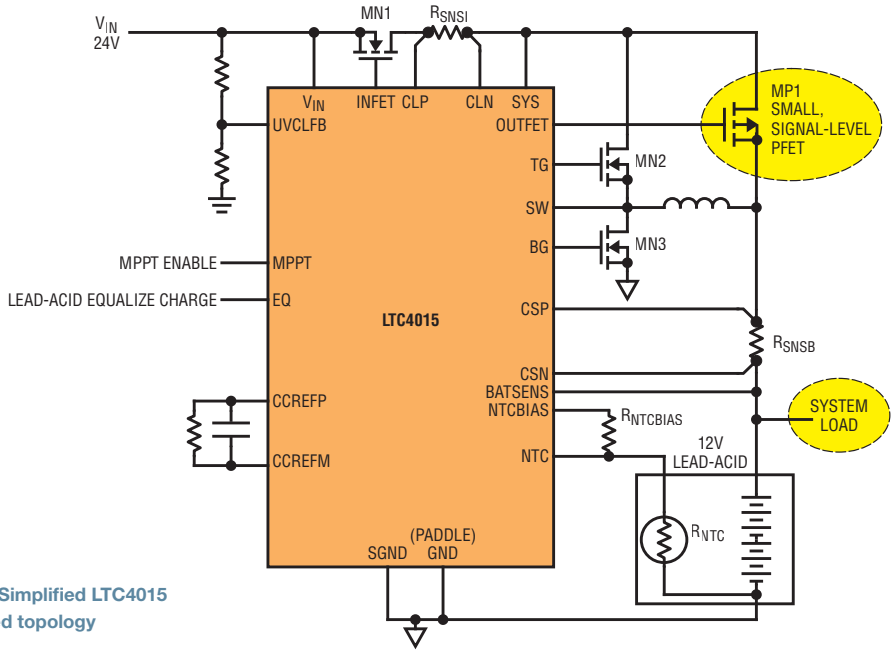

Figure 4. Simplified LTC4015 battery fed topology

The LTC4015 can serve as the power management backbone of battery powered applications, and is especially adept at charging batteries and supporting loads when a solar panel is the input power supply.

system would be able to monitor and report on the sum of system load current and battery current, but no current readings would be available in "battery-only" operation (no input supply).

- This battery fed topology also impacts the termination algorithms, especially the current-based C/x termination. Instead of terminating when the battery current falls below a programmed threshold, the LTC4015 charge algorithm terminates when the combination of the load current and the battery current falls below that threshold. If the charge cycle were to terminate, then all system load current would be drawn from the battery until a recharge cycle begins.
- Finally, the ideal diode-OR PowerPath topology (Figure 1) provides power to the system as soon as input voltage is available even if the battery is heavily discharged. In the battery-fed topology of Figure 4, the input supply will need to charge the battery to a voltage greater than the minimum system voltage before the system can operate.

A corollary to this final disadvantage is that the battery must be able to supply the full load current at all times. Because the MPPT algorithm and the battery series resistance (BSR) algorithm will temporarily and periodically disable the switching regulator, the battery must be able to supply the full system load during these times. This is particularly critical when charging Li-ion chemistries. The LTC4015 Li-ion charge algorithms include a pre-charge phase. If a system load can discharge the battery below the pre-charge threshold, and that same load exceeds the pre-charge current, it is possible that the battery will be drained even with an input supply present. This could permanently damage the battery.

Because of these disadvantages, careful consideration should go into deciding between the standard diode-OR topology and the battery fed topology.

MPPT AND LOW INPUT POWER

Despite its carefully designed MPPT algorithm (including the special corner cases described above) and ability to operate in different topologies, there is one scenario where the LTC4015 cannot maximize solar panel output regardless of topology.

The LTC4015 battery charger function requires a minimum amount of current to operate, which varies depending on the application (switching MOSFET selection, compensation, etc.) If the maximum input current available from the solar panel is above 2mA to 3mA but below the minimum level required to operate the charger (approximately in the range of 5mA to 20mA) then the battery may actually be discharged slightly by the charger.

Under these conditions—for example, a very dimly lit, but not completely dark, solar panel—the worst-case battery drain current is generally less than 10mA. The condition persists as long as the available input current remains in the range described. If the available input current falls lower, then the battery discharge returns to near normal battery-only mode levels—see the data sheet for details.

For typical solar panel applications, this condition is generally short-lived and infrequent, requiring no mitigation. For example, a short period of time before sunrise and after sunset may result in some extra battery drain. Nevertheless, as described in the data sheet, if this condition is a concern, it can be mitigated by disabling the charger (setting suspend_charger = 1) whenever the battery charge current falls below 1% of full-scale (IBAT ≤ 218) and retrying by writing suspend_charger = 0 periodically (e.g., every 60 seconds). Optionally, this retry can be limited to only occur when V_{IN} is above a known threshold.

CONCLUSION

The LTC4015 can serve as the power management backbone of battery powered applications, and is especially adept at charging batteries and supporting loads when a solar panel is the input power supply. Its integrated ideal diode-OR controllers and ability to measure and regulate input current, battery current, input voltage and output voltage, enable it to maintain high battery charging performance and maximum power point tracking for solar panel input supplies. ■

MOSFET Safe Operating Area and Hot Swap Circuits

Dan Eddleman

"Is this MOSFET's SOA (safe operating area) adequate for my application?" This is the most frequently asked question by designers of hot swap circuits. When evaluating a MOSFET's suitability for a specific application, it helps to have an intuitive understanding of SOA. After a short review of the SOA specifications found in MOSFET data sheets, this article presents a simple explanation of SOA, with a focus on MOSFET thermal behavior. With this understanding, designers can confidently use SOAtherm, a tool included with LTspice, to accurately evaluate MOSFET SOA in their circuit simulations.

WHAT IS SOA?

Every MOSFET data sheet includes an SOA plot, which describes the maximum time a MOSFET can be exposed to a specific voltage and current. Figure 1 shows the SOA plot from NXP Semiconductor's data sheet for the PSMN1R5-30BLE 30V 1.5mΩ N-channel MOSFET. Consider the condition where 10V at 100A is applied to the MOSFET. Looking at the corresponding point on the SOA plot, we see that it falls between the 1ms and 10ms lines. The plot indicates that you can apply 10V and 100A for at least 1ms without damaging the MOSFET if the case (tab) is held at 25°C.

Of course, holding the case at 25°C is impossible, requiring an unattainably perfect heat sink, but fortunately, for short duration events (less than 10ms–100ms), the case temperature does not rise significantly. Hot swap circuits only see significant drain-to-source voltages during short duration events—start-up, input supply steps and output overcurrent conditions—so this 10ms–100ms time limit is typically satisfied.

Derating for higher case temperatures (above 25°C) is discussed in the "Transient Thermal Impedance" section below. For events longer than 10ms, see the "Beyond 10ms" sidebar.

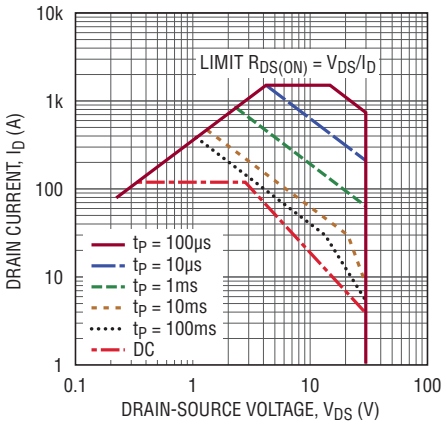

Figure 1. Safe operating area; maximum allowable time for a pulse at a fixed voltage and current

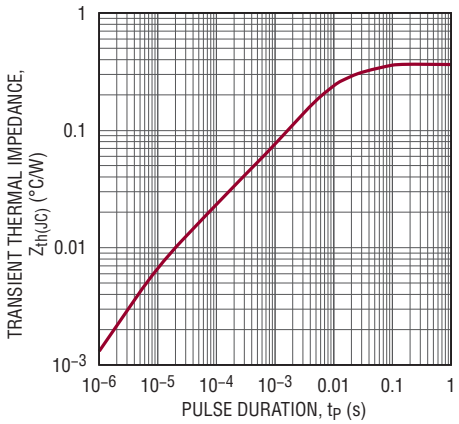

Figure 2. Transient thermal impedance from junction to mounting base as a function of pulse duration

STEADY STATE

Before delving deeper into transient SOA events, it is helpful to step back and review the more familiar steady state (DC) limits. MOSFET data sheets specify the maximum silicon temperature (typically 150°C or 175°C) as well as θ_{JC} and θ_{JA}, the thermal resistance from silicon (junction) to the bottom of the package (case), and the thermal resistance from silicon (junction) to the environment (ambient), respectively. (Sometimes $R_{th(JC)}$ and $R_{th(JA)}$ are used as alternative names for θ_{JC} and θ_{JA}.)

From the PSMN1R5-30BLE data sheet, θ_{JA} = 50°C/W and θ_{JC} = 0.3°C/W. θ_{JA} is used to calculate the temperature rise from ambient to the MOSFET's silicon die with a PC board configuration specified in the data sheet. In the PSMN1R5-30BLE data sheet, θ_{JA} is specified with a minimum footprint on FR4 PC board. Assuming your PC board is identical to the one the MOSFET manufacturer used to specify θ_{JA}, the silicon die temperature is:

$$T_{JUNCTION} = T_{AMBIENT} + \theta_{JA}\,[°C/W] \bullet Power[W]$$

For example, with a θ_{JA} of 50°C/W and an ambient temperature of 75°C, the die temperature will be 125°C when 1W is dissipated by the MOSFET.

For most modern MOSFETs with exposed metal tabs, θ_{JA} is primarily determined by the PC board layout rather than the MOSFET itself (although the exposed pad shape and size play a role). Because θ_{JA} is highly dependent on the PC board layout and airflow, the manufacturer's specified θ_{JA} is only suitable for rough estimates.

θ_{JC} is often a more useful metric, as it describes the MOSFET behavior without the influence of the PC board layout. To determine silicon temperature, use the following:

$$T_{JUNCTION} = T_{CASE} + \theta_{JC}[°C/W] \bullet Power[W]$$

With 1W of power dissipation, the silicon temperature is only 0.3°C above the case temperature. When using this formula, the case temperature (T_{CASE}) must be determined by physical measurement or through a thermal simulation of the PC board. Obviously, the PC board layout, airflow, and heat sinks are critical factors when calculating the steady state conditions.

TRANSIENT THERMAL IMPEDANCE

Most MOSFET data sheets also include a transient thermal impedance plot. The "single pulse" transient thermal impedance ($Z_{th(JC)}$) is the temperature rise produced by a time-limited power pulse. The longest time point on the transient thermal impedance plot always matches the θ_{JC} specification, because θ_{JC} is, by definition, the steady state (infinite time) thermal impedance. Figure 2 shows the transient thermal impedance from the PSMN1R5-30BLE data sheet. For the purposes here, only the "single pulse" curve is important.

The transient thermal impedance plot may be used to calculate the temperature rise for a power pulse of any duration. For example, assume a MOSFET drain-to-source voltage (V_{DS}) of 12V and a drain current (I_D) of 100A. The power dissipated by the MOSFET is 12V • 100A = 1.2kW. If we look at the transient thermal impedance plot at 1ms, the thermal impedance is 0.075°C/W. The silicon junction temperature is:

$$T_{JUNCTION} = T_{CASE} + Z_{th(JC)}\ [°C/W] \bullet Power[W]$$
$$= 0.075°C/W \bullet 1.2kW + 25°C$$
$$= 115°C$$

for a 1ms, 1.2kW pulse with a fixed case temperature of 25°C.

For moderate V_{DS} voltages (below the Spirito region, see sidebar), MOSFET manufacturers generate the SOA plot from the transient thermal impedance plot. In other words, these two plots are alternate expressions of the same information. The SOA plot shows the time it takes for the silicon die to reach its maximum junction temperature (150°C or 175°C) for each V_{DS} voltage and I_D current combination. Be aware that the SOA plot is only valid for a case temperature of 25°C and must be derated for higher case temperatures, including the case

Spirito Effect

Years ago when the maximum current in a hot swap design was less than 10A, it was easy to find a MOSFET to satisfy most applications. Two things have changed in the last decade. First, supply currents have increased significantly with 100A or more becoming common. Second, MOSFET manufacturers have been hard at work improving the resistance specifications of MOSFETs ($R_{DS(ON)}$) when they are fully turned on. Ironically, this has reduced the available SOA at higher drain-to-source voltages in what has been termed the "Spirito effect." Professor Paolo Spirito[1] explained that as MOSFET manufacturers have increased transconductance to improve the on-resistance, there is a greater tendency for the MOSFETs to fail by forming unstable hot spots.

Two primary factors compete to determine if hot spots cause MOSFET failure. One factor is the MOSFET's ability to dissipate power without a rapid increase in temperature. (This is reflected in the transient thermal impedance curve.) The second competing factor is the tendency of MOSFET cells to "run away" by stealing more current from neighboring cells as they get hotter. This second factor is dominated by the temperature coefficient of the MOSFET's threshold voltage, which drops with increasing temperature, causing "current crowding" in hotter cells. (MOSFET transconductance falls with increasing temperature due to reduced carrier mobility in the MOSFET's conduction channel. It somewhat counteracts the current crowding effects, but can be safely ignored in this explanation.)

Inside a MOSFET package, there is a silicon die that contains an array of MOSFET cells with their gates, drains and sources connected in parallel. As some cells become hotter than others, their threshold voltages decrease relative to the cooler cells, causing the hotter cells to conduct more current. If the competing factors cited above reach an unstable condition, certain cells may thermally run away, drawing more and more current until they self-destruct.

The Spirito effect is observed primarily at high V_{DS} voltages since a given change in cell current results in a greater change in power at high V_{DS}, resulting in an increased tendency of cells to thermally run away. Similarly, the Spirito effect is most pronounced at lower currents where there is more time for the MOSFET cells to thermally run away. (At higher currents, the average die temperature reaches the 150°C or 175°C before any cells exhibit significant thermal runaway.) For this reason, the high V_{DS} and low I_D region of the SOA plot, where the Spirito effect is dominant, is sometimes referred to as the "Spirito region" and is highlighted in the PSMN1R5-30BLE SOA in Figure 3.

Spirito Region

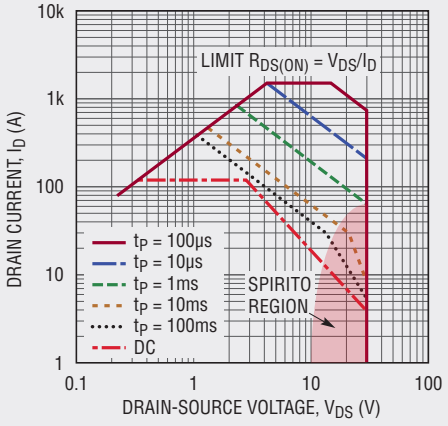

Notes

[1] G. Breglio, F. Frisina, A. Magri, and P. Spirito, "Electro-Thermal Instability in Low Voltage Power MOS: Experimental Characterization," IEEE Proceedings ISPSD 1999, Toronto, p233

temperature rise that occurs from the pulse itself. (See "Beyond 10ms" sidebar.)

Knowing that the maximum junction temperature of the PSMN1R5-30BLE is 175°C, and using a case temperature of 25°C, we can calculate the maximum allowable time at 1.2kW.

$T_{JUNCTION} - T_{CASE} = Z_{th(JC)}[°C/W] \bullet Power[W]$

$175°C - 25°C = Z_{th(JC)} \bullet 1.2kW$

$Z_{th(JC)} = 0.125°C/W$

Looking at the transient thermal impedance plot we find that $Z_{th(JC)}$ crosses 0.125°C/W at roughly 2ms, which also matches the SOA plot.

With an understanding of transient thermal impedance plots, we can calculate the allowable time for case temperatures other than 25°C. In the previous 1.2kW example, the allowable time was 2ms with a 25°C case temperature. Now, assume the case temperature is 85°C:

$T_{JUNCTION} - T_{CASE} = Z_{th(JC)}[°C/W] \bullet Power[W]$

$175°C - 85°C = Z_{th(JC)} \bullet 1.2kW$

$Z_{th(JC)} = 0.075°C/W$

Looking at the transient thermal impedance plot we find that $Z_{th(JC)}$ crosses 0.075°C/W at 1ms, significantly less than the 2ms we found for a 25°C case temperature.

Because thermal behavior is linear, we can use the transient thermal impedance plot to determine the temperature rise for any power shape. While it is possible to do this calculation using convolution, it is easier to model thermal behavior in an electrical circuit simulator such as SPICE. In particular, the SOAtherm tool in LTspice can be used to model the MOSFET thermal behavior.

SOATHERM THERMAL MODELING IN LTspice PREDICTS MAXIMUM MOSFET DIE TEMPERATURE

A designer armed only with MOSFET data sheet SOA plots faces a difficult challenge in predicting a MOSFET's suitability for a hot swap design. Fortunately, MOSFET thermal behavior (and SOA) can be modeled in circuit simulators such as LTspice.

The SOAtherm symbol included in LTspice includes a collection of MOSFET thermal models that simplify the task of predicting MOSFET maximum die temperature over time, even in the Spirito region. The thermal model reports the temperature of the hottest point on the MOSFET die without influencing the electrical behavior of the MOSFET model.

For better or worse, the SOAtherm models are based on the MOSFET manufacturers' data sheets, and as such are only as accurate as the manufacturers' data itself. With that in mind, design with plenty of margin since the SOA curves provided by MOSFET manufacturers are usually typical numbers without sufficient derating to account for part-to-part variation.

Using SOAtherm

To use SOAtherm, place the SOAtherm-NMOS symbol on top of a MOSFET in an LTspice simulation (Figure 3). The voltages at the Tc and Tj pins of the SOAtherm-NMOS symbol indicate the case temperature and silicon junction temperature respectively. (Refer to the SOAtherm-NMOS tutorial for more information about using this model, including how to adjust the ambient temperature settings and other parameters.)

After running the simulation, the silicon and case temperature can be observed in the waveform viewer (Figure 4). In the waveform shown here, the MOSFET silicon junction temperature rises from 25°C to

Beyond 10ms

For most MOSFETs, the case temperature does not rise significantly during transient events lasting less than 10ms, because it takes time for the heat to move through the MOSFET silicon and copper. At roughly 10ms, the heat begins to reach the PCB.

If the MOSFET's copper tab is small, the temperature of the MOSFET begins to rise faster as the heat reaches the PCB. For packages where the copper tab is larger (i.e., D2PAK packages) the heat begins to move outward into portions of the copper tab that are still cool. As a result, packages with more copper perform better in high SOA applications (hot swap designs, linear amplifiers, etc.) than MOSFETs with less copper, even if their transient thermal impedance and SOA plots appear similar.

Think of the copper as a reservoir that helps to limit the MOSFET temperature rise during events in the 10ms–10s timeframe. The transient thermal impedance plot and the SOA plot are often deceptive, because they are created by assuming the case temperature is fixed at 25°C by an impossibly perfect heat sink.

The figure here shows the simulated thermal characteristics of a Power-SO8 package and a D2PAK package soldered to a PCB with a 1oz copper plane on the top layer. The figure also includes the thermal impedance curves of the type found in MOSFET data sheets where the case temperature is fixed.

At 1ms, the heat is concentrated within the silicon die. The D2PAK silicon is cooler with a thermal impedance of 0.075°C/W compared to 0.14°C/W for the Power-SO8, primarily due to the larger silicon die in the D2PAK. At 10ms, the heat begins to reach the bottom of the copper tab, and the temperatures start to diverge. At 100ms, the Power-SO8 die has a temperature rise of 4.2°C/W, while the temperature rise of the D2PAK is only 0.6°C/W. Clearly, the extra copper in the D2PAK saves the day.

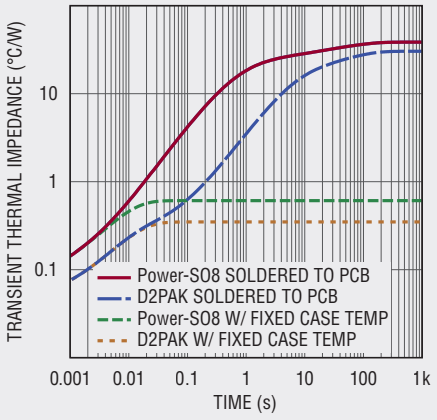

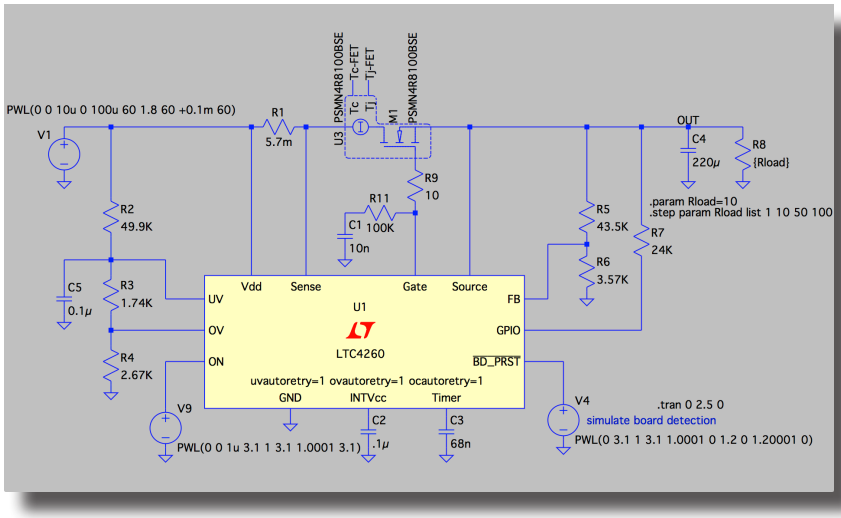

Figure 3. SOAtherm simplifies hot swap MOSFET selection by evaluating SOA in an LTspice circuit simulation.

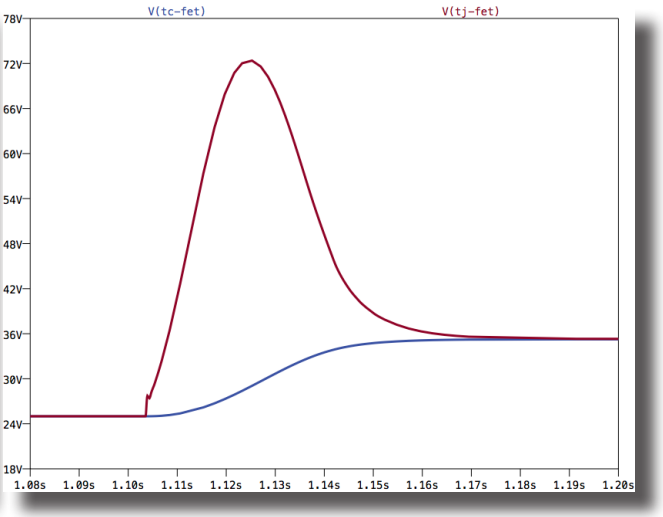

Figure 4. SOAtherm waveform. Voltage corresponds to °C

72°C. The case temperature rises from 25°C to 35°C. (A 1V rise on the Tc or Tj pin is equivalent to a 1°C temperature rise.)

Remember to Simulate Special Circumstances

There are several important special circumstances that should not be overlooked when using SOAtherm to determine if MOSFET SOA limits may be exceeded.

- Input supply step. For example, the SOA requirements of a –48V telecom application, where the input supply may quickly step from –36V to –72V can require a MOSFET with significant SOA capabilities. When supplies are pre-regulated, or are well controlled to eliminate such steps, the SOA requirements are reduced.
- Start-up into a load. The downstream circuitry may turn on and draw current before the supply is fully ramped-up, or a component such as a capacitor may fail in a resistive short. Simulating a resistive load at the output can indicate when a MOSFET might unexpectedly be subjected to a condition requiring significant SOA.
- A short circuit at the output occurring during otherwise normal operation. You never know when the user is going to drop a paperclip into the chassis, and a hot swap circuit (or fuse) prevents a call to the fire department.

EXAMPLE USING THE LTC4226 WIDE OPERATING RANGE DUAL HOT SWAP CONTROLLER

The LTC4226 is a dual hot swap controller that drives external N-channel MOSFETs in applications with supply voltages as high as 44V.

In the circuit in Figure 5, the LTC4226 provides current limit and circuit breaker features for a 12V supply and a 5V supply. The circuit breaker timer is configured with the capacitors connected to the FTMR1 and FTMR2 pins. When the voltage across either sense resistor is between 50mV and 86mV, the corresponding capacitor at FTMR1 or FTMR2 is ramped up with a 2µA current.

Because the current limit is not engaged until the sense resistor voltage reaches 86mV, the power dissipation in the MOSFET is negligible as long as the current remains below 86mV/5mΩ = 17.2A. When the current exceeds that level, current limit is engaged and the FTMR1 or FTMR2 pin ramps up with 20µA. The appropriate channel's MOSFET is turned off when the corresponding FTMR pin reaches 1.23V, setting a maximum time before the MOSFET is shut off. In this example, 100nF capacitors configure a 6.2ms current limit timeout for both channels.

With the LTC4226, the worst-case MOSFET power dissipation occurs when the output is shorted to ground. As a result, determining the required SOA is straightforward. (With hot swap controllers that feature current foldback or power limiting, more effort is required to determine the worst-case loading condition.) Referring to the SOA plot in Figure 1 for the PSMN1R5-30BLE, it can be seen that 6.2ms is well inside the SOA limit at 17.2A and 12V. An SOAtherm simulation confirms that the total junction temperature rise is less than 50°C. The same simulation shows a negligible case temperature rise of roughly 5°C, which would be expected from this rather large D2PAK package during a short 6ms event.

The 5V supply in this example application uses a powerPAK-SO8 package for the MOSFET, which is smaller than the D2PAK used for the 12V supply. A smaller package may be used for the 5V supply because the worst-case power dissipation of the 5V supply's MOSFET is 17.2A • 5V = 86W versus the 17.2A • 12V = 206W worst-case dissipation of the 12V supply's MOSFET. An SOAtherm simulation of this circuit

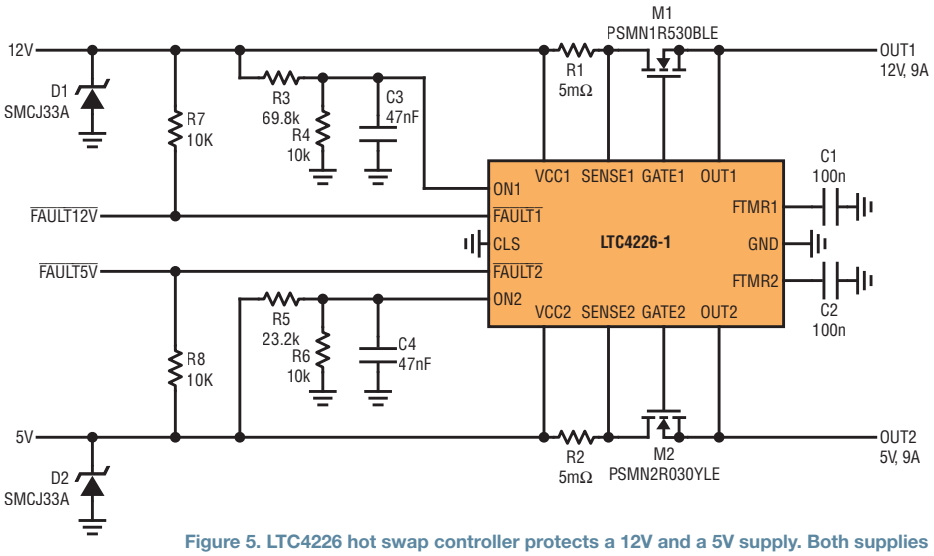

Figure 5. LTC4226 hot swap controller protects a 12V and a 5V supply. Both supplies provide 9A steady state current and up to 17.2A during transients

predicts a junction temperature rise of 40°C, including a case temperature rise of 30°C. The larger case temperature rise is explained by the smaller size of the powerPAK-SO8 package (and correspondingly less copper) relative to the D2PAK used for the 12V supply.

The above calculations and simulation help to verify a circuit design and MOSFET selection, but the ultimate test must be done in the lab with an assembled circuit. Because the worst-case SOA requirement of the LTC4226 occurs with an output short, the lab test is as simple as quickly applying the input power supply with a grounded output. A good technique is to hot-plug the LTC4226 circuit into a live power supply to simulate an actual hot swap event. Alternatively, an output short can be applied while the input supply is fully powered. To determine if the circuit has extra margin, swap the timer capacitor to a larger value and test again.

The circuit in Figure 6 shows a technique for using two parallel MOSFETs when a single MOSFET may not satisfy the SOA requirements of an application. In general, it is not advisable to use parallel MOSFETs to increase the SOA capability of a circuit. Mismatches between the MOSFETs, especially mismatches in the threshold voltages, may result in one MOSFET thermally running away and conducting all of the current. Nevertheless, the circuit in Figure 6 safely uses parallel MOSFETs by implementing independent current limit in each channel, preventing either MOSFET from running away.

Additionally, the cross-coupled PNPs, Q1 and Q2, only allow the circuit breaker timers to activate when both MOSFETs are conducting their full current. Without the cross-coupled PNPs, one channel's circuit breaker timer could activate if it was delivering a greater share of the load current.

CONCLUSION

As the power levels required in hot swap applications have increased, so have concerns regarding MOSFET safe operating area. Frequently, the most challenging aspect of designing a high power hot swap circuit is determining whether a specific MOSFET is capable of supporting the application. At a minimum, the circuit designer must be comfortable interpreting MOSFET SOA plots. As power levels increase and approach the limits of existing MOSFET technology, an understanding of transient thermal impedance plots and the ability to simulate this behavior in SPICE circuit simulations are invaluable tools in the hot swap circuit designer's arsenal. ■

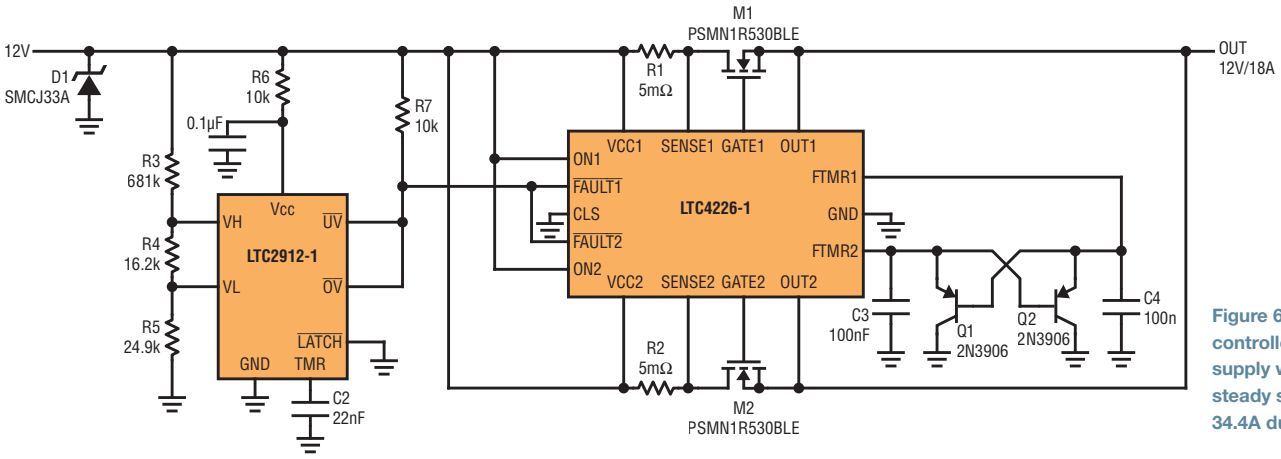

Figure 6. LTC4226 hot swap controller protects a 12V supply while providing 18A steady state current and up to 34.4A during transients

LT8315 12V HIGH INPUT VOLTAGE ISOLATED FLYBACK CONVERTER

The LT8315 is a high voltage flyback converter with integrated 630V/300mA switch. No opto-isolator is needed for regulation. The device samples the output voltage from the isolated flyback waveform appearing across a third winding on the transformer. Quasi-resonant boundary mode operation improves load regulation, reduces transformer size, and maintains high efficiency.

www.linear.com/solutions/7750

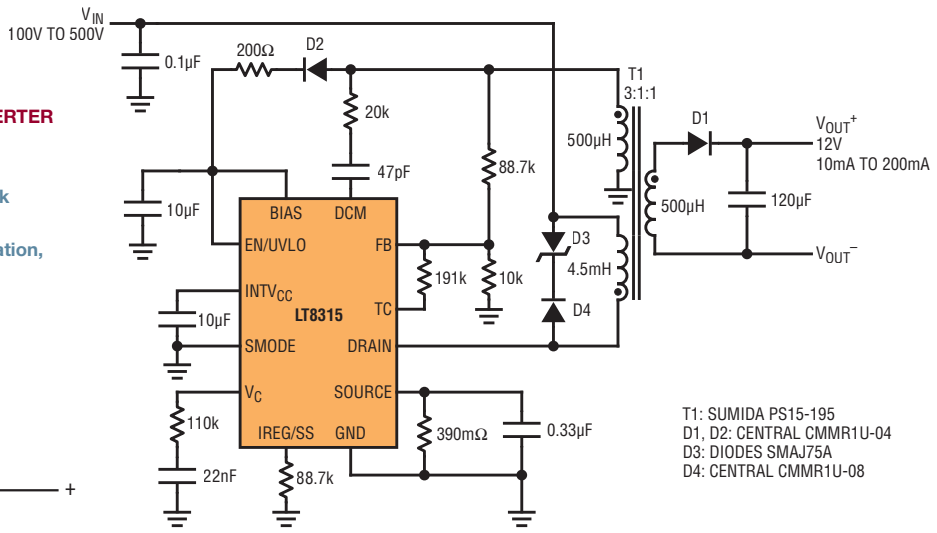

LT4320 IDEAL DIODE BRIDGE

The LT4320/LT4320-1 are ideal diode bridge controllers that drive four N-channel MOSFETs, supporting voltage rectification from DC to 600Hz typical. By maximizing available voltage and reducing power dissipation, the ideal diode bridge simplifies power supply design and reduces power supply cost, especially in low voltage applications.

www.linear.com/solutions/7766

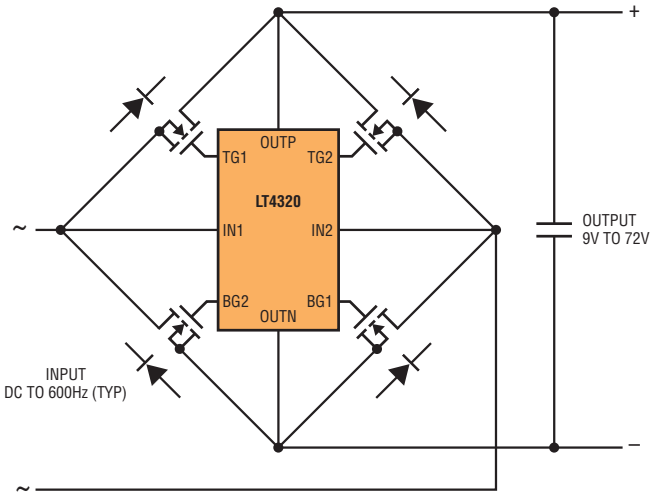

LTC3119 PHOTOVOLTAIC PANEL INPUT LEAD-ACID CHARGER WITH TEMPERATURE CORRECTION (1MHz)

The LTC3119 is a high efficiency 18V monolithic buck-boost converter that can deliver up to 5A of continuous output current. Extensive feature integration and very low resistance internal power switches minimize the total solution footprint for even the most demanding applications. A proprietary 4-switch PWM architecture provides seamless low noise operation from input voltages above, equal to or below the output voltage.

www.linear.com/solutions/7741

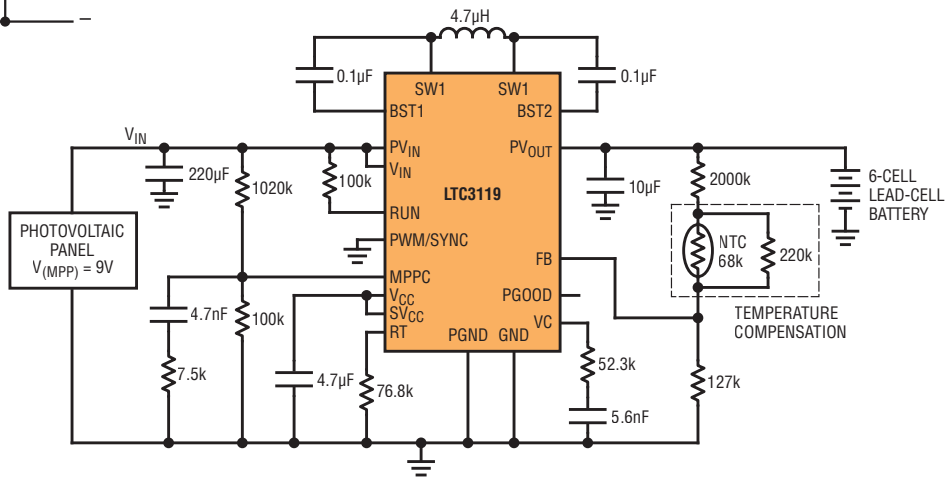

Manage Your LT Journal Subscription

Going Green? Switch from the printed magazine to our online version. When the next issue of LT Journal is published, you'll receive an email digest with links to the full online version. Please visit www.linear.com/mylinear and login to manage your delivery options.

Unsubscribe? We're sorry to see you go. Please visit www.linear.com/mylinear and login to manage your subscription.

If you do not have a MyLinear account and would like to unsubscribe, please visit www.linear.com/unsubscribe, enter the exact name and mailing address on the mailing label and click "Unsubscribe."