



feature

ANALOG & POWER

IBA-Based Power Solutions Demand System-Level Protection

Bob Thomas
 bobtho@cisco.com
 Cisco Systems Inc.

In many computing and communication applications, the intermediate bus architecture (IBA) with non-isolated point-of-load (niPOL) converters continues to displace conventional distributed and centralized power solutions. Key drivers of this trend include the increasing number of system voltages, their higher output currents, tighter regulation requirements, and lower total system cost. While meeting these needs, many IBA/niPOL solutions don't include certain protection mechanisms that come standard in previous conventional approaches.

For example, output overvoltage protection (OVP) and overtemperature protection (OTP) are very desirable in power-system designs. Without them, a failure in the system could lead to smoke and an acrid smell, loud noises as components fail, and visible damage to the product. Many niPOL converters, either as purchased assemblies or discrete designs, don't include comprehensive OVP and OTP circuitry. It's possible to add this protection to each niPOL converter, but a more elegant approach would be to solve it at the system level.

SYSTEM OVERVIEW • To understand power protection at the system level, let's review a typical power-system block diagram for an IBA/POL solution (Fig. 1). Here, the 12-V intermediate bus is gener-

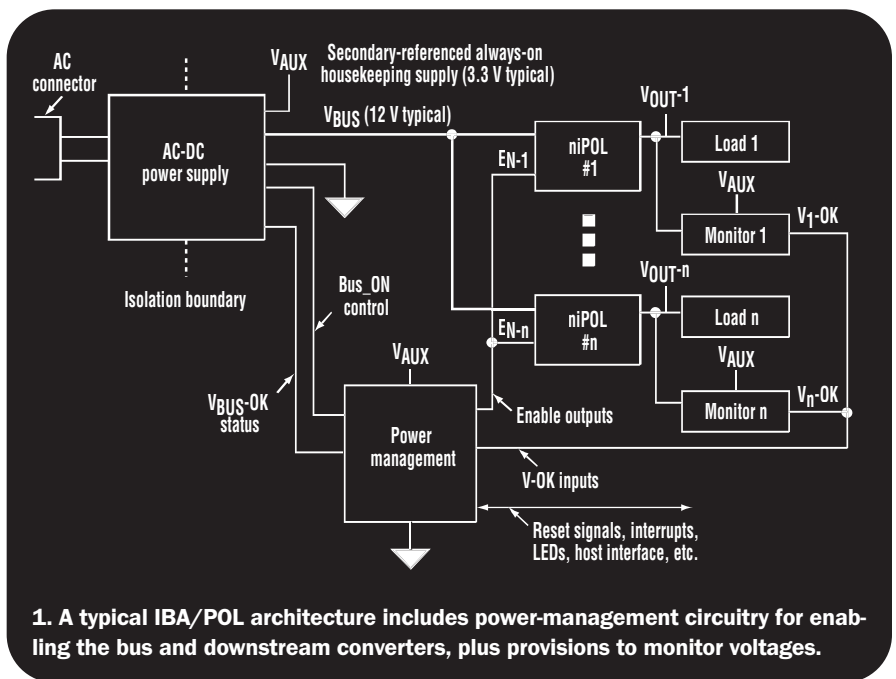
ated from an ac-dc power supply. In addition to the main output, there's a low-power always-on auxiliary supply for the power-management circuitry. This circuitry controls the intermediate bus and the individual niPOL converters with enable signals that set the on/off state, and V-OK signals to monitor the voltages. Other functions in the power-management circuitry may include voltage sequencing, reset signals, status LEDs, and other interfaces to the host system.

Because the intermediate bus provides power to all downstream converters (as well as fans, disk drives, or other loads), it typically is a high-power output.

Depending on the system, this bus can deliver between 200 W and 1 kW. If a fault condition allowed this power to be concentrated in a single device, such as a damaged semiconductor, the localized heating would likely lead to burning and other undesirable effects.

The key feature to exploit is the power-management circuitry's ability to shut down the intermediate bus. If some fault causes excessive temperature or voltage somewhere in the system, that fault's ability to cause damage goes away upon the power source's removal. The only reasonable response to a critical overtemperature or overvoltage event is to shut down the power source. Therefore, the design solution provides signals to the power-management circuitry, which enable it to latch-off the intermediate bus in response to one of these fault conditions.

OVERTEMPERATURE PROTECTION • A centralized power system, such as a multi-output ac-dc power supply, has



1. A typical IBA/POL architecture includes power-management circuitry for enabling the bus and downstream converters, plus provisions to monitor voltages.

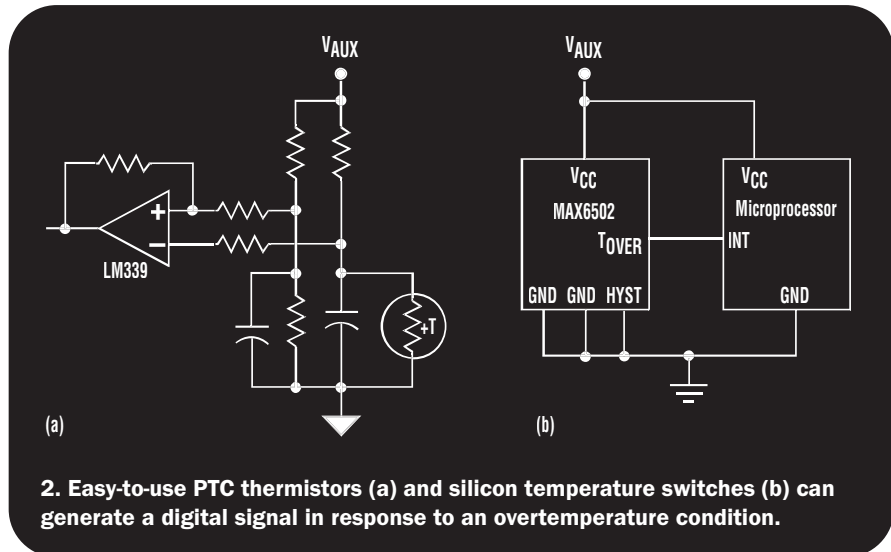
internal temperature sensors and fan-speed monitors to protect against internal or external thermal fault conditions. A traditional isolated dc-dc brick converter also supplies OTP by using one or more thermal sensors placed at critical locations on its assembly. These products are thoroughly tested to ensure that the output shuts down before the converter suffers permanent damage.

By contrast, most niPOL converters don't include OTP as a standard feature. Some controller ICs will shut down if their die temperature is too high. However, this may not detect and protect against excessive temperature in power components like the MOSFETs. This is especially true for discrete niPOL designs, where the power train may be thermally isolated from the controller.

To provide OTP at the system level, let's consider the causes for power-system overtemperature events. Assuming the converter is designed to protect itself from an overcurrent condition, the two key remaining conditions are excessive ambient temperature and insufficient airflow. Most system designs include ambient temperature sensors and fan-speed monitors. The power-management circuitry can monitor these signals. If temperature and fan-speed thresholds are exceeded, the power-management controller latches off the bus converter and asserts a fault condition.

Also, the system may have components like processors, disk drives, and ASICs that are more sensitive to thermal faults than the power converters. The system design already may have thermal monitors for these devices that can supply additional thermal-fault information to the power-management circuitry. Verification typically happens as part of the thorough system qualification testing.

If more thermal monitoring of the power system is required, it can be performed with small, low-cost solutions. A positive-temperature-coefficient (PTC) thermistor can be placed at a critical location on the board. When connected to a comparator in the bridge configuration as shown, the circuit can be powered from the auxiliary supply with no reference voltage required (Fig. 2a). Another simple solution is a silicon temperature sensor that supplies a digital signal when the temperature threshold is exceeded (Fig. 2b). These signals can be fed to the power-management circuitry to indicate a thermal



2. Easy-to-use PTC thermistors (a) and silicon temperature switches (b) can generate a digital signal in response to an overtemperature condition.

fault condition. The temperature thresholds should be chosen carefully to protect the components while avoiding any “nuisance tripping” during supported operating conditions.

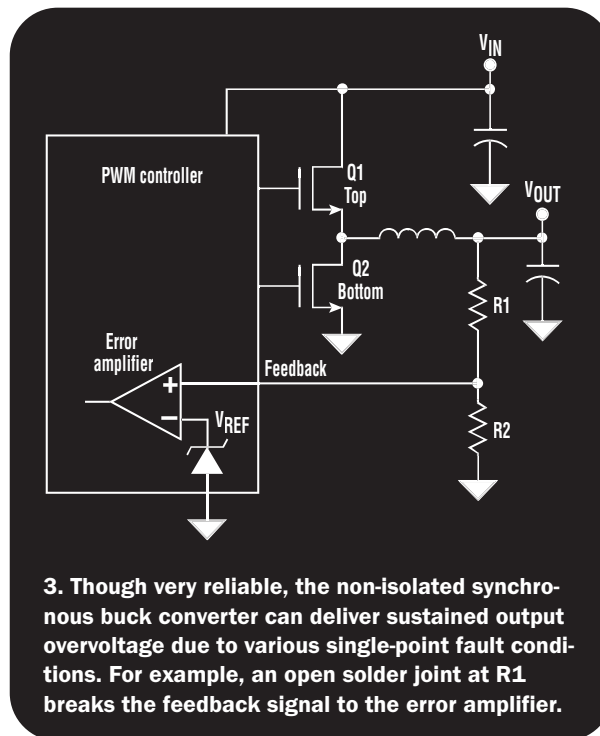
The most effective solution may combine the existing system thermal sensors and additional thermal sensors placed at strategic locations within the power system. Thermal threshold detectors are easy to use, and their outputs can be diode-ORed to minimize the complexity of the power-management circuitry.

Careful attention must be applied at the design and qualification stage to avoid false shutdown conditions, while ensuring that thermal fault conditions don't lead to excessive product

damage. As part of the power-system design, identify the components and locations where high power dissipation is expected. While switching MOSFETs likely will make up most of the power loss, FET drivers, controllers, and other devices also should be considered.

As each location is identified, determine which thermal sensor in the system would be able to monitor it. Because thermal faults can be caused by a lack of airflow, the thermal-sensing path should not require forced-air. If necessary, consider adding more thermal sensors at locations that can't be effectively monitored by existing devices. If they are later determined unnecessary, they can be “no-loaded” in production.

To validate the effectiveness of the thermal monitors and ensure that they never cause false shutdowns, system validation testing is



3. Though very reliable, the non-isolated synchronous buck converter can deliver sustained output overvoltage due to various single-point fault conditions. For example, an open solder joint at R1 breaks the feedback signal to the error amplifier.

performed. This is typically done by placing thermocouples on the sensors and on the power-dissipating components in the system. Running the system under “worst-case” conditions helps monitor these temperatures. The product is operated at the highest supported ambient temperature and in a condition that consumes the maximum power. Assuming that the component temperatures don’t exceed their derated maximums under this condition, all thermal sensor temperatures should fall below their minimum trip values. It’s important to add design margin to ensure that the thermal sensors never trip under supported operating conditions.

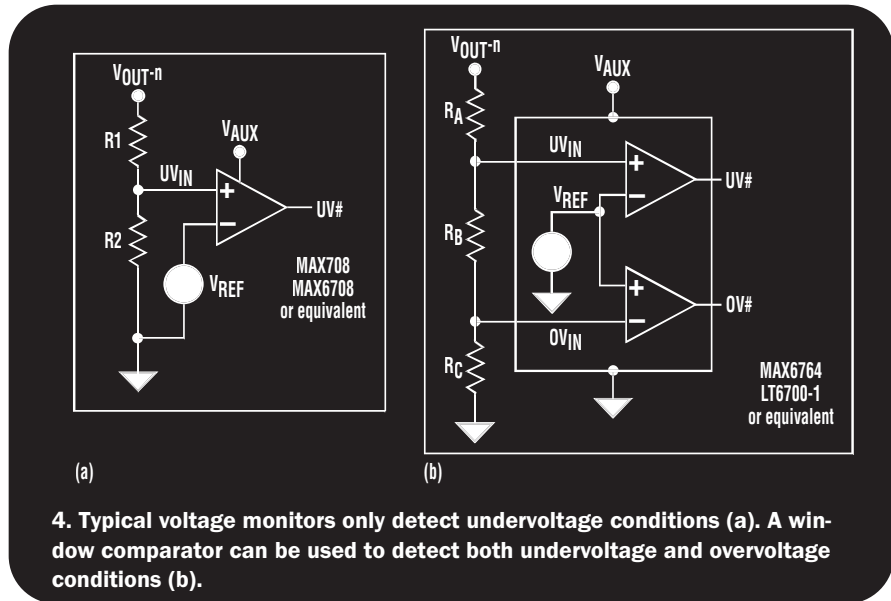
Upon completion of system thermal validation, it’s a good idea to perform thermal fault testing. Here, the unit is operated beyond its supported ambient temperature rating, and airflow is restricted by obstructing vent holes. Under these various fault conditions, the temperatures are monitored until thermal shutdown occurs. In all cases, a tripped thermal sensor should shut down the system power before there’s permanent damage.

OUTPUT OVERVOLTAGE PROTECTION • Due to shrinking feature sizes, modern low-voltage IC technology is particularly susceptible to overvoltage conditions. An electrical overvoltage on the supply, even if brief, can cause the semiconductor junctions to break down. This can create a low-resistance path through the device, which then becomes a localized heater.

If the resistance isn’t low enough to trip a converter’s overload protection circuit, the failed device may go into thermal runaway. What results is smoke, flame, burnt boards and components, and a nasty smell. If this happens in the development lab, it’s an event. If it happens to the customer, it can be a nightmare. Therefore, it’s vital to detect and adequately respond to any output overvoltage condition before such damage can occur.

Purchased ac-dc power supplies and isolated dc-dc units have features to prevent sustained overvoltage conditions. First, the isolated nature of the power train prevents any shorted MOSFET from causing an output overvoltage. Second, a separate reference is used to monitor the voltage at the output of the device. If the output voltage becomes too high, a separate signal crosses the isolation boundary to latch off the pulse-width-modulation (PWM) controller. In short, these products are designed and tested so that no single-point failure will cause a sustained output overvoltage condition.

Though quite reliable, the non-isolated synchronous-buck converter (the most common niPOL) can exhibit a sustained overvoltage condition in response to several single-point-failure conditions (Fig. 3). First, consider what would happen if one of R1’s solder joints were to open. This breaks the feedback signal to the error amplifier, causing the controller to think the output voltage is too low. In response, it will go to maximum duty cycle, causing



4. Typical voltage monitors only detect undervoltage conditions (a). A window comparator can be used to detect both undervoltage and overvoltage conditions (b).

almost the entire input voltage to appear across the output.

Next, consider if MOSFET Q1 were to fail as a short circuit. This would drive the switching node high and attempt to impress the input voltage across the output. Most controllers would respond to this condition by turning on Q2, either by lowering the duty cycle or through its internal overvoltage detection. Large supply currents would then flow through the two MOSFETs in the hopes of overloading the input. But if this is a very high-power bus, the outcome may be uncertain. Finally, if the reference inside the controller were to fail, this also could allow the output to exceed the intended set point.

As with overtemperature protection, more circuitry can be added to each niPOL converter to prevent these failure modes from causing an output overvoltage. In response to a disconnected feedback or failed reference, an additional reference and comparator can be used to monitor the output signals and, if necessary, shut down the controller. In response to a shorted MOSFET, an input disconnect switch or in-line fuse with an SCR can be used to disrupt input power to the converter. Yet these circuits negatively impact the power-system design in terms of their cost, space, and power loss. Once again, a system-level approach can be used to realize the more elegant solution.

As seen in Figure 1, each output already has a voltage-monitoring circuit. In most cases, designers choose one of the common voltage-monitoring ICs available in the market, as shown in Figure 4a. Here, the voltage being monitored is compared with an internal reference voltage. A single digital output determines if the monitored voltage is below or above the reference voltage. For example, a 1.8-V supply being monitored will cause this signal to become active once the output reaches 1.71 V, or 95% of its nominal value. Because this circuit gives only one piece of information (above or below a threshold), it’s only used to monitor output undervoltage (UV#). A typical design holds the system in reset unless all supply voltages are above their UV thresholds.

But by adding a comparator and a resistor, the circuit becomes

the window comparator of Figure 4b. The window comparator provides not only the existing UV# signal, it also supplies an overvoltage (OV#) signal to detect if the output voltage is too high. Multiple suppliers offer this component in the same size package as the single comparator, so only one additional resistor is required. This OV# detection also could be realized with a separate reference and comparator.

As with a thermal fault condition, the only appropriate action in response to an output overvoltage is to power off the system. Therefore, if the OV# signal were to become active, the power-management circuit would respond by immediately latching off the intermediate bus converter and asserting a fault state. This interrupts any fault power that could cause extensive damage. It also may prevent damage to the semiconductors and other loads, allowing the board to be repaired.

The OV# trip voltage should be set high enough to ensure that it's not falsely asserted during startup, load steps, or other transient conditions. But it should assert before the absolute maximum voltage of the semiconductors it's powering. A small capacitor could be placed in parallel with R_C to filter out any noise or to add a small delay.

This monitoring circuit, which is independent of the niPOL

converter, has its own reference. It's placed at a convenient location on the system board, typically close to the main load. Since the outputs are digital, they can be routed back to power-management circuitry with much less concern for noise coupling. Also, if the OV signal is an open-collector, all OV monitors can be diode-ORed together to create one master OV signal to the power-management controller, minimizing design complexity.

Power-system designers must ensure the total power system behaves as expected when put together. This includes response due to various faults. Though many faults create benign conditions, such as the product ceasing to operate, some can cause very dramatic events. Concentrated heating and power dissipation under fault conditions should be avoided whenever possible.

By using the "always on" auxiliary supply and existing power-management circuitry, overvoltage and overtemperature protection can be realized at the system level with relatively minor cost, component, and design impact. Disabling the high-power intermediate bus removes the energy that might otherwise inflict significant damage to the system.

ED Online 10219

BOB THOMAS is a power-system designer at Cisco Systems Inc. (www.cisco.com), San Jose, Calif. He holds a BSEE from the University of California, Davis.