 By Bob Neidorff, Texas Instruments High Performance Analog Products

MAINTAIN GOOD POWER TO ALL MODULES WHILE REMOVING BAD MODULES AND INSERTING NEW MODULES INTO A RUNNING SYSTEM.

A universal approach to hot-swapping power in telecommunication systems

TODAY'S MOST IMPORTANT SYSTEMS can't tolerate downtime. We never want to shut down telephone exchanges, data-network equipment, cellular-communication facilities, or hospital information systems. However, even the most robust systems with redundant datapaths, power sources, and storage can experience random instances of failure. It can come in the form of component failure; stress-induced failure from unpredictable events, such as lightning strikes; and the need for upgrading systems on the fly. As a result, designers engineer these systems for at least limited operation with a failed module and for replacement of failed modules without shutting down the system. This practice is known as *hot-swapping*.

Hot-swapping has many implications. For example, software must adapt to adding and removing modules, and data buses must tolerate transient errors caused by connecting and disconnecting live bus transmitters and receivers. This article focuses on a third consideration: maintaining good power to all modules while removing bad modules and inserting new modules into a running system.

POWER IMPLICATIONS OF HOT-SWAPPING

Figure 1 shows a common redundant-power-distribution architecture for high-availability systems (**Reference 1**). Because every component is a potential source of failure, engineers design every part of the system for hot-swapping. One exception is the backplane itself. The system's backplane, or superstructure, is so intertwined with every function, that replacing the backplane would require removal of every element—every board, every bracket, and every wiring harness. So the first priority in system engineering is to obtain the highest possible backplane reliability, which means using a backplane with no components except for well-engineered wiring, redundant interconnect components, and easily replaceable parts, such as fuses.

This system supplies power to each module from two high-reliability power-distribution busses (**Reference 2**). In large systems, these power sources are a long distance from the rack. To minimize voltage-drop-related problems, power conductors are very thick, and the power distribution voltage is high. In the telecommunications industry, the standard for power distribution is nominally −48V, but the exact voltage level varies widely due to changes in load current, resistance and inductance in the power-distribution network, and status of the power source—normal, brownout conditions, or operating from batteries. A negative supply minimizes corrosion in the presence of leakage paths, because negative voltages repel negative ions, which tend to corrode metals.

To efficiently convert low-current, high-voltage power to low voltage at high current, each board or module includes dc/dc converters. Even though these dc-to-dc converters use sophisticated high-frequency switching techniques, they require a low-impedance source at their input for transient response, stability, and protection against dropouts. Long inductive-power-distribution cables are inadequate for this task, even with remote sensing, so the input to each dc/dc converter must include large electrolytic capacitors.

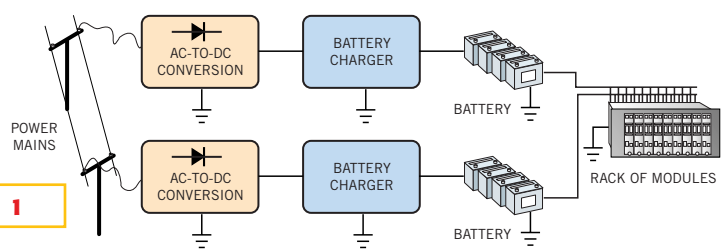

Figure 1 High-availability systems commonly employ this redundant-power-distribution architecture.

If you plug simple boards with large electrolytic capacitors into a powered backplane, large inrush surges will occur. These surges can cause the backplane voltage to droop, which can reset or otherwise interfere with adjacent boards. Surges can also abuse connector pins, overstress capacitors, and interfere with data by generating EMI. A large enough surge could shut down the whole system. To prevent these surges, each board in the system includes a hot-swap circuit that limits inrush current.

Hot-swap circuits slowly ramp power to the module to prevent disruptive power dips on the backplane. In addition, hot-swap circuits can disconnect a module if the module attempts to consume excessive power, retry the module after a fault-clearing delay, and allow communication of module status and commands to and from a system monitor.

In addition to ramping power to the electrolytic capacitor, the hot-swap circuit drives one or more dc/dc converters, which in turn drive a range of loads, such as complex digital processors, lasers, and fans. However, the dc/dc converter manages the unique power requirements of these varied components, so the task of hot-swap-circuit design is mainly a challenge of applying power to a large capacitor.

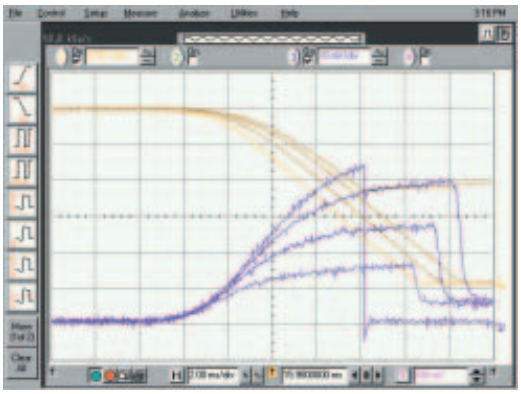

Figure 2 These oscilloscope traces show controlled-voltage slew rate powering four different load capacitors: 50 μF, 100 μF, 150 μF, and 200 μF, plus a fixed resistor. The yellow traces show load voltage ramping down to −48V, and the purple traces show load current stabilizing at a level proportional to the load capacitor, then dropping to normal operating current. In this example, turning the driver on into a 200-μF load produces too much inrush current and the circuit latches off.

REAL-WORLD POWER LEVELS

Communications systems distribute power at a nominal −48V. However, dc/dc converters allow for operation over a wide dc-voltage range, perhaps −36 to −72V (**Reference 1**), to allow for different modes of operation, voltage drops in the power-distribution bus, and temperature variations.

In addition to this operating range, there are often voltage surges on the power bus. One of the worst causes of surges is the sudden change in load current when a fuse disconnects a defective, excessively high-current module from the inductive power bus. One power-surge standard allows for maximum voltages of −75V for 10 msec, −100V for 10 μsec, and −200V for 1 msec (**Reference 3**).

The hot-swap circuit is directly exposed to these extreme power surges and can help manage the surges by disconnecting the load from the power bus if the input exceeds a maximum safe operating voltage. However, the worst surges are too severe for the most rugged hot-swap circuit. Therefore, many systems include transient suppressors (metal-oxide varistors, Transorbs, and others) to absorb the largest, short-duration surges.

ACTIVE HOT-SWAP TECHNIQUES

Hot-swapping requires power FETs with electronic drive to control inrush current. Various ICs can drive power FETs. Some control the voltage slew rate on the load by limiting inrush current to $I=C_{LOAD}\,dV/dt$. If the load capacitance is well-known and dominates load impedance, controlling voltage slew rate can control inrush current. However, designers must optimize the slew rate of each implementation for the expected load capacitance.

These ICs often contain current-sense circuits with current limiting, logic inputs to turn on and off the load, and logic outputs to report on the state of the load. Current-limiting circuits act like circuit breakers with controlled response time and accurate trip currents. When load current exceeds a programmed maximum, the IC disconnects the load. If an overcurrent occurs during turn-on, the surge resulting from the overcurrent is minimal, because the current limit catches the surge as the power ramps up. However, if overcurrent occurs after the turn-on slewing is complete and the power FET is fully enhanced, high inrush surges are still possible, because of the delay required to ensure that the fault is real and also because of the time required to discharge the capacitance on the power-FET gate.

The most universal approach to hot-swapping is directly controlling load current using an LCA (linear current amplifier). The LCA combines current sensing with a high-gain amplifier tailored to drive a power FET. When you insert or turn on the board in a system with an LCA, a current command drives the LCA input to a level representing maximum load current. The LCA regulates load current to an exact level while the load capacitor is charging, independent of the load capacitor's size. This technique is universal because one hot-swap implementation works with a range of loads and automatically optimizes load-capacitance charging time. Though controlling the load-voltage slew rate reduces inrush current, load characteristics ultimately determine the exact inrush, requiring redesign for each load.

TABLE 1—CONTROLLED-VOLTAGE-SLEW-RATE HOT-SWAP VERSUS CONTROLLED-CURRENT HOT-SWAP

	Controlled-voltage slew rate	Controlled current
External parts	Requires two extra capacitors and a resistor; current-sense resistor optional	Requires current-sense resistor
Stability	Inherently stable	Relies on internal compensation for stability
Load-capacitor dependence	Must be tailored to the particular load capacitance	Independent of load capacitance
Load-impedance dependence	Load voltage is linearly ramped; load current can surge if the load impedance is nonlinear	Load current is always under control; independent of load impedance
EMI	Load-current surges can interfere with other systems	Implements controlled dI/dt for lowest EMI
Fault behavior	Shuts down in a transient or a sustained fault	Regulates in a transient, then times out if the fault is sustained

Figure 2 shows a controlled-voltage slew-rate hot-swap controller applying power to four capacitive loads, and **Figure 3** shows a controlled-load-current hot-swap controller applying power to the same four loads. With the controlled-voltage-slew-rate controller, load capacitances larger than 150 µF cause inrush current to exceed the trip point, forcing the circuit to shut down. The controlled-load-current design turns on smoothly regardless of load and charges the load capacitor with the same peak current.

In addition, the controlled-load-current hot-swap circuit can satisfy momentary demands for high load current, allowing systems to operate at peak performance without overstressing components. A timer in the hot-swap circuit allows controlled, short bursts of high current, and disconnects in the presence of sustained demand for high load current.

In some cases, system-noise issues require limiting the load-current slew rate to a set maximum. A designer can easily program controlled-load-current hot-swap circuits for controlled slew rate by using a capacitor to control the input to the LCA.

For these reasons, the hot-swap circuit containing an LCA is universally preferred. It controls inrush current independent of load capacitance, allows controlled brief inrush, disconnects when inrush time exceeds a programmed maximum, and easily controls load-current slew rate.

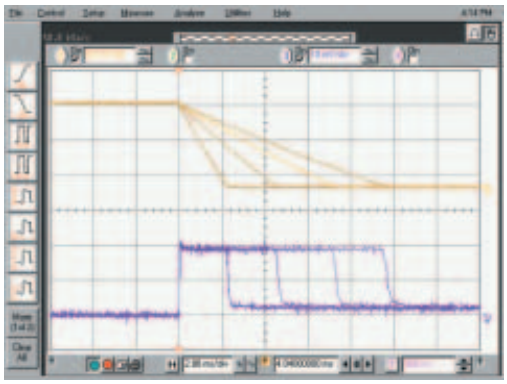

Figure 3 These oscilloscope traces show controlled load current applying power to the same four loads as Figure 2. In this example, the yellow traces show that load-voltage slew rate varies with load capacitance, but the purple traces show that peak load current never exceeds a programmed value. All four loads are successfully powered.

CONTROLLED-CURRENT HOT-SWAP

Controlled current is not without drawbacks, however. The controlled-voltage slew rate's advantage is that it is inherently stable. The controlled-voltage slew-rate implementation doesn't require closed-loop control and simply relies on open-loop current into a capacitor to set slew rate (**Figure 4**).

By comparison, controlled-current hot-swap places the power FET and current-sense resistor in a feedback loop (**Figure 5**). This loop must be compensated in all operating modes to prevent instability. If the loop becomes unstable or marginally stable, output current could overshoot and prematurely trip current-limiting circuits. Fortunately, modern controlled-current hot-swap ICs contain a compensated LCA with careful consideration of loop stability under all possible conditions, removing stability concerns (**Reference 4**).

Controlled-current hot swap requires either a current-sense resistor or another current-sense element. Most systems use current sensing of some sort, so this requirement isn't typically a drawback, but it does put another constraint on the technique.

HOT-SWAP-IMPLEMENTATION DETAILS

Hot-swap ICs controlling the load-voltage slew-rate drive power-FET gates with a constant current and rely on a gate-to-drain capacitor to control slew rate. In this way, load slew rate is set at $dV/dt = I/C_{GD}$. In a high-voltage system, this capacitor must be a high-voltage capacitor. Any power-FET gate-drain capacitance adds to the external capacitor value, so, to maintain good control, the added capacitor should be much larger than the internal capacitances of the power FET.

Once the load voltage reaches full supply voltage, the drain stops slewing, but the gate current continues charging the external capacitors to some steady-state V_{GS}—typically, 12V. In the event of a fault, the IC must quickly discharge the power FET's internal capacitance and external gate capacitors from 12 to 0V to prevent damage and interference to adjacent modules. To keep fault-response time low, the external capacitor should be as small as practical. However, this situation leads to an imperfect compromise, because it conflicts with the need for an external capacitor large enough to maintain tolerance.

Hot-swap implementations that use controlled-voltage-slew-rate detect the instant of load overcurrent, set a fault latch, and shut down the load. This implementation is safe, but the resulting system is subject to shutdown whenever a transient produces an instantaneous overcurrent—a situation that might occur if you insert a second supply into the system, the mains turns back on after a period of operating from battery power, or the input voltage surges from a fuse-clearing transient.

Supply transients and turn-on events can couple from the drain to the power FET's gate through the external capacitor, causing a momentary current surge. To prevent this situation, practical implementations include a resistor in series with the external gate-to-drain capacitor and a second external capacitor from gate to source (**Figure 4**). However, you must also discharge this second capacitor in a fault, which degrades fault-response time.

By comparison, controlled-current hot-swap systems use an op amp, specifically an LCA, to keep load current under control at all times. The input to the LCA is a voltage command proportional to the load current. The LCA adjusts gate voltage on the external FET to the level necessary for maintaining load current at the input command voltage divided by the current-sense resistor.

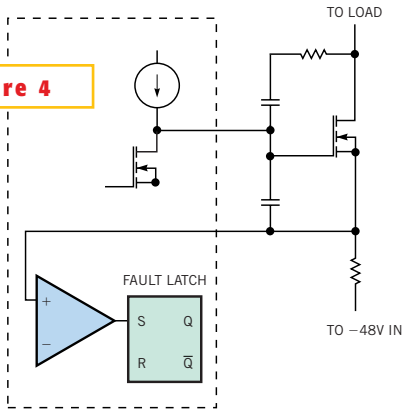

Figure 4 A controlled-voltage slew-rate hot-swap implementation employs two external capacitors and a resistor around the power FET.

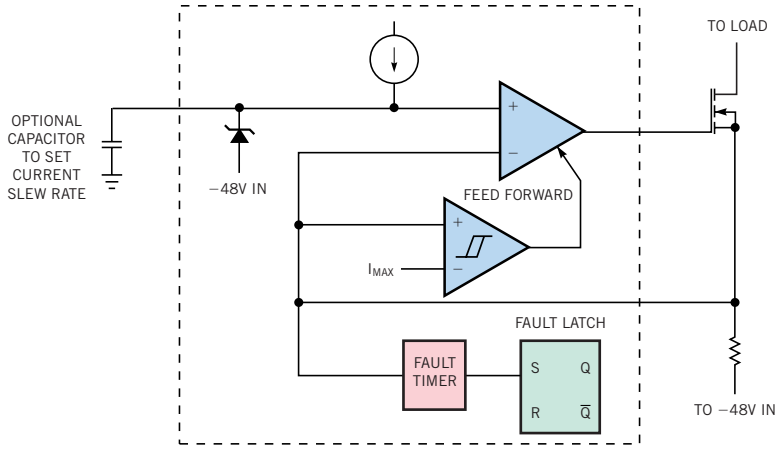

Figure 5 **A controlled-load-current hot-swap implementation places the power FET and current-sense resistor in a feedback loop.**

The voltage command is zero when the system is switched off. When you turn on the system, the input-voltage command is stepped from zero to the desired maximum current. The output current ramps up to the fully regulated current as quickly as the LCA can drive the external FET, usually 100 μsec. After that initial turn-on ramp, the LCA maintains load current at the commanded level until the load capacitor charges to the supply voltage, after which the LCA saturates, overdriving the power FET and applying approximately 12V from gate to source.

Reduced load-current slew rate can be advantageous in reducing system EMI, so, if desired, you can reduce turn-on load-current slew rate by adding a capacitor to the LCA's input (**Figure 5**). However, using a controlled load-current IC requires neither an external high-voltage gate-to-drain capacitor nor an external gate-to-source capacitor.

In a controlled-current hot-swap system, the LCA is typically compensated for low loop bandwidth to prevent stability problems. This design feature permits free choice of power FETs without concern for loop compensation. It also means that the current loop itself is too slow to respond to current surges caused by an instantaneous change in the input voltage. These surges might occur when you insert a second power supply into a system, when you restore main power after battery operation or after a fuse-clearing event. To deal with these transients, controlled-current hot-swap ICs contains a comparator for sensing abrupt overcurrent conditions and a feedforward path in the LCA to quickly bring the loop back into regulation.

Compared with the system using controlled-voltage slew rate and two external gate capacitors, the system using controlled current has much lower total gate capacitance. The only load on the LCA is the power FET's own internal capacitance, meaning that the controlled current implementation is inherently much faster at responding to a fault.□

References

1. White, RV, "Computers In The Central Office—A Primer On Powering Equipment from -48V," APEC '98 Conference Proceedings, pg 902.

2. Hunter, PL, and TW Johnson, "Fault Tolerant Hot Pluggable Power System Design," APEC '94 Conference Proceedings, pg 815.

3. "A Module of Telcordia LSSGR, FR-64."*GR-513-CORE Issue 1*, September 1995, Section 4.2.3.

4. Texas Instruments TPS2398/TPS2399 Product Data Sheet, http://focus.ti.com/lit/ds/symlink/tps2398.pdf.

5. Texas Instruments Application Note SLUA283, "Universal Telecommunications Hot Swap Device Family," http://focus.ti.com/lit/an/slua283/slua283.pdf.

Author's biography

Bob Neidorff is an engineering manager for high-performance analog products at Texas Instruments (Manchester, NH), where he designs and tests power-management ICs. He holds bachelor's and master's degrees in electrical engineering from the Massachusetts Institute of Technology (Cambridge) and enjoys table tennis, metalworking, and woodworking.