INTELLIGENT LOAD SWITCHING HELPS IMPLEMENT RELIABLE HOT SWAP SYSTEMS

A Lattice Semiconductor White Paper

March 2010

Lattice Semiconductor
5555 Northeast Moore Ct.
Hillsboro, Oregon 97124 USA
Telephone: (503) 268-8000
www.latticesemi.com

Introduction

Almost invariably, servers, disk arrays and other high-availability systems are subject to the requirement that functional modules be replaceable on-the-fly without the need for powering the system down. Replacing modules while the system is in operation is commonly known as *hot swapping*. A key factor in being able to provide hot swap capability is the appropriate management of the local power systems on each of the interchangeable modules.

To support hot swapping, a subassembly such as a printed circuit board must reliably perform several operations, and power management is among the most critical. When inserted into the larger system, a board's *hot swap controller* must ensure that power at the connector is stable before being distributed to other onboard systems and commanding the board to initialize. For boards drawing little current, power can be switched onboard with digitally switched MOSFETs. Higher levels of current (e.g. 10's of amperes) require more sophisticated switching strategies to avoid inducing transients in the bus-level power supplies and damaging the current-switching MOSFETs. Maximizing system-level performance requires close coordination of the low-level functions of switch control with those of the top-level hot swap management functions, such as sequencing and fault detection. While it is possible to implement such systems with hardwired circuitry, it is often more simple and cost-effective to do so using programmable system components.

Hot Swap Power Switching

The power pins on a hot swappable module's connector typically are not connected directly to the module's internal power buses. The more usual arrangement is for the local board power to be isolated from the bus power with MOSFETs or other types of power switching devices, as shown in the schematic of Figure 1. This circuit controls the connection of a single +12V bus power supply to board-side power through MOSFET M_1, and is based on Lattice Semiconductor's ispPAC-POWR-1014 power management device. Some of the individual functions performed by the other parts of this circuit include:

1. Voltage monitoring – through resistive dividers R_1/R_2 and R_7/R_8
2. Current Sensing – through R_{SENSE} and a ZXCT1009 differential amplifier
3. High-side MOSFET drive – the CHARGE_PUMP signal from the ispPAC device is a square-wave used to develop a high voltage (> +12V) across C_2 that can be used to fully switch N-channel MOSFET M_1 on. M_1's gate voltage is controlled by the SHUT_DOWN signal buffered by Q_2

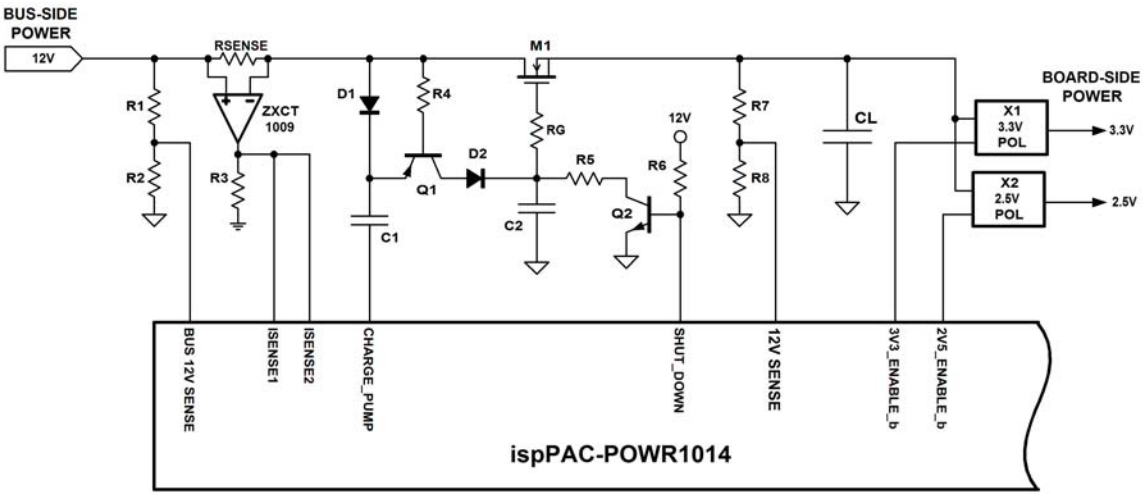

Figure 1 – A hot swap controller typically employs a power MOSFET (M1) to connect the bus power to the internal board power systems after the board has been positively seated in its socket. Lattice's ispPAC-POWR1014 provides intelligent control that can be tailored to the application's requirements through its programmable logic and analog functions.

For modules consuming small amounts of power with minimal internal supply capacitance, one could power up the module simply by switching M_1 into its low-resistance on-state (hard-switching) quickly. In the case of modules with greater power requirements, however, this will result in a large turn-on current transient through M_1, because when C_L is in the process of charging it will appear as a momentary short circuit between the supply rail and ground. The resulting current transient creates two problems. The first is that it can cause voltage drops in the bus power supply, potentially affecting the operation of other modules sharing bus power. A second problem with this scheme is less subtle: the current transient may damage the MOSFET, resulting either in reduced long-term reliability or outright failure.

When charging a capacitor (C) up to a voltage source (V) through a MOSFET or other purely resistive device, the total amount of energy dissipated in the MOSFET is $CV^2/2$, the same as that which ends up being stored in the capacitor. This is independent of the value of the MOSFET's on resistance or the time needed to charge. While the total amount of dissipated energy is not negotiable, the rate at which it's dissipated – the instantaneous power - can be controlled. For example, using a MOSFET with a small on resistance results in high power dissipation for a short duration, while a device with a larger on resistance will experience lower power dissipation over a longer duration. Making the appropriate tradeoff between maximum power dissipation and the time needed to charge a module's local capacitance is a key part of implementing an effective hot swap design.

The amount of time a MOSFET can safely dissipate a given amount of power usually is specified in the device's data sheet by its safe operating area (SOA) chart (Figure 2). The SOA chart specifies the maximum amount of time the MOSFET can safely remain biased under various combinations of drain-to-source voltage (V_{DS}) and drain current (I_D). Superimposed on this SOA chart are “trajectories” of V_{DS} and I_D corresponding to the control schemes to be discussed.

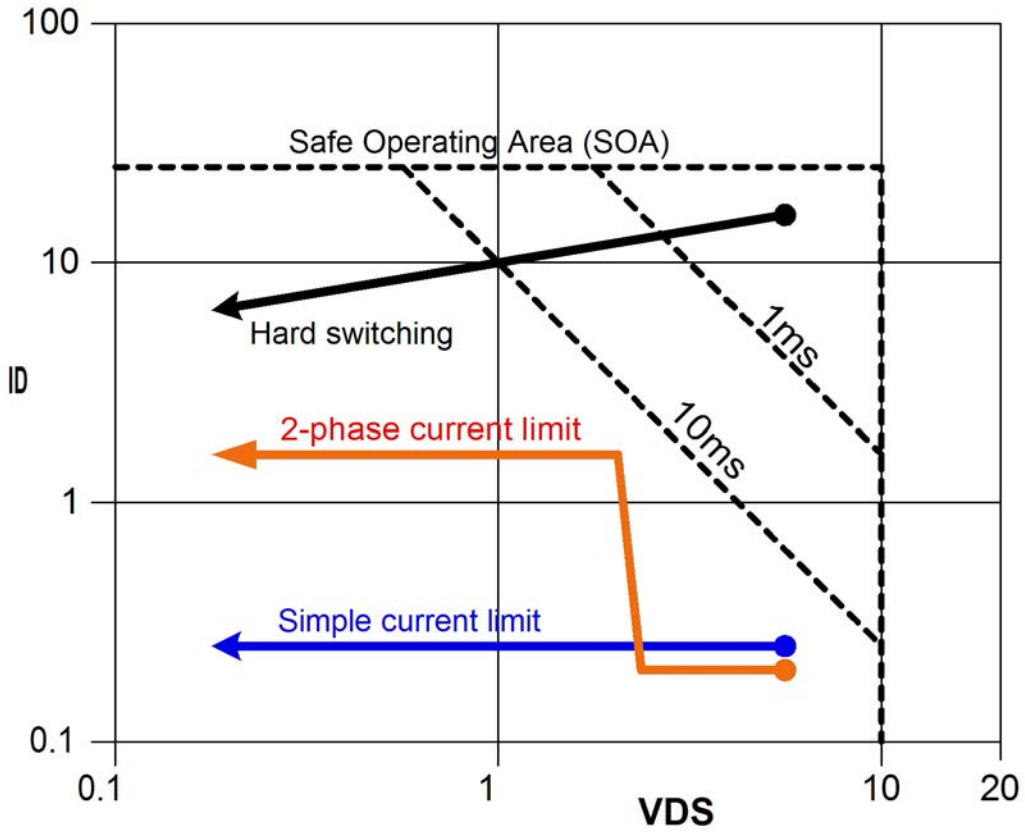

Figure 2 – The Safe Operating Area (SOA) chart for a MOSFET shows the device's safe endurance limits for combinations of drain-to-source voltage and drain current. Different control strategies can be used to avoid excessively stressing the MOSFET when used as a hot swap load switch.

If one switches the MOSFET on quickly (shown by the black “Hard switching” curve of Figure 2), it is subjected initially to a maximum value of V_{DS}, with I_D limited only by the channel resistance and parasitic impedances such as PCB trace resistance and inductance. As the load capacitor charges, the MOSFET's operating point moves left and down to more benign conditions. If the operating point does not shift quickly enough, the MOSFET may be damaged or destroyed. And even if one selects a MOSFET with sufficiently high power dissipation capabilities, the problem that the initial current surge could disrupt bus-side power supplies still remains.

Controlling Current Surge

One common technique used to avoid the problems resulting from abruptly turning on the MOSFET is to ramp up the MOSFET's gate voltage gradually, at a rate that is slow enough so that the load capacitor's voltage will track with a minimal V_{GS}. This ensures that the operating point will remain in a low-current region near the bottom of the SOA chart. This strategy can be readily implemented with the circuit of Figure 1 by appropriately selecting the value of C_2.

While simple to implement, ramp-up up rates in the above scheme must be heavily margined to accommodate component variation in both the MOSFET and power-bus load capacitance. For low to moderate current applications, specifying a slightly larger MOSFET may not be much of an additional expense, and may well be justified by cost-savings from simplified control. In other cases, where large on-board capacitances must be charged, this approach can result in significant delays between the time when a module is plugged into the larger system and when it is ready to begin operation.

With current-sensing hardware, it becomes possible to maintain a constant current trajectory through the SOA by using negative feedback control. By providing precise regulation of drain current, the MOSFET switching trajectory (Figure 2, simple current limit) can be set to a higher current level than would be prudent in the previous case of simply ramping up gate voltage in open-loop fashion. Because it is often necessary to

monitor current for diagnostic purposes, current-sensing hardware may already be available in a design, in which case only the control logic need be added.

Linear current-mode control, however, can be tricky to implement reliably, with the potential for instability and uncontrolled oscillation. An alternative is to use a hysteretic control scheme (Figure 3), in which the MOSFET current is maintained between a lower and an upper threshold. In a hysteretic control scheme, MOSFET gate voltage is ramped up until the drain current reaches a predetermined upper threshold. At this point, the gate voltage is ramped down until drain current falls below a predetermined lower threshold. The process then repeats, with drain current varying between the two threshold levels.

While hysteretic control could be implemented with a small number of discrete components, it can also be realized by merely reprogramming the ispPAC Power Manager in the original circuit of Figure 1. Each of the Power Manager device's voltage monitor pins supports dual comparator functions with independently programmable high and low voltage thresholds. Programming the device's FET driver outputs to a higher current provides a faster but still controlled rise time for MOSFET gate voltage and corresponding drain current. An additional advantage of using a programmable device to manage the current control process is that it is straightforward to fully integrate the hot swap control logic with that needed for normal board operation. For example, the Power Manager device can be programmed to permit higher currents during the short

time the board is initializing, and then to seamlessly shift into normal operating mode, where the MOSFET is completely turned on and current is monitored at a lower threshold to detect board fault conditions.

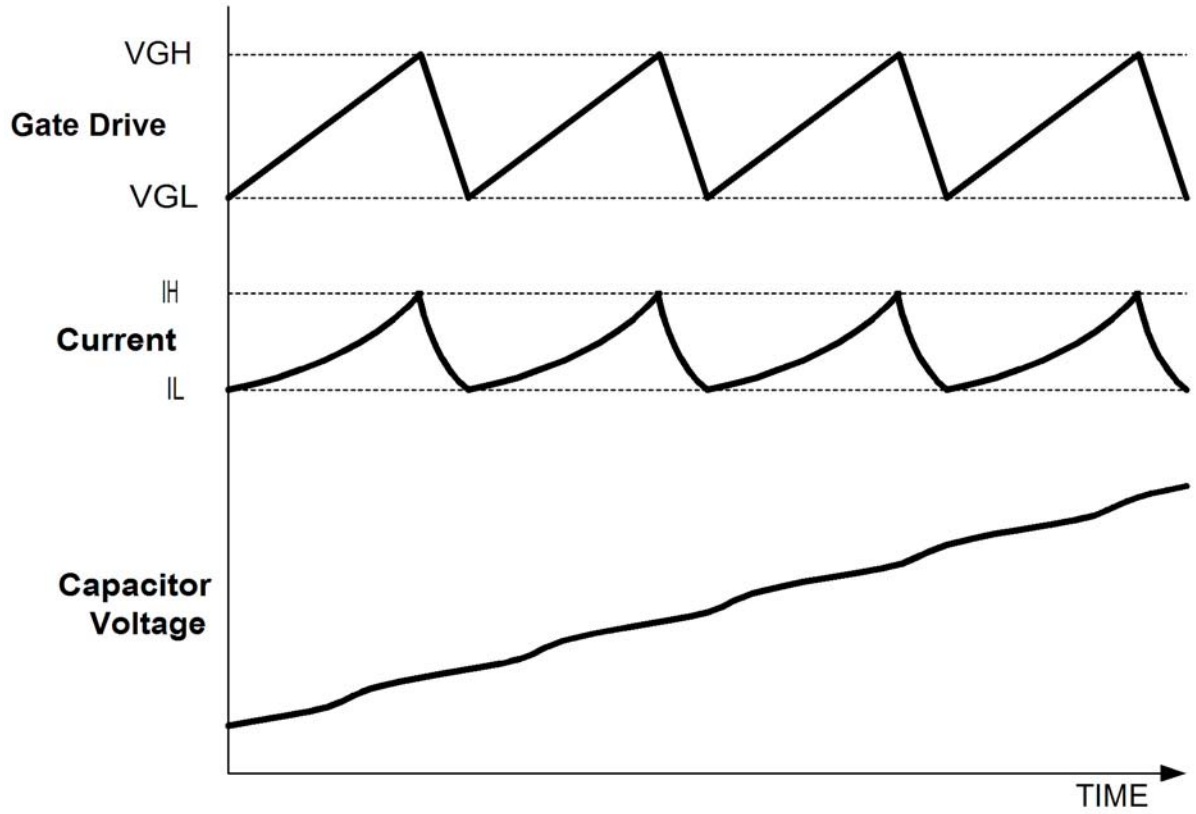

Figure 3 – Hysteretic current limiting cyclically allows MOSFET current to increase to a maximum allowable level, and then throttles it back to a slightly lower level. This technique provides many of the advantages of linear current control while sidestepping many of the potential stability issues.

Optimizing Switching Performance

The programmable nature of the Power Manager device supports the designer in realizing more highly optimized control techniques at little or no additional cost. One example of such a technique is charging up board capacitance in two separate phases – a low current initial phase and a high current final phase, as shown in Figure 4. The value of adding this complexity is that it optimizes charge rate so that the MOSFET's operating point more closely follows the constraints of the device's SOA curve (2-phase current limit plot in Figure 2). This provides two benefits over the constant

current charging scheme described previously. First, by switching to a higher current midway through the charging cycle, less time is needed to bring the load capacitor up to operating voltage. Second, this scheme only operates the MOSFET at higher current levels when the VDS across the device is relatively low, minimizing power dissipation. This in turn allows the designer to specify a smaller, less expensive MOSFET for a desired level of performance.

While this technique is straightforward to implement with a programmable hot swap controller such as an ispPAC Power Manager device, it would require significant additional hardware and design effort if realized using a fixed-function hot swap controller. Programmability also makes it straightforward for the designer to vary control parameters, making it easier to adapt the same basic circuit for use in multiple applications with few or no hardware changes.

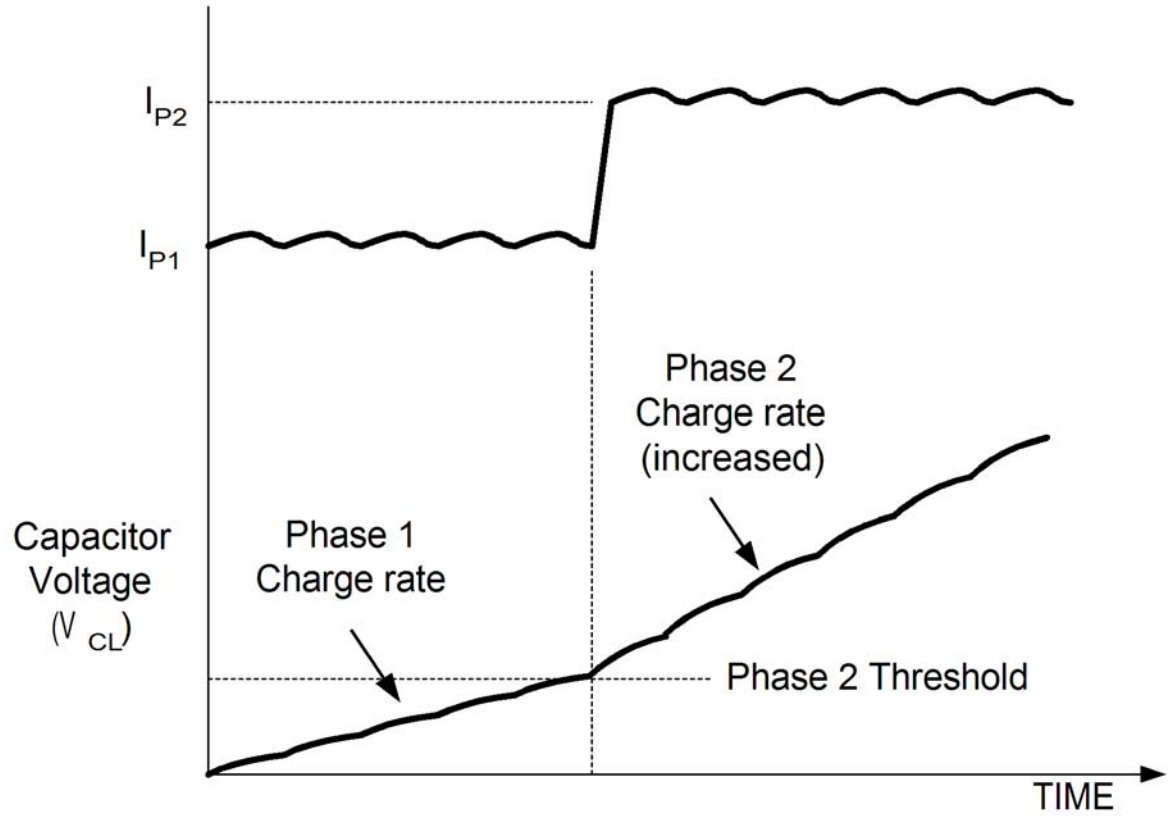

Figure 4 – A two-phase switching strategy first charges the on-board capacitors at a low current when the switch MOSFET is subject to a high V_{DS}, and then increases the current when the capacitors are partially charged and the MOSFET experiences lower V_{DS}.

Summary

Supporting hot swapping features requires sophisticated power management techniques. While fixed-function controllers may be useful in the most simple designs, programmable solutions provide the designer with the flexibility needed to meet increasingly demanding applications. This article has described one example of how the performance of a load switch can be enhanced through the use of programmable control.

###