AdvancedTCA Hot Swap Controller Eases Power Distribution

Design Note 353

Mitchell Lee

Introduction

AdvancedTCA is a new modular computing architecture developed by the PCI Industrial Computer Manufacturers Group for use in central office telecom environments. PICMG 3.0 defines, among other things, the electrical and mechanical attributes of the backplane, connectors and removable cards in these systems.

System power is supplied by the –48V dual battery feed typical of telecom installations, and ATCA borrows many of its related specifications from established telecom standards.

Power Requirements

Each removable card or front board, is designed for live insertion into a working system. A power draw of up to 200W per front board is allowed, placing the maximum load current in the 4A to 5A range.

As is common in these types of systems, card-centric inrush limiting and current and voltage monitoring are highly desirable to sanitize the incoming battery feeds and minimize power plane disturbances. The LTC®4252A Hot Swap™ controller is a good match for –48V, 0W to 200W applications.

Circuit Solutions

Figure 1 shows a complete circuit designed to handle up to the maximum available power. The LTC4252A's accurate current limit is set to provide at least 5.5A under all conditions, a comfortable margin for 200W, yet trips off just under 7A to preserve fuse integrity in the presence of nuisance overloads.

Over and undervoltage monitoring are both implemented in this circuit. UV thresholds are set at –37V turning on and –33.3V turning off, as measured after the ORing diodes. OV turns off at –74.7V and back on at –73.2V,
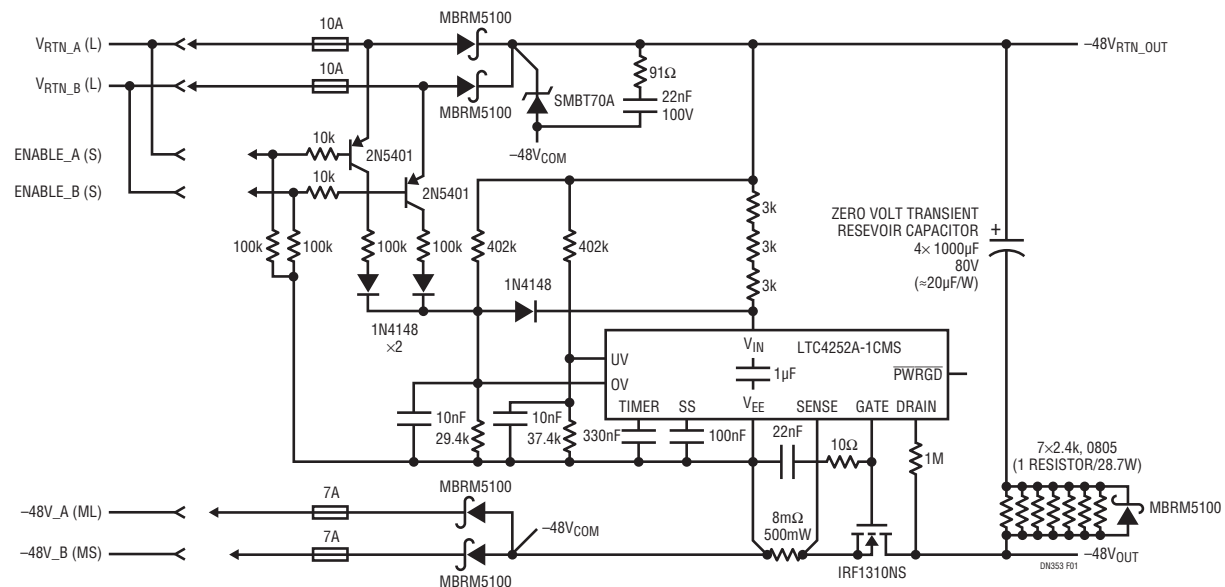
Figure 1. 200W AdvancedTCA Hot Swap Controller Circuit

again as measured after the ORing diodes. This assures operation over the full range of –43V to –72V, as well as under conditions of input surge to –75V and transients to –100V in accordance with ATCA specifications.

Once insertion is detected, the LTC4252A pauses for 230ms to allow for contact bounce, then soft-starts the load using a ramped current scheme. Inrush current is increased gradually until the MOSFET is fully on.

Current overloads detected by the SENSE pin and 8mΩ shunt are dealt with by three distinct levels of response. If a small, sustained overload of 7A or more is detected, the TIMER pin delays for 5.7ms before shutting down. If the overload exceeds 7.5A the LTC4252A throttles back the MOSFET and holds the current to that value. Again, after a 5.7ms delay the circuit shuts down. If the overload is severe, a strong and very fast amplifier quickly corrects the MOSFET gate voltage, bringing it down to near threshold for the device. Then the LTC4252A's current limit circuit takes over and maintains 7.5A for the duration of the 5.7ms TIMER delay period.

The LTC4252A also monitors the voltage drop across the MOSFET and reduces the TIMER delay to as little as 1.8ms as the voltage stress increases. This keeps the MOSFET comfortably within its safe operating area in the presence of hard faults.

Zero Volt Transient

The so-called Zero Volt Transient requirement is a legacy of earlier telecom equipment standards stipulating uninterrupted system operation during the course of a 5ms input voltage dropout. An energy of 1J is needed to sustain a 200W load during this interval.

In addition to the energy storage requirement, the connector pin configuration presents a special design challenge. Front board insertion is detected by two short pins ENABLE_A and ENABLE_B. Rather than simply looping through the backplane, these pins connect to V_{RTN_A} and V_{RTN_B} which complicates their use for insertion detection.

If ENABLE_A and ENABLE_B were used to detect insertion by directly driving the LTC4252A's UV pin resistors, the MOSFET would turn off at the first sign of a voltage drop, regardless of the cause. The LTC4252A would initiate a new start-up cycle when the input was restored, complete with debounce and soft-start. If

recovering from a zero volt transient, the 1J storage mechanism would deplete long before the MOSFET recovered, interrupting front board operation.

Here extraction is inferred from the difference between each ENABLE and its associated V_{RTN}, thereby ignoring input dropouts. A PNP transistor pulls up on OV in the event of an ENABLE disconnect while V_{RTN} is powered, shutting down the LTC4252A and permitting safe extraction with no connector damage. In contrast the base and emitter terminals remain shorted during a zero volt transient; no signal reaches the OV pin and the MOSFET remains on. Power flows uninterrupted to the load when the input voltage is restored.

Energy Storage

The accepted method of energy storage to satisfy the 1J requirement is a bulk reservoir capacitor which is charged through resistors. Several benefits arising from the use of the resistors are not immediately obvious. First, transient overvoltages are blocked from the capacitors allowing the use of diminutive, 80V rated units. Second, an energy equal to that stored in the capacitors is dissipated in the charging path. If Q1 alone shouldered this burden a much larger and more expensive high SOA device would be necessary. Instead this charging energy is dissipated in a few resistors.

Computing Energy

The bulk storage reservoir capacitance is calculated knowing a minimum input voltage at the start of a zero volt transient (specified as –43V at the input of the front board), decaying to a worst-case UV detection voltage of –34V. Diode losses result in a load voltage of –41V decaying to –34V over the 5ms dropout period.

In reckoning the necessary capacitance, a common mistake is to apply the familiar $E = (1/2)CV^2$ energy equation, but using ΔV for the voltage term. This is incorrect and leads to an alarmingly big capacitance of over 40,000µF. Fortunately the value is far smaller and using the correct formula:

$$E = \frac{1}{2}C\left(V1^2 - V2^2\right) \qquad (1)$$

is found to be just 3,800µF, for 1J of available energy. Reduce this value in direct proportion to the load in lower power applications.

For applications help,
call (408) 432-1900