SLAC-PUB-13589
April 2009

DIGITALLY CONTROLLED HIGH AVAILABILITY POWER SUPPLY

David MacNair, Stanford Linear Accelerator Center, Menlo Park, CA

Abstract

This paper will report on the test results of a prototype 1320 watt power module for a high availability power supply. The module will allow parallel operation for N+1 redundancy with hot swap capability. The two quadrant output of each module allows pairs of modules to provide a 4 quadrant (bipolar) operation. Each module employs a novel 4 FET buck regulator arranged in a bridge configuration. Each side of the bridge alternately conducts through a small saturable ferrite that limits the reverse current in the FET body diode during turn off. This allows hard switching of the FETs with low switching losses. The module is designed with over-rated components to provide high reliability and better then 97% efficiency at full load. The modules use a Microchip DSP for control, monitoring, and fault detection. The switching FETS are driven by PWM modules in the DSP at 60 KHz. A Dual CAN bus interface provides for low cost redundant control paths. The DSP will also provide current sharing between modules, synchronized switching, and soft start up for hot swapping. The input and output of each module have low resistance FETs to allow hot swapping and isolation of faulted units.

OVERVIEW

This paper describes a module intended to be power DC magnets in a N+1 redundant configuration. The prototype is a 1320 watt (33A 40V) module, but the design can be easily scaled from 100 watts to 100 kW using the same digital control board. The reliability of the module is maximized by controlling all electrical and thermal stress.

Oversized components reduce conduction losses and improve efficiency. The digital control provides cycle by cycle voltage and current control to eliminate voltage or current overloads in transient conditions.

ALTERNATING BUCK REGULATOR

The ideal topology is the synchronous buck regulator. This provides the maximum power output for a given voltage rated FET. Replacing the free wheeling diode with a FET significantly reduces power loss. This also allows for 2 quadrant operation since the FET can conduct current in either direction. A pair of modules can provide bipolar output (four quadrant) when the load is connected between the two ouputs. The non linear behavior at the transition of continuous diode conduction is eliminated, improving operation at low currents.

While synchronous regulators are common at low voltages and powers, the difficulty of synchronizing the switching of large devices has prevented their use at high power. If one device is turning on before the other turns off, massive currents will flow through both devices. In reverse conduction, if a device turns off prior to the other turning on, the current is forced into the intrinsic FET body diode. The typical high power FET diode has a slow recovery time, and a "snap off" recovery. Figure #2 shows the gate voltage and drain voltage of a IRFB3077 during turn off from a 30 amp reverse current. The extremely fast transition of the body diode causes the lead inductances to produce very large voltage transients inside the device.

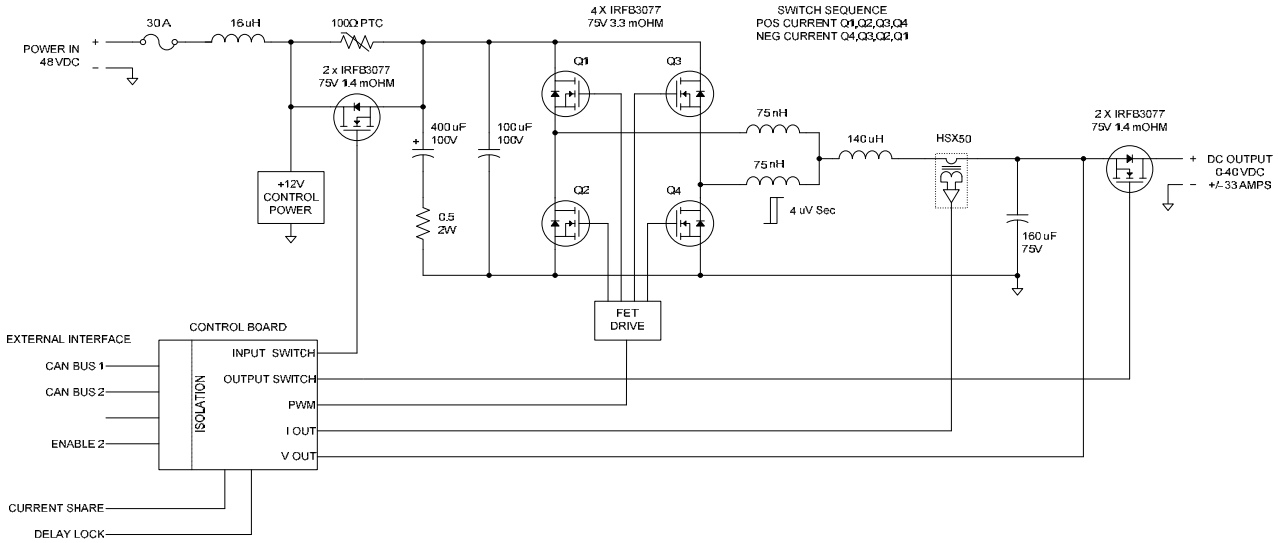

Figure 1. Two Quadrant DC-DC Converter with input and output FET switches

Presented at the 2009 Particle Accelerator Conference, 05/04/2009 -- 05/08/2009, Vancouver, Canada

Work supported by the U.S. Department of Energy under contract number DE-AC02-76SF00515.

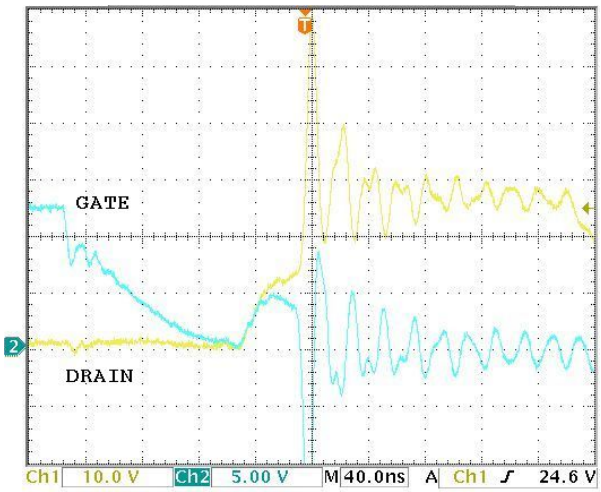

Figure 2: FET Diode Turn off 30 amps

A common approach is to increase the gate resistance of the FETs to slow down the switching speed, but at the cost of increased switching losses. The approach used for this module is to use a saturable inductor in series with the FET. The inductor saturates at below 2 amps and provides a delay of 80 nSec before conducting in the opposite direction. This allows time for the diode to fully recover. With the body diode recovery effects eliminated, the timing of the FET turn off is no longer critical. All FETs are turned off 100 nSec before any FET is turned on.

When a forward conducting FET turns off, the compliment FET on the same half bridge is turned on, leaving the current in the saturable inductors unchanged. When a reverse conducting FET turns off, the compliment FET on the other bridge turns on, forcing the current through the saturable inductors to give the body diode time to recover. The controller insures a 200 nSec minimum pulse to completely reverse the inductor field.

As shown in figure 1, the order the FETs are switched depends on the polarity of the output current. For a positive output current, Q2 and Q4 conduct in the reverse direction, so the sequence is Q1,Q2,Q3,Q4. For negative output current, Q1 and Q3 conduct in the reverse direction, so the sequence is reversed.

Because of the saturable inductors, the FETs turn on with zero current. The only loss is in the hysteresis of the ferrite cores. The ferrites used require around 4 uV-sec to completely reverse the field in the cores, resulting in an 80 nSec delay at 48 volts. The control board provides a 200 nSec minimum on time, which insures that the cores will always be completely reversed. Five TC6-4A11 cores were used with a single turn. The small diameter provides a low saturation current for minimum core loss and maximizes cooling area.

The module must operate smoothly from zero to full voltage at any current. The topology allows for hard switching of the FETs with relatively small losses in the saturable ferrites and snubbers. Using oversized FETs with very low on resistance, and designing the magnetics for low resistance, allowed the module to exceed 98% efficiency at full load.

The module has FETs to modules with shorted components to be isolated from the input and output busses. The module uses two 3 milliohm FETs in parallel for each switch to reduce the conduction losses. The switch FET gates are driven by a photovoltaic devices to eliminate the need for isolated power. For a 10 mA input, they provide 8 volts to the gate. The turn on time is more then 1 millisecond, while a small FET reduces the turn off to less then 100 microseconds.

CONTROL BOARD

The control board for the module uses a Microchip DSP (dsPIC60F12A) for control and communication. A serial 16 bit ADC is used to digitize the output voltage, and a serial 12 bit ADC digitizes the output current. The board is design to operate with switching frequencies of 20 to 100 kHz. For each switching cycle, the complete PID algorithm is run with the data from both ADCs. The DSP uses 4 of its 8 PWM outputs to control the FET gate drivers.

The magnet current will be regulated by an external redundant controllers. They will communicate the desired voltage to the module using the one of the redundant CAN bus interfaces. The modules use the CAN bus to transmit status information back to the current controller. The current regulator can update the desired voltage at up to 720 times per second. This will allow current regulation bandwidth of up to 100 Hz. The CAN bus address filters allows multicast and unicast messages. Multicast messages are used to set the voltage of all modules together, while unicast can enable or disable an individual module.

Figure 3 shows the PID (proportional, integral, derivative) algorithm used by the DSP. The voltage control includes two derivative terms (dV/dT and dI/dT) for increased flexibility. The current control also has a proportional term and two derivative terms. The integrator is shared by both to allow cycle by cycle switching between modes.

The module was tuned to a 1 KHz bandwidth for the voltage loop. The bandwidth was limited by the 1 KHz resonant frequency of the output filter. The step response of the voltage is shown in figure 3.

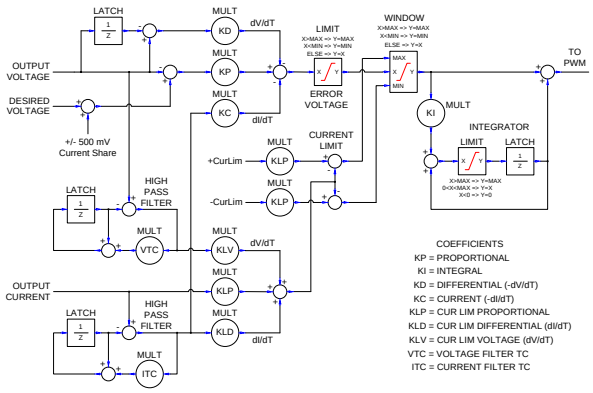

Figure 3: Control Algorithm

The DSP has an internal 12 bit ADC with 16 channel multiplexer. These are used monitor the input and output voltages, and voltage drop on each FET isolation switch. This ADC also reads the module output current, and the average current of all the parallel modules. These two readings are used to balance the output currents between modules by adding a small offset to the desired output voltage.

The DSP outputs a 1 uSec pulse every cycle on a bus shared between all parallel modules. An event capture register records the time of the first pulse received from another module. Another register counts all the pulses received in a cycle. This allows the module to calculate the desired delay time, and adjust its operating frequency to maintain the delay to the next pulse. From 2 to 8 modules can phase shift their switching to minimize the output ripple.

The DSP allows for detection failed components in the module. The input FET switch has a 100 ohm PTC thermistor across it to charge the input capacitors. This allows the module to do a simple self test prior to turning on the output switch. The module is run at 400 millivolts output to check for shorted components. If the unit draws excess current through the thermistor, the unit will abort the turn on. The output FET switch prevents the voltage from appearing on the output terminals of the module.

If the output capacitor of a module shorts, it will cause the output voltage on all modules to drop to near zero and switch to current limit mode. When the DSP enters current limit with an output voltage below 500 millivolts, it will open the output FET switch. It the voltage does not exceed 500 millivolts after 20 milliseconds with the 700 millivolt drop of the FET diode, the unit is disabled. With the faulted module isolated, the remaining modules will enable the output switch and resume normal operation.

The DSP provides a soft turn on and turn off sequence to allow for the addition or removal of modules from an operating system (hot swap) without disturbing the other modules. When a module is enabled, it is set for the same output voltage as the other modules, but with a 1 amp current limit. Once the module voltage has been matched to the other modules, the current limit is ramped up to 40 amps. If a module is disabled, its current limit is ramped down to 1 amp before the module is turned off.

The DSP UART is used for tuning of the digital control loops. The DSP collects 512 data samples from four variables, out of 16 available, and transmits the data for display in a Matlab program. Figure 4 shows the Matlab GUI panel. The program allows any of the control coefficients to be modified and tested. The coefficients are then saved to the EEPROM in the DSP.

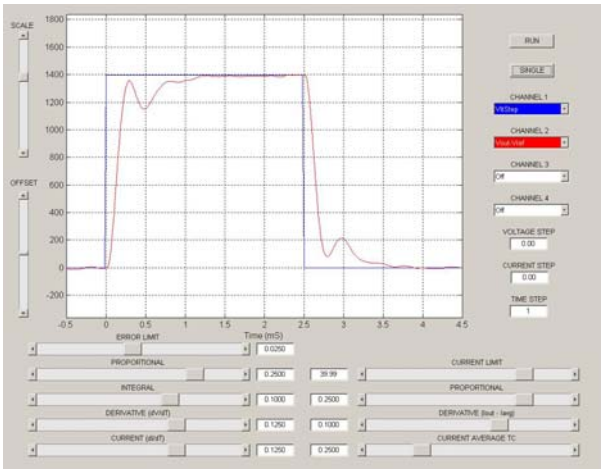

Figure 4: Matlab Tuning Panel

CONCLUSION

A design of scalable power module has been demonstrated. The module provides input and output isolation, while reaching 98% efficiency at full load. The digital control loops provide flexibility to maximize the bandwidth of the module, while switching seamlessly between voltage and current control. The module can autonomously synchronize switching and share current with other parallel modules.

The DSP allows the modules to detect failures and isolate the failure from the system. Special turn on and turn off sequences minimize output voltage transients for hot swap applications.

REFERENCES

[1] D MacNair "DIGITALLY CONTROLLED HIGH AVAILABILITY POWER SUPPLY", LINAC 2008 TUP021