

Use High-Density Power Components for High Availability

Modern COTS systems require a variety of voltages and have varying power needs, but they must also maintain high availability in mission-critical systems. The best choice to achieve all three: COTS power components.

Keith Nardone, Product Marketing Manager,
Military Products, Vicor Corporation

High reliability has long been a priority for military electronic systems. In recent years, with the increasing proliferation of electronic systems in military units in the air, aboard ship and on the ground, the requirement for high system availability is also becoming more widespread. Specifiers of COTS power products are following the lead of many commercial OEMs who have discovered that, for their customers, uninterrupted operation and continuous access to information are becoming mission-critical components of their business operations.

High-density power components (Figure 1) offer intrinsic high efficiency and reliability; nevertheless, in practical, real-life applications, they can fail. Fortunately, when high system availability is a requirement, these modular components can be easily configured into distributed power architectures, fault-tolerant arrays, and hot-

swappable cards—important design approaches to avoid failure of the total power solution.

Taking the Heat: High Efficiency and Reliability

Component power modules, usually built in highly automated, advanced manufacturing facilities, are produced with higher inherent quality than conventional power supplies. Field MTBFs in the tens of millions of hours are not uncommon. Using advanced switching techniques and topologies, power modules generally produce much higher conversion efficiencies than discrete designs, resulting in much less heat being generated.

In spite of high efficiency and reliability to begin with, manufacturers of modular DC-DC converters employ a range of strategies to cope with the unavoidable heat generation. Using synchronous rectification in low voltage units instead of diode rectification, for example, is one popular strategy to improve efficiency. Diode rectification uses a Schottky diode, which has a small resistance and an essentially constant voltage drop. Consequently, the dissipated power is roughly proportional to the current through the diode ($P_{\text{diss}} = EI$).

Synchronous rectification, on the other hand, operates a little differently and introduces additional cost and com-

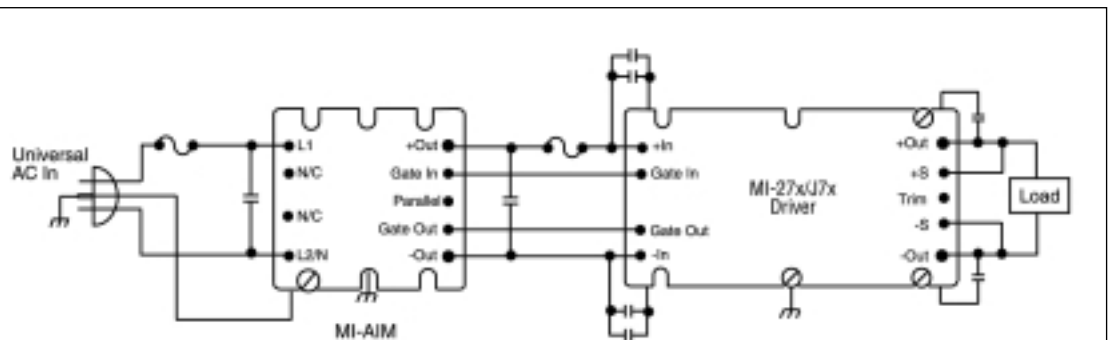


Figure 1

Typical connection diagram of a simple COTS power supply with input module (MI-AIM) interfacing directly with AC mains to provide line rectification, EMI filtering, transient protection and inrush limiting for the DC-DC converter module (MI-27xJ7x).




Figure 2

DC-coupled single-wire interface. All paralleling (PR) pins are directly connected to one another. The interface supports current sharing but is not fault tolerant.

plexity. It employs a MOSFET switch to accomplish rectification; the power dissipated in this case is roughly proportional to the square of the current ($P_{diss} = I_2R$). At lower currents, the MOSFET will generate less heat than the diode and is the better choice.

After a crossover point, perhaps about 20 Amps, diode rectification will generate less heat loss than the MOSFET, and it becomes the better choice. High current synchronously rectified converters overcome this limitation by paralleling several switches to reduce the current through each. This, however, further increases complexity and lowers reliability.

In addition, because thermal management is a continuing concern for DC-DC converter manufacturers, it receives constant attention in the form of incremental design improvements. In some designs, integrated control circuits are molded in epoxy and include a copper plate that is soldered directly to the baseplate. This provides a very low thermal impedance path to the baseplate for improved cooling. Other designs achieve their thermal management objectives with no baseplate and no potting at all.

Other innovations employed to enhance thermal management involve the transformer. One approach uses plated-cavity cores that both confine the

magnetic flux and help conduct heat away from the transformer, thereby increasing the power-handling capability of the power train. Open construction of the transformer can also provide a better thermal path to the baseplate or heatsink.

After all of the design improvements are factored in, however, the generation of heat remains inevitable. Since product reliability and operating life are inversely proportional to operating temperature, design approaches such as distributed power architectures (a more benign environment) can contribute to heat management.

Design Approaches to Achieve High Availability

Beyond mitigation strategies to deal with heat, when high availability is a requirement other approaches become very important, including fault-tolerant arrays (to deal with unavoidable failures), current sharing (to minimize the dynamic response needed) and hot swap (to avoid shutdown).

Distributed Power

As mentioned earlier, a distributed power architecture typically distributes higher voltages at much lower current levels through the system than a centralized system. This architecture reduces power losses during distribution, which results in smaller and lower cost current conductors. Thermal issues are also managed better.

By spreading power conversion around the system, heat is dispersed and hot spots are removed. This can reduce the need for heatsinking or fans. As an added benefit with distributed systems, the level of partitioning is unrestricted: from on board, card-level conversion to one (or more) modules powering several cards. Generally, certain tasks—such as thermal management and hot-swapping—become simpler as the number of load partitions increases, contributing to higher availability.

The lower the current, the lower the I_2R losses and the smaller the conductor size requirements, particularly when the application includes a number of relays, switches and connectors. In high-power applications, such as high-end data pro-

cessing, the savings from lower power losses and smaller conductors can be substantial. In addition, 150 or 300 Volts can be obtained directly from rectification and filtering of the AC line, again at substantial cost savings. Universal (85-264 VAC) input or power factor correcting front ends also produce a 300 Volt output, which is one of the de facto “standards” for intermediate bus voltages.

The requirements of safety standards, however, typically conflict with the selection of a higher intermediate voltage. The SELV (Safety Extra Low Voltage) directive—a commercial requirement in most countries—restricts the voltage to which personnel may be exposed. The maximum SELV is usually 60 Volts. Because of the telecommunications industry’s historic use of 48 Volts, this has also become a “standard” for power distribution, since many power components are available for this voltage.

Fault-Tolerant Power

The primary design approach to achieving high availability is fault-tolerant arrays. The first and foremost requirement for fault tolerance is redundancy; that is, the existence of at least one extra or “redundant” converter in the system. Such a system of converters is commonly referred to as an N+M array,




Figure 3

Using an AC-coupled single-wire interface, all PR pins are connected to a single communications bus through capacitors. This interface supports current sharing and is highly fault tolerant.

where N converters are required to satisfy the power requirements and M additional modules provide redundancy. Fault-tolerant power systems may incorporate some level of redundant front ends and DC-DC converters. Redundancy at the front end can be provided in a number of ways, such as:

- by independent front ends
- by a separate DC input from a battery plus an AC mains connection with appropriate switching in case the AC mains connection is lost.

Redundancy can also be built in at the converter level so that if a failure does occur, there are still sufficient converters available to provide full power to the system. The designer must also decide if the system must be shut down for replacement or be capable of replacement on-line (be hot swappable). This decision rests largely on the level of availability and reliability required. If the unit must have 100 percent availability, hot swap is a requirement. If MTBF is the major issue, then the module does not have to be replaced on-line. Providing a redundant module significantly increases the MTBF of the power system.

An increasing number of features that simplify application in a redundant parallel array are available. The most significant of these is current sharing. All modules in the array must be capable of supplying undisturbed power in the event of shutdown or failure of one module, in spite of the sudden change in load current demanded of each. To satisfy these criteria, it is essential that the individual converters share the load current in order to minimize the dynamic response required of each.

Current Sharing

DC-coupled single-wire paralleling (Figure 2) involves the paralleling of two or more identical modules, each containing intelligence. Internal circuitry actively adjusts the output voltage of each supply so that all units deliver equal currents. This approach supports a level of redundancy, but it is susceptible to single-point failures that can, at best, defeat current sharing and, at worst, destroy every

module in the array. The major reason for this weakness is the single-wire, galvanic connection between modules.

One converter architecture, zero-current switching, uses a unique load share scheme that overcomes these disadvantages. Specifically, this architecture enables an AC-coupled single-wire parallel bus that gives the design engineer a clear edge in power system design, while offering additional functionality.

With this topology, any module can assume control of the array. The module that has the highest output voltage transmits a pulse on the parallel bus to which all other modules on the bus synchronize. Since similar zero-current switching converters transfer the same amount of energy on each switching cycle, synchronized converters will inherently share current. If, due to transient events or a module failure, another module's output voltage becomes higher, it will transparently take command of the array with no perturbation of the output bus.

This capability of synchronous current sharing in democratically-controlled arrays offers power architects new opportunities to achieve simple, non-dissipative current-share control. It provides options that simplify current sharing and eliminate the need to sense the current from each individual module and adjust each control voltage. A pulsed signal also gives designers the option to add capacitors or transformers between parallel pins, providing DC-blocked coupling. Such coupling prevents certain internal or external failure modes from propagating to other modules in the array, thus providing an increased level of fault tolerance. Additional advantages of this architecture include excellent transient response and no loop-within-a-loop control problems.

Using an AC signal at the parallel (PR) pin (a bi-directional port on each module used to transmit and receive information between modules) also provides the system designer with the opportunity to add a level of reliability not previously possible. Instead of simply connecting all PR pins together, the designer has the option to capacitively couple them (Figure 3) and hence avoid the same potential failure inherent in all DC-

coupled, single-wire paralleling schemes: the failure of a single module can affect current sharing or even destroy other modules. Adding a capacitor from each module to the common parallel bus eliminates this failure mode.

Hot Swap

Most applications today that require fault tolerance or redundancy also require hot-swap capability to ensure continuous system operation. Hot-swap cards must be designed to keep any primary referenced potentials (or secondary-side circuitry capable of delivering large amounts of energy) from coming into contact with the user. It is also essential that when a module fails, the failure is detected and identified by an alarm or notice to provide service.

The design must also protect the input voltage and output voltage buses from transients during swap-out. Typically, in the event of a short-circuit fault on the input bus caused by a faulty unit, large capacitors provide the necessary hold-up time until the fuse on the faulty unit has blown. However, these large capacitors can cause the input bus to sag when the replacement card is inserted. This capacitance can be decoupled with a series resistor and a shunt switch (an FET or relay). To minimize efficiency loss, the shunt switch is open on initial plug-in and then closes to short out the series resistor.

Output considerations during hot-swap are similar. ORing diodes, used to protect the bus from an output short in a module, will also isolate the capacitance on the hot-swap card to avoid discharging the same capacitors on the redundant cards already in the system. If ORing diodes are not used on the output, the output capacitance on the hot-swap card will have to be isolated in a similar manner to that used for the input bus. ■■

Vicor
Andover, MA.
(978) 470-2900.
[www.vicr.com].