

Building Block Modules Power Supercomputers

A combination of similarly packaged building block modules provides power for supercomputer voltage regulator modules, processors, memory and 12-V legacy loads.

SUPERCOMPUTERS CAN dissipate up to 25 kW per rack in a data center, where half of the electricity bill covers the supercomputer and the other half is for air conditioning. Therefore, the system design's processing power must be maximized while its operating power is minimized. Minimizing the operating power requires relatively low-power computer components as well as an efficient power management system.

Efficient power management systems are installed in the new Blue Gene/P and Power 6 series of supercomputers developed by IBM. At the U.S. Department of Energy (DOE) Argonne National Laboratory, the IBM Blue Gene/P high-

performance computing system is now the fastest supercomputer in the world for open science use. At the same time, it is highly power efficient, with up to 371 Mflop/s/W.

Chosen to provide the power management subsystem were Vicor Corp.'s V•I Chip power component building blocks, which can scale up for high power with high distribution and conversion efficiency. In addition, the V•I chip system is flexible enough to enable lower power systems to use the energy-saving benefits and maximize processing performance per rack.

To achieve high efficiency, the V•I chip approach separates regulation and voltage transformation, resulting in:

- *Reduced power distribution losses*
- *A reduction in duplicated functions in the power conversion path*
- *Reduced power dissipation at the point-of-load while increasing total system efficiency.*

One of the Vicor building blocks is the BCM™ bus converter, which provides isolated, unregulated power using a voltage transformer/current multiplier (Fig. 1). It uses a Sine


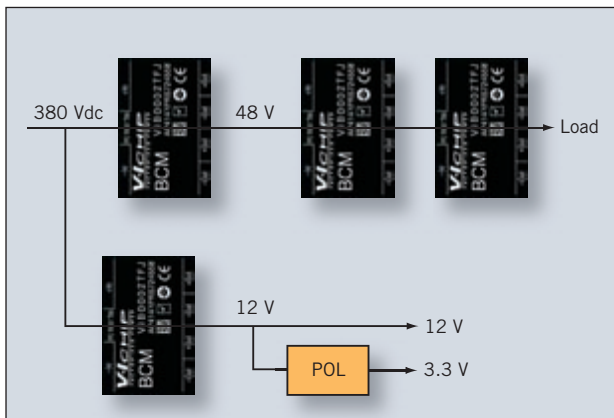


Fig. 1. Sine Amplitude Converter (SAC) employed in the BCM bus converter provides isolated, unregulated power using a voltage transformer/current multiplier.

Fig. 2. Using BCMs to generate 48-V rails for main powertrain and 12 V for legacy loads and point-of-load (POL) converters.



Amplitude Converter (SAC™) topology with zero-voltage switching (ZVS) and zero-current switching (ZCS). With 384 Vdc applied, it can produce either 12-V output at 300 W and 95.5% efficiency or 48-V output at 330 W and 96.5% efficiency (Fig. 2). The BCM is housed in a 1.1-in.² package (1.28 in. ×

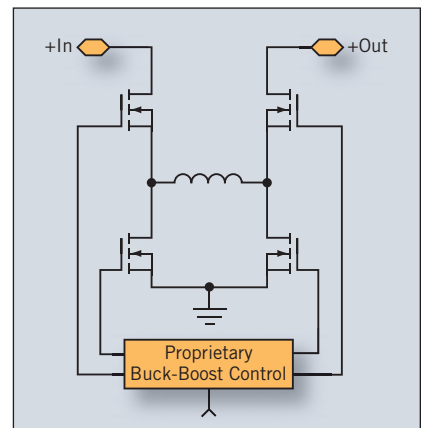


Fig. 3. PRM buck-boost converter accepts an unregulated 48-V input and produces a regulated 26 V to 55 V that is applied to the downstream VTM.




Fig. 4. VTM SAC supplies an isolated, regulated voltage transformation direct to the load.



Fig. 5. An array of BCMs in IBM's p575 provides 5.1 kW at 11 V for the VRMs in the computer.




Fig. 6. Efficiency vs. output power shows the scalability of modules, which adds to the flexibility of the building blocks.

0.87 in. × 0.26 in.) that has a power density capability of more than 1,000 W/in.³.

Another building block is the PRM™ nonisolated regulator with ZVS buck-boost topology that switches in excess of 1 MHz (Fig. 3). It provides 320 W in a 1.1-in.² package, has a power density of 1,100 W/in.³ and is 97% efficient at 320-W out.

The third module is the VTM™, an isolated voltage transformer that employs the SAC topology with ZVS and ZCS and switches at more than 1 MHz (Fig. 4). It can handle up to 100 A in a 1.1-in.² package and is 96% efficient at 300-W output. Furthermore, it accepts a regulated 26 V to 55 V from the PRM and is capable of producing 0.8 V to 55 V at up to 100-A load.

As seen in Fig. 5, 17 paralleled

BCMs rated at 300 W with 350-Vdc input deliver 5.1 kW at 11 V and 96% efficiency in the “p575.” The BCM outputs power downstream voltage regulator modules (VRMs) in the supercomputer.

The efficiency versus output power graph of Fig. 6 shows the scalability of parallel arrays of similar building blocks. The “eco-array” senses load demand and enables BCMs as required. This optimizes the low-load efficiency of parallel arrays of modules.

In addition, the PRM and VTM combination gets into the act, as shown in Fig. 7. One PRM and two paralleled VTMs accept 48 V_{IN} and provide 1.2 V_{OUT} at 200 A for the massively paralleled power PC processors in the Blue Gene P. A similar PRM and two paralleled VTM combination accepts 48 V_{IN} and produces 1.8 V_{OUT} at 160 A for the system memory.

Optimizing efficiency and power density in mainframe systems oftentimes comes with a trade-off in design flexibility. A 380-Vdc architecture for “big iron” with loads of hundreds of amps at less than 1 V is unlikely to be optimal powering plug-and-play rack-mount server motherboards. V•I chips can scale up for high power with high distribution and conversion efficiency. However, their flexibility enables lower power systems to utilize the energy-saving benefits and maximize processing performance per rack. ⏻




Fig. 7. One PRM and two paralleled VTMs provide 1 V at 200 A for the Blue Gene P processor and 1.8 V at 160 A for the memory arrays.