

INDUSTRY STANDARDS LEAD PUSH TOWARD ENERGY-EFFICIENT COMPUTING

ENVIRONMENTAL CONCERNS AND RISING ENERGY COSTS ARE SPURRING INDUSTRY AND GOVERNMENT GROUPS TO DEVELOP REQUIREMENTS FOR HIGH-EFFICIENCY AC AND DC POWER CONVERSION, LEADING TO ENERGY-EFFICIENT SERVERS. MEETING THE NEWEST SPECIFICATIONS WILL DEMAND KNOWLEDGE OF COMPETING POWER-CONVERSION TOPOLOGIES, COMPONENTS, AND DESIGN.

BY LEE HARRISON • PERITUS POWER

In addition to environmental concerns, the increasing cost of electricity is driving data-center managers to more energy-efficient installations. As utility bills become the primary expense for data centers, electricity costs now outweigh real-estate costs, with power consumption per data center ranging from 2 to 22 MW. In 2007, the Internet accounted for 9.4% of total US electricity consumption and 5.3% of global electricity consumption. Networking equipment, such as modems, routers, hubs, and switches, accounted for about 25% of the electricity demand in an average office. If the computers and servers in an infrastructure require 200 kW, then the networking components in that infrastructure

need 50 kW. In addition, 45% of the power a data center consumes is for air-conditioning and cooling. In modern data centers, performance per watt has become more critical than performance per processor.

Every year, computers and servers demand higher performance. Server and PC manufacturers respond with larger and faster disk drives, faster memory, multicore processors, and more I/O devices. Although this approach satisfies customer demand, the new designs often require more power and additional cooling, continuing the trend toward higher power consumption. The average power consumption per server has increased from 150 to 250W in 2000 to

450 to 800W in 2009, and the average power per rack of servers in 2000 was nearly 1 kW, rising to 6 to 8 kW in 2006, and should top 20 kW in 2010.

Environmental concerns are not the only ones behind the push to energy efficiency. Electronic equipment that runs cooler has lower failure rates and increased reliability. Data centers save money when servers require less power because less power means decreased cooling and air-conditioning costs. Large data farms benefit from relatively small reductions in electricity cost. Google, for example, may be able to save nearly \$1 million a year in utility bills by reducing power consumption by 2 to 3%. Some

industry analysts believe that financial benefits are more important than climate issues; in a time of reducing costs across all businesses, this belief may be one factor in the drive toward energy efficiency.

The ac/dc power-conversion step in the overall power chain for server farms can yield some of the most significant gains in power efficiency. Many industry and government groups, including 80 Plus, EPRI (Electric Power Research Institute), the CSCI (Climate Savers Computing Initiative), and the US Environmental Protection Agency, have





developed requirements for high-efficiency ac/dc power supplies. **Table 1** shows the current targets that the CSCI and Energy Star define for single-output power supplies. **Table 2** shows multiple-output supplies (**Reference 1**).

The Silver category is challenging for some power vendors, but the Gold and next-generation Platinum levels are more difficult to reach. Some power-supply vendors have achieved the Gold standard, and some servers are currently available with Gold power supplies, but the Gold standard is unlikely to become well-established until the end of 2010. To meet those standards, manufacturers have taken different approaches, including using interleaved PFC (power-factor control), bridgeless PFC, and resonant topologies.

INTERLEAVED PFC

The power industry has long used interleaving techniques in multiphase-buck-converter designs. This design commonly meets fast load-transient demands for DSPs and other processors. Some designs interleave multiple synchronous power stages to increase power delivery to the load and decrease input and output capacitance, maximizing the benefits of smaller output inductors in each phase of the design. Due to the ripple-current cancellation effects at the output capacitors, this topology can realize transient demands in excess of 350A/ μ sec.

Interleaving PFC boost stages benefits from the same principles. PFC boost encompasses CCM (continuous conduction mode), DCM (discontinuous conduction mode), and CRM (critical conduction mode). High-efficiency designs are likely to use CCM in the server area because this mode of operation suits power supplies with more than 350W of output power. Although CCM eases the design of EMI (electromagnetic-interference) filters, the approach generally requires a larger boost inductor than do alternative design techniques.

AT A GLANCE

▣ The ac/dc power-conversion step in the overall power chain for server farms can yield some of the most significant gains in power efficiency.

▣ To meet industry standards, manufacturers have taken different approaches, including using interleaved PFC (power-factor control), bridgeless PFC, and resonant topologies.

▣ Thanks to its 0V switching losses, higher switching frequencies, and smaller footprints, resonant-converter topology may be able to achieve Energy Star Platinum standards.

▣ For the near future, silicon will remain the dominant switching semiconductor material, and gallium nitride will start to make inroads over the next year.

Interleaved CCM is not without its own risks: The design suffers from inherently higher switching and reverse-recovery losses in the rectifier. You can partially overcome these losses, however, by using silicon-carbide diodes. In addition, you cannot typically adjust the frequency in CCM designs because CCM operates as an average-current-mode, PWM (pulse-width-modulated)-control, fixed-frequency design, forcing average input current to be proportional to the rectified ac line cycle.

Energy Star specifications are the primary reason for the move to interleaved PFC. As enclosures and footprints for power become smaller, however, the better EMI performance helps reduce the size of EMI filters. Power density is excellent for interleaved PFC designs, but circuit design can be difficult; the design challenges include its many modes of operation, the need to guarantee safe operation at start-up and shut-down, autoranging of wide-input sources, and fault conditions.

To achieve high efficiency at loads less than 15%, particularly around the 5 to

10% range, you will likely need some form of phase shedding, pulse skipping, or burst modes. These approaches can be fully autonomous or selectable through software and DSPs, depending on load. No matter which approach you choose, problems arise because other functions can be occurring simultaneously. Although using pulse skipping and burst modes lets you attain increased efficiency at light loads, they wreak havoc when you are trying to accurately measure PFC, and they can cause EMI issues and strange harmonic effects in racks and across servers that connect to the same ac source supply.

Measuring PFC at light loads requires the use of a line-impedance-stabilization network and a capacitor between the ac source and the power supply. Investigations are ongoing that aim to reveal the effects of measuring PFC under these conditions. Determining the reason for strange harmonic effects in high-efficiency power conversion also requires further investigation. Thus, efficiency comes at a price. You may be able to improve efficiency, but doing so introduces new side effects.

Implementing PFC in burst or skipping modes also causes problems for the digital-sensing input-voltage, -power, and -current outputs. Although fully digitally controlled power supplies can likely cope with these problems, it's difficult to obtain valid data from a control signal that is constantly changing. Many designs turn off power reporting at these lighter load conditions, but Energy Star is determined to tighten the already difficult-to-meet specifications, so designers must find approaches for reporting data at light loads. Digital techniques in PFC have recently produced impressive performance at light loads but at the expense of THD (total harmonic distortion).

BRIDGELESS PFC

Bridgeless PFC offers loss savings of 0.4 to 1.5W at high-load conditions,

TABLE 1 TARGETS FOR SINGLE-OUTPUT POWER SUPPLIES

Loading (%)	CSCI Silver and Energy Star Version 1			CSCI Gold and Energy Star Version 1			CSCI Platinum		
	Efficiency (%)	PFC ($\leq 1000W$)	PFC ($> 1000W$)	Efficiency (%)	PFC ($\leq 1000W$)	PFC ($> 1000W$)	Efficiency (%)	PF ($\leq 1000W$)	PF ($> 1000W$)
10	75	0.65	0.8	80	0.65	0.8	82	0.65	0.8
20	85	0.8	0	88	0.8	0.9	90	0.8	0.9
50	89	0.9	0.9	92	0.9	0.9	94	0.9	0.9
100	85	0.95	0.95	88	0.95	0.95	91	0.95	0.95

TABLE 2 TARGETS FOR MULTIPLE-OUTPUT POWER SUPPLIES

Loading (%)	CSCI Bronze		CSCI Silver		CSCI Gold	
	Efficiency (%)	PFC	Efficiency (%)	PFC	Efficiency (%)	PFC
20	82	0.8	85	0.8	87	0.8
50	85	0.9	88	0.9	90	0.9
100	82	0.95	85	0.95	87	0.95

depending on output power. The benefits of a low-power-line mode are much greater and suit countries that still run servers in low line conditions. Standard bridge rectifiers suffer in this application because they use series-connected semiconductors, but the approach remains a low-cost and reliable choice for ac rectification. In some designs, standard bridge rectifiers run at dangerously high temperatures. Airflow is typically not ideal in the location of a 1U server power supply, and some customers of power-semiconductor manufacturers have requested an increase in thermal capabilities to 175°C from the previous limit of 150°C.

Bridgeless approaches have not become mainstream designs due to their higher cost and challenges, including complex gate-drive circuits and the

difference between the input and the output ground references. Prototypes of bridgeless designs have shown high-EMI common-mode switching elements at the negative bulk-capacitor connection, and, with virtually no low-frequency path back to the ac source, EMI is greater. Various semiconductor companies, including STMicroelectronics and On Semiconductor, now offer controllers for bridgeless rectification, although some operate only in high line conditions. These units operate at fixed frequencies of approximately 100 to 150 kHz, but the \$20 to \$60 cost to implement bridgeless front ends is too high for a 1W power saving.

For now, CCM designs with carefully used silicon and light-load power-saving techniques have sidelined bridgeless

designs. Unless there is a major breakthrough in the cost, complexity, and confidence of bridgeless technology, this scenario may remain so for some time to come. In addition, if dc-based power-distribution infrastructures for data farms become popular, ac/dc bridgeless technology may become obsolete before it gets off the ground.

RESONANT TOPOLOGIES

Thanks to its 0V switching losses, higher switching frequencies, and smaller footprints, resonant-converter topology may be able to achieve Energy Star Platinum standards. Moreover, parasitics, which are problematic for alternative topologies, become advantageous in resonant topologies (Table 3). PWM designs lose power at high switching frequency, and resonant-mode topology begins where PWM topologies end. There are many options available in waveform shaping to eliminate switching losses, including ZCS (zero-current switching); ZVS (zero-voltage switching); and quiresonance, which uses only a part of the sinusoidal waveform.

Power designers commonly choose

TABLE 3 COMPARISON OF RESONANT POWER SUPPLIES

Converter	Advantages	Disadvantages
Series resonant	Behaves as a current source and best suited to high-voltage, high-power designs Provides good handling of overload and fault conditions Requires small or no output filter	Poor regulation at light- and no-load conditions Continuous-current mode below resonance causes high component stress and premature failure if not carefully designed
Parallel resonant	Behaves as a voltage source and suitable for low-voltage outputs, good no- and light-load regulation Discontinuous mode of operation is similar to standard PWM buck conversion	Requires snubbers, consuming efficiency Does not self-protect in overload or short-circuit conditions, requiring additional protection Constant resonant circulating current at all load conditions can cause problems
Zero-current switching, quiresonant, fixed on-time	Low turn-off losses in silicon Recycles any remaining leakage current	Power components contain high peak currents
Zero-voltage switching, quiresonant, fixed off-time	Virtually eliminates switching losses Recycles any remaining leakage current	Operates above resonant range Power devices require at least three times the supply voltage; multiresonant techniques offset this problem
Series-parallel resonant	Excellent performance at light-load, no-load, and short-circuit or fault conditions	Produces parasitic elements that can be troublesome, requiring additional passive components
Multiresonant, zero-voltage-switching Class E	Virtually eliminates switching losses, virtually no frequency shift compared with other methods	Can cause stability issues and generates high peak voltage and current across semiconductors
Clamped PWM	Simple control loop with low-voltage stress Current-mode capable and fixed frequency of operation	Turn-on losses can be high; operates best at medium- to full-load conditions
PWM, zero-voltage switching	Fixed-frequency design with maintained switching losses at low levels across all power devices High efficiency at high switching frequency	Creates problematic switching noise at light loads and does not achieve zero-voltage switching in the light-load mode of operation
LLC resonant	High efficiency with excellent line/load regulation Outperforms series-resonant mode	Performance is superior to other modes of operation, but design is critical Possible N+1-redundancy concerns



LLC (inductor-inductor-capacitor) converters to implement resonant topologies. These devices offer good regulation and frequency control, and they boast efficiency as great as 97% for the LLC stage, further assisting compliance with the Platinum standard. LLC topologies also have reduced inductance in the primary transformer, excellent frequency control over a wide load range, excellent line and load regulation down to zero loads, and ZCS control of secondary rectifiers.

MATERIAL AND MAGNETICS

For the near future, silicon will remain the dominant switching semiconductor material, and gallium nitride will start to make inroads over the next year (**Reference 2**). Gallium-nitride promises a tenfold reduction in on-resistance for 50V devices by 2013, but no more than 10% of applications will be widely implementing it before 2015. Using planar transformer technology and integrated magnetics with power components as part of the transformer assembly can achieve 0.5 to 2W savings. Planar magnetics offer lower switching and cop-

per losses. They also offer a low profile, high power density, and high frequencies, so designers can build them into the main PCB (printed-circuit board) or assemble them as subassemblies for daughtercards.

Integrated magnetics can provide a 50% smaller footprint than you can achieve using separate magnetic and rectification stages. Thanks to lower commutation loops, the approach also offers lower, better-controlled high-frequency power-supply noise spikes and EMI, and reduced copper losses guarantee increased efficiency. Generally, integrated magnetic designs offer repeatability and cost reductions in manufacturing, and the tolerances of magnetic characteristics and output stages show closer results to each other than those of discrete stages. This technology will find its way into future high-efficiency designs; the option of increased power density alone is a big enough reason to adopt the technology. Efficiency gains may even become secondary to this requirement.

The magnetics industry has seen fewer improvements in technology than the semiconductor industry. Passive

components account for more than 75% of power-supply real estate; you must target these components when looking to reduce your design's size and increase its power density. Magnetic components, including EMI filters, account for the biggest real-estate consumer, and bulk-capacitor electrolytics follow closely.

DIGITAL CONTROL

Digitally controlled power supplies have seen improvements in efficiency due to improvements in digital control and optimization and sensor accuracy for voltage and current reporting from the power supply. The recently released Energy Star requirement requires $\pm 10\%$ accuracy with a cutoff at $\pm 10W$ —that is, the accuracy need never be better than $\pm 10W$. The Tier 2 specification, which will come out in October 2010, will require $\pm 5\%$ accuracy with a cutoff at $\pm 5W$, making this requirement one of the most difficult to meet. The voltage-reporting accuracy is relatively easy, but the current-reporting accuracy is difficult because the sensing component must have an extremely low value to

avoid consuming additional power and thus lowering efficiency. Further, at low loads, the current waveform for determining the current reporting is too distorted to be meaningful. You typically use the PFC current to sense input current. SNRs (signal-to-noise ratios) can be poor, and it is not unusual to need signal amplifiers to improve the signal. However, in digitally controlled power supplies, the DSP can calibrate the current-sensing element. You typically use a rectifier separate from the main power rail to sense input voltage. A primary-side DSP using power-correction factors can sense peak or average voltages and report the status. The accuracy you can achieve with this technique ranges from 1 to 5%.

FOR MORE INFORMATION

80 Plus
www.80plus.org

Climate Savers Computing Initiative
www.climatesaverscomputing.org

Electric Power Research Institute
www.epri.com

Energy Star
www.energystar.gov

Environmental Protection Agency
www.epa.gov

On Semiconductor
www.onsemi.com

STMicroelectronics
www.st.com

You can report input power as a function of input current and input voltage. Again, designers typically use calibration to prevent errors, and you must build the aging of components into this calibration. Without digital filtering and carefully controlled algorithms, these requirements would be almost impossible to meet. Even with today's technology, it is a difficult task.

Digital control assists not only with the latest energy-sensor-reporting requirements but also with the control of power conversion and optimization of efficiency. It allows direct control over phase control, burst or skipping modes, switching-frequency control, transient response, voltage and current limits, dead times, and sequencing of internal waveforms and voltages. Optimization of switching frequencies is a complex task, which becomes a four-way juggling act involving power-supply performance, efficiency, magnetic-component footprint sizes, and power-supply-enclosure limitations. **EDN**

REFERENCES

1 Climate Savers Computing Tech

Specs, www.climatesaverscomputing.org/tech-specs.

2 Brier, Michael A, "GaN Based Power Devices: Cost-Effective Revolutionary Performance," International Rectifier, 2008, www.irf.com/pressroom/articles/560PEE0811.pdf.

AUTHOR'S BIOGRAPHY

Lee Harrison is director of Peritus Power (www.perituspower.com). Previously, he worked at Sun Microsystems as a power-system architect and technical leader for the Sparc and x86 platforms, developing the technology and strategy for power conversion with Emerson Network Power, Delta, Lineage, Power One, and FDK from 2000 to 2009. He provides input to the Environmental Protection Agency on power-related issues and has been a voting member of the Climate Savers ac/dc work group for the last four years. Before joining Sun, Harrison was an engineering manager for a UK-based defense power-supply company, specializing in high-density, low-profile dc/dc and ac/dc power conversion and nuclear-protected electronics. His LinkedIn profile is located at www.linkedin.com/pub/lee-harrison/0/382/569.