ISSN: 1863-5598

ZKZ 64717
01-11

Bodo's Power Systems®

Electronics in Motion and Conversion

January 2011

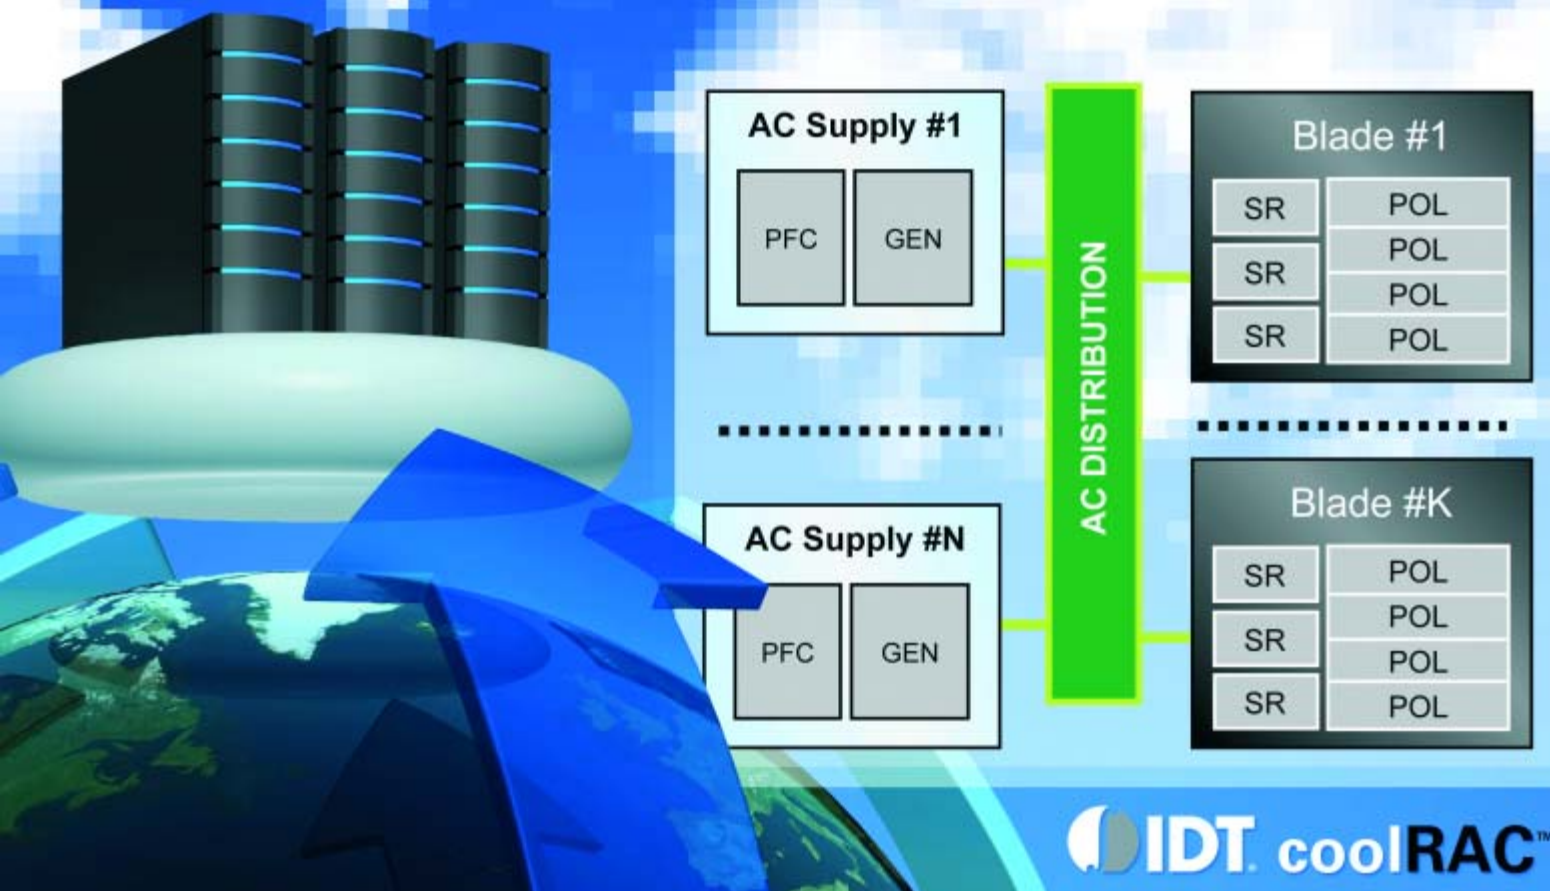

IDT coolRAC™

Bodo´s Power Systems®

Energy Efficiency Becomes the Focus for Data Centre Servers

This architecture is also favorable for future improvements

The significant growth in power consumption for information technology (IT) data centres has created an excellent opportunity for companies to develop products focused on energy efficiency.

By Arman Naghavi, Vice President and General Manager of the Analog & Power Division, Integrated Device Technology, Inc.

Data centre power consumption is growing at an annual rate of 16%. Driven by increasing demand for cloud computing, social networking, multi-media, and utility computing services, government and commercial organizations will spend nearly $8 billion (US) per year to power and cool data centres. The majority of power is consumed by server systems and cooling.

The traditional power architectures of these servers are reaching the technological limits of energy efficiency. Some efficiency improvement approaches can require a redesign of the entire data centre power infrastructure to reduce power consumption; this solution comes at a prohibitive cost. Data centre building and infrastructure redesign can easily cost tens of millions of dollars.

Companies, such as IDT, have approached the problem by focusing on the total power efficiency of the server system, without requiring any significant data centre infrastructure or building modifications. The result is IDT's high efficiency server power architecture, known as coolRAC (RAC is an acronym for Resonant Alternating Current). This technology allows data centres to realise significant reductions in server power consumption.

Total server efficiency is measured from the AC input, across the server rack and to the processors on the motherboard. Typical server platform architectures have a total of 80% power conversion efficiency and often use high-cost components. This means that approximately 20% of the power consumed is simply lost as thermal waste. IDT coolRAC technology can improve this conversion to over 90% efficiency, reducing wasteful power losses and the associated cooling requirements. Since coolRAC utilizes the existing data centre electrical infrastructure, this technology can be easily retrofitted into existing centres without expensive building improvements.

Data centre power usage is increasingly the focus of facility managers searching for ways to increase processing capabilities and at the same time, reduce costs and become "green" by minimizing their carbon footprints. Rack-mounted computing servers, with multiple processors and minimal down-times, typically account for over 70% of the data centre's energy consumption. For example, a large centre equipped with 50,000 servers can consume upwards of 12MW of power. By improving the efficiency of the servers using IDT's coolRAC technology, direct power savings of over 1MW are achievable. At a conservatively estimated $0.10/kWh rate, this equates to over $800,000 a year in direct power cost savings. Furthermore, it facilitates the expansion of the data centre's computing capabilities.

IDT's coolRAC solution utilises standard input power (3-Phase 220/480V AC) and uses high frequency alternating current (HF AC) distribution. Below is a discussion of the advantages and challenges of this technology along with details on power distribution efficiency improvements and comparisons

Architecture	Efficiency per stage				
	PSU	Backplane	Intermediate Bus Converter	POL	TOTAL
12V DC	92.0%	97.0%	n/a	89.0%	79.4%
48V DC with isolated POL	93.0%	99.0%	n/a	87.0%	80.1%
48V DC with Intermediate Bus Converter	93.0%	99.0%	97.0%	89.0%	79.5%

Table 1: Efficiency per stage

Architecture	Losses % per stage				
	PSU	Backplane	Intermediate Bus Converter	POL	TOTAL
12V DC	38.9%	13.4%	n/a	47.7%	100.0%
48V DC with isolated POL	35.2%	4.7%	n/a	60.1%	100.0%
48V DC with Intermediate Bus Converter	34.1%	4.5%	13.5%	47.9%	100.0%

Table 2: Losses % per stage

with existing technologies. Test results from a prototype demonstration are provided, and the next steps in the development process are outlined.

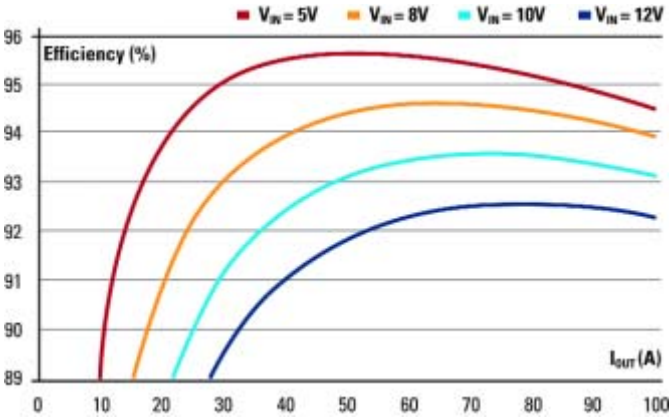

Figure 1: POL buck converter efficiency with different Vin

High frequency AC distribution coolRAC concept

To make the most impact in overall system efficiency, the first step is to address the efficiency of point-of-load (POL) converters. In today's architectures, buck converters have proved to be the simplest, most efficient, and the most cost-effective POL converter topology. However, modern digital IC loads require the POL converter to supply an output voltage approaching 1V DC. As a result, a buck converter with a 12V DC input voltage will have a high conversion ratio and thus will be unable to achieve high efficiency for a 1V DC output. The importance of improving POL efficiency is shown in Tables 1 and 2 with comparisons of efficiency and power losses for different architectures at specific stages.

To decrease the conversion ratio and therefore increase POL efficiency, the POL input bus voltage must be decreased. Figure 1 shows dramatic POL efficiency improvement with no other changes, except for decreasing the input bus voltage in steps from 12V down to 5V.

The efficiency improvement could be further realised by using silicon switches with a lower voltage rating. Lower-rated devices would provide even higher efficiency than shown above in Figure 1 at similar costs. This is due to the lower price per silicon area or, in other words, a lower drain-to-source on-resistance (Rdson) per silicon area. Likewise, depending on the application, a customer may choose to save cost of silicon where high efficiency is not critical and still realise a significant efficiency improvement compared to a 12V input bus.

Considering the industry trend toward lowering input supply voltages for modern digital ICs like CPUs, GPUs, ASICs and memory, reducing POL input voltage has become imperative for highly efficient POL converters. For today's industry applications, a 5V DC POL input bus voltage appears to be the optimal choice for server systems. This low input voltage significantly reduces the voltage conversion ratio and increases power efficiency. Also, 5V rated power components are readily available in the market today. Peak efficiency of 96% for a 5V DC-to-1.2V DC buck converter shown in Figure 1 was demonstrated with 30V rated silicon switches and could be further improved by the use of 12V rated devices.

However, delivering input power to the POL at lower voltages increases losses in distribution media, making a 5V DC distribution system impractical in most cases. Sophisticated backplane layouts including expensive connectors, backplane PCBs and busbars are required to achieve a modest 2% power loss for a traditional 12V DC distribution architecture for high-end server systems. Changing the distribution voltage to 5V will increase power losses more than five times.

IDT's coolRAC technology delivers low voltage to the POL in the most energy efficient way, solving a fundamental conflict between maximizing efficiency for the POL converter (dictating LOW voltage distribution) and minimizing losses in the power distribution system (dictating HIGH voltage distribution).

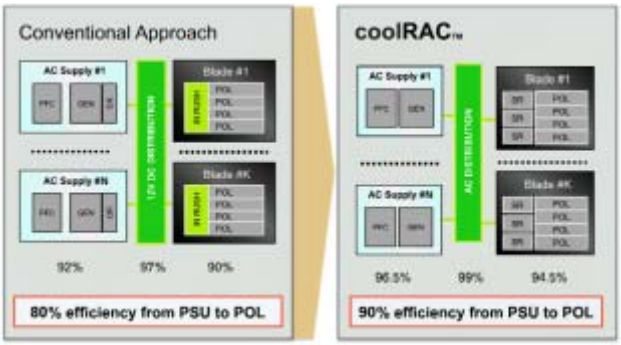

Figure 2: Traditional 12V DC distribution architecture and new proposed High Frequency AC distribution architecture, where: PFC – Power Factor Correction stage, GEN – high frequency AC Generator, SR – Synchronous Rectifier, POL – Point of Load Converter, OR – OR-ing circuitry, and INRUSH - Inrush circuitry.

Figure 2 shows the comparison between the most popular 12V DC distribution architecture and the IDT coolRAC solution. A redundant system with N front-end power supply units (PSUs) powering K computing units is shown.

Clearly, both architectures have the same number of power conversion stages. By simply moving synchronous rectifiers (SRs) from the front-end PSUs to the blade motherboards allows the use of a reasonably high distribution voltage (for example, 50V peak-to-peak). This minimizes losses in the distribution media. The magnitude of the AC distribution voltage defined by the transformer's turns ratio is very flexible and allows an optimal and scalable solution for server systems with different power levels.

Evidently, the redundant HF AC distribution architecture raises additional challenges for the designer compared to traditional DC distribution systems; for example, the generators' synchronization and current share that is required for features such as hot swappable power supply redundancy. This is the main reason why the concept of distributing HF AC was not widely accepted by the market. Some server vendors have released HF AC distribution on server platforms, but have not utilised the redundancy architecture.

The advantages of HF AC distribution topology have been understood by experts for many years, yet it has not been feasible to implement due to challenges in controlling multiple AC generators in parallel and implementing an efficient and reliable control algorithm for synchronous rectifiers. However, recent advances in analog to digital conversion and digital control techniques, such as those achieved by IDT, are now making this topology attractive to the market.

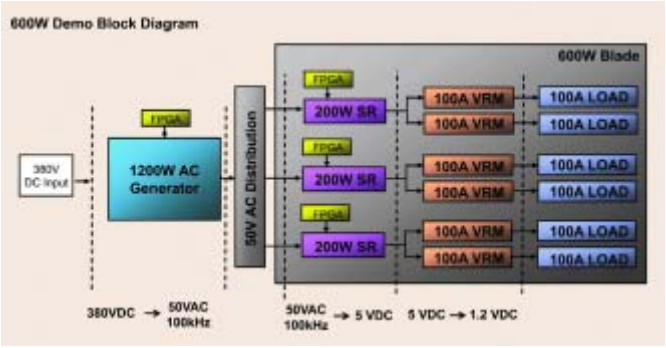

Figure 3: "Proof of Efficiency" 600W Demo Block Diagram

The additional benefits of HF AC distribution compared to DC distribution, based on fundamental differences between AC and DC currents, reveal a possibility to use reactive power components for regulation in addition to active power dissipative components. By using resonant topologies in the coolRAC architecture described, there is a possibility to provide hot plug and redundancy capability without the need for OR-ing and inrush control circuitries.

Of course, all these advantages could be realised only after developing new digital and analog control algorithms and implementing them into new control ICs to make this new HF AC distribution architecture feasible and cost effective.

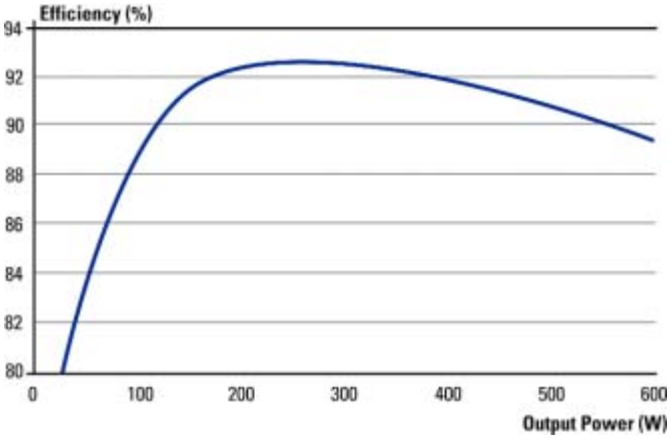

Figure 4: Demonstrated HF AC distribution system efficiency

Proof of efficiency demonstration test results

To prove the architecture's feasibility and efficiency improvement, IDT built and successfully evaluated an experimental test vehicle comprised of one 1200W AC Generator, three 200W SRs, and six 100W POLs. The test vehicle block diagram and demonstrated efficiency results are shown in Figure 3 and Figure 4 respectively.

Efficiency was measured from the 380V DC input to the 1.2V DC POL output, taking the power factor correction (PFC) efficiency out of consideration for the proof of efficiency study. A 1200W AC Generator and 200W SRs were controlled using digital algorithms implemented with field-programmable gate array (FPGA) daughter cards attached to each unit (shown as FPGA in Figure 3). In the future, each control card will be replaced by IDT-developed proprietary ICs. The targeted total system efficiency goal was set to 90% peak (as calculated above in Figure 2). Assuming that the PFC is 98% efficient, the headroom for the rest of the system efficiency goal was calculated to be 92% from the 380V DC PFC output to the 1.2V DC POL output voltage.

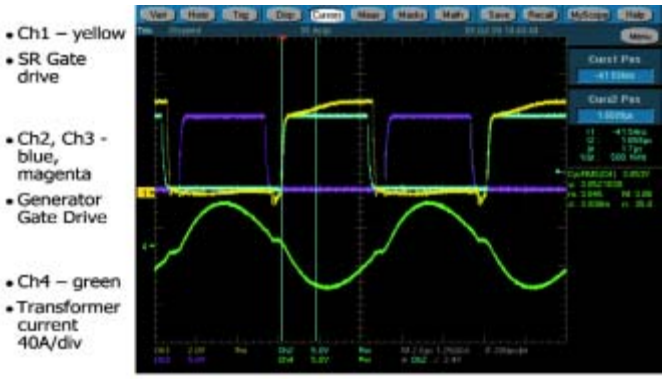

Figure 5: HF AC waveforms at full load

Measured peak efficiency was 92.5%; which exceeded the expectation by 0.5% and lead to the conclusion that the entire HF AC distribution system peak efficiency of 90.5% is achievable with the IDT's technology. Compared to traditional 12V DC distribution, IDT's CoolRAC architecture can deliver an efficiency improvement of more than 10% (See Table 3)

Electromagnetic interference between HF AC power distribution and computing signals was not the subject of the study, but observed current and voltage waveforms showed very smooth edges because of the resonant properties of the proposed energy transfer architecture. As shown in Figure 5, AC bus current waveforms are actually sinusoidal for a wide range of loads and don't contain spectrums that interfere with computing signals. Of course, possible EMI issues should be investigated further during a system-level signal integrity study.

Current development

Based on results demonstrated by the proof of efficiency demonstration, IDT is working on the next step in the development; a 3kW x N Redundant HF AC power distribution server rack. Targeted parameters are shown below:

- AC power supply unit (PSU) - PFC + Generator
- Input: 480V AC - 3-phase 4 wire input
- Output: 3000W ~100kHz 50V AC
- Power Factor Correction: 0.99 PF min
- Efficiency: 97.5% efficiency at full load
- Paralleling: Capable of operating in N+1 redundant operation
- Sharing: at least 10%
- Form Factor: length: 15-18", width: 4.0", height: 2.0"
- Cooling: Forced-air

- Synchronous rectifiers (SRs):
- Input: 50V AC
- Output: 3 options: 100W 5V DC, 200W 5V DC, 300W 5V DC
- Efficiency: 98% efficient
- Cooling: Forced-air

Architecture	Efficiency per stage				
	SBox	Backplane	IBC	POL	TOTAL
12V DC legacy distribution	92.0%	97.0%	n/a	90.0%	80.0%
CoolRAC 100KHz 50V p/p HF AC distribution	96.7%	99.0%	98.5%	96.0%	90.5%

Table 3: Efficiency per stage

The following tasks will be accomplished for this demonstration:

- Control of resonant synchronous rectifier (SR)
  The commonly used self-driven SR control scheme doesn't allow for the best efficiency. There are two issues that need to be solved: reliable gate drive synchronization with AC input voltage, and preventing energy transfer from the output capacitor to the AC bus.
- AC generators synchronization
  To enable paralleling of AC generators, precise synchronization of their clocks is required, desirably, without additional signal lines.
- AC generators current sharing
  Generators' quasi-square wave outputs mismatch could result in significant equalizing currents. Potential loss of efficiency or failures is possible. An effective and robust control system is needed.
- PFC stage control with fast transient response and minimized ripple
  Traditional dual loop PFC control has fundamental conflict between feedback speed and power factor. Because the coolRAC approach has an unregulated second stage, traditional control doesn't provide optimal performance. A control scheme with fast load transient response without power factor degradation is required.

Conclusion

Data centre electrical power requirements continue to increase to meet the growing demand for computing power. Since the server is the largest portion of data centre energy consumption, an increase in server efficiency would make the most impact to save power, reduce cooling requirements and lower costs.

IDT's coolRAC intermediate bus architecture provides a server power solution that dramatically increases efficiency and reduces thermal waste, thus cutting power usage without a large scale data centre redesign. This solution achieves 90% total system efficiency from the power source to the 1.2V DC load while utilising currently available power circuit components. Further improvement in silicon devices, such as 12V rated MOSFETs for POL converters, will allow for additional improvement to achieve even higher levels of total system efficiency.

IDT's new topology can be implemented in a very cost-efficient way. At commercial volumes, the costs of implementing would be lower than a traditional DC distribution architecture. The solution will maintain, and even reduce, the number of components, while simplifying the power distribution across the server rack. The manufacturing costs of the system, and its cost vs. performance ratio, will be significantly improved as well. Also, no additional cost needs to be incurred outside the server rack – no retrofit or adjustment to facility layout, power system, or HVAC resources in the data centre.

This architecture is also favorable for future improvements. Because key elements of the control mechanisms are to be implemented in advanced silicon ICs, the topology creates an opportunity for data centres to fully leverage the continuing advances in silicon process cost and performance to realize energy savings. Such benefits in power management are not accessible to data centres using the traditional DC distribution architecture.

Over eight million servers are installed each year around the world and the total power consumption of server systems is estimated to be approximately over 60 billion kW per year. As data centre facility managers consider their organizations' growth requirements and costs, server power consumption will be a key factor in determining the most efficient and cost-effective system.

www.idt.com