A 99.7% Efficient Series-Stacked Architecture for Rack-Level Power Delivery in HDD Storage Servers

Ping Wang, Yenan Chen, Minjie Chen
Department of Electrical Engineering
Andlinger Center for Energy and the Environment
Princeton University, Princeton, NJ 08540, USA

***Abstract*—Distributing power with a 48 V-54 V voltage bus on the server rack level is becoming an increasingly popular solution for future energy efficient data centers. A 48 V to 5 V dc-dc converter with an efficiency of around 90% is typically needed for storage servers which host a large number of hard-disk-drives (HDDs) operating at 5 V. In this paper, we present a novel series-stacked power electronics architecture, based on the differential power processing (DPP) concept, that can simultaneously supply power to a large number of 5 V HDDs from a 50 V dc voltage bus with extremely high energy efficiency. With lower component count and smaller magnetics, the proposed architecture offers significantly improved power density while achieving a peak energy efficiency of 99.5%, which represents over 10x reduction in loss, 3x reduction in size, and 3x reduction in cost as compared to a traditional 48 V-5 V dc-dc converter. This architecture can dramatically reduce the power consumption on the rack level in future data centers, and can open new opportunities for power and computer architecture co-design.**

I. Introduction

U.S. data centers currently consume more than 90 billion kilowatt-hours of electricity a year and produce as much carbon emission as the entire airline industry. Data centers nowadays are still using a traditional power architecture which was designed for small-scale server system many years ago – each server is connected to a high voltage ac bus through an ac-dc grid-interface converter. Multiple cascaded power conversion stages are needed to support a variety of dc voltages needed for CPUs, memories, and HDDs (usually 1.8 V∼12 V), which limits the full end-to-end system efficiency. Delivering power in dc can reduce the number of power conversion stages and thus improve the overall system efficiency. One future trend is to place the uninterruptible power supply (UPS) at the 48 V∼54 V dc bus and distributed power at this voltage bus on the rack level. This architecture offers fewer power conversion stages and is compatible with the existing 48 V telecom ecosystem.

Conventional power delivery architectures typically utilize numerous dc-dc converters to step the 48 V∼54 V down to 5 V (Fig. 1a) as needed by the HDDs. All the power consumed by the 5 V HDDs need to be processed by one or more dc-dc converters. In storage servers, since the HDDs or SSDs are highly modular with uniform voltage rating (5 V or 12 V), there are opportunities to connect many HDDs or SSDs in series to the dc bus and eliminate the numerous dc-dc converters (Fig. 1b). Only the differential power between HDDs, which is only a small fraction of the total consumed power, needs to be processed by power electronics, which can significantly improve the system efficiency.

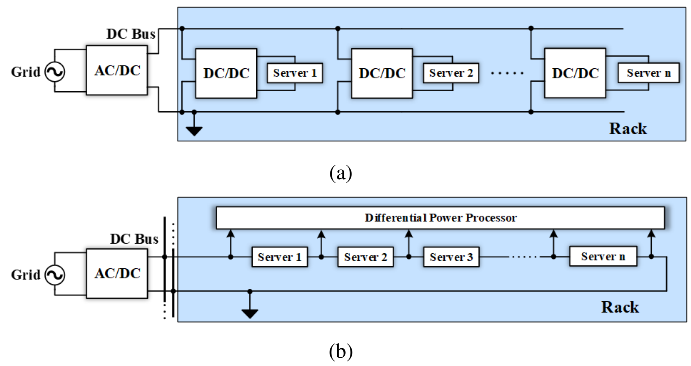

Fig. 1. Rack-level power delivery with a 48 V bus: (a) parallel architecture with numerous step-down converters; (b) proposed series-stacked architecture.

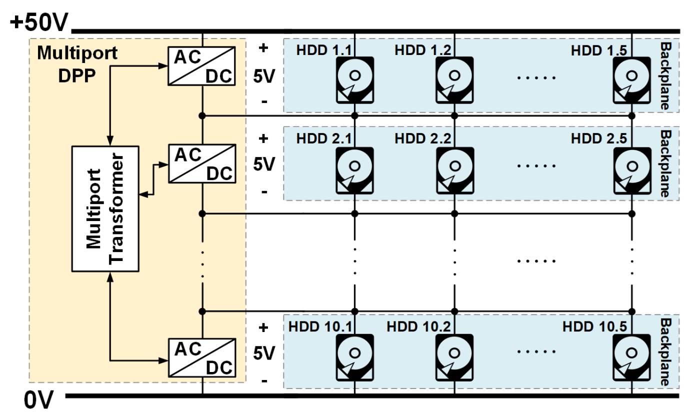

Fig. 2. Electrical diagram of a 50 HDD 300 W storage server. The system has 10 series-stacked 5 V voltage domains. Each voltage domain has 5 HDDs.

II. Series-Stacked Rack-Level Power Delivery

Fig. 3a shows the schematic of a multiport ac-coupled (MAC) converter that can perform differential power processing (DPP) for a large number of series stacked HDDs. A 300 W 50 HDD storage server with 10 series-stacked voltage domains has been built and tested (Fig. 3b). The proposed multiport DPP converter connects each voltage domain to a multi-winding transformer through a dc-ac unit. Each voltage domain supplies 5 V to five 2.5-inch HDDs, and ten of them

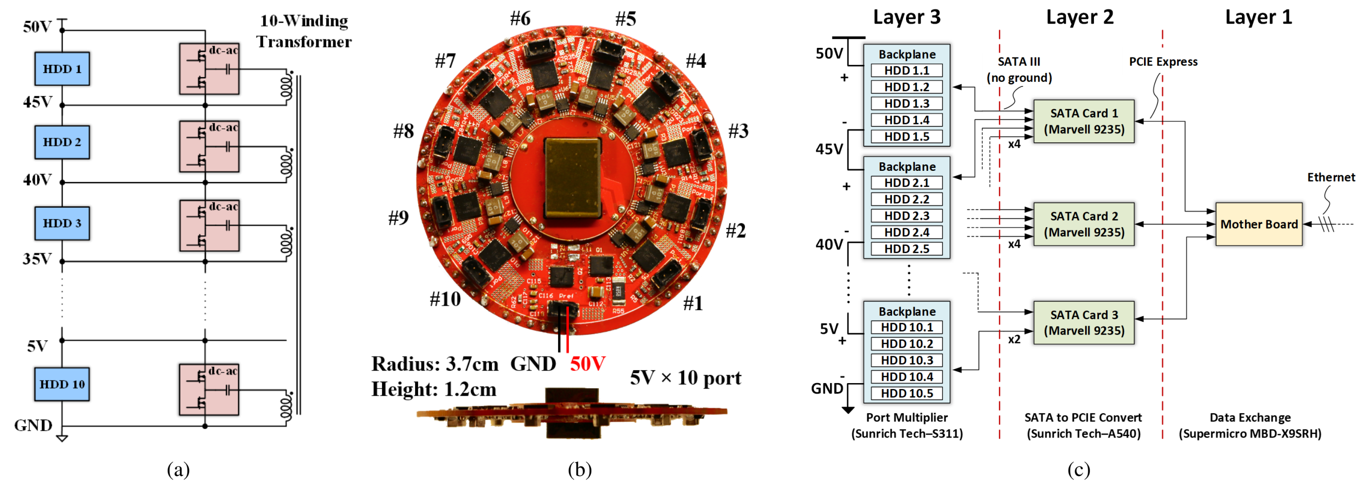

Fig. 3. (a) Circuit topology of the MAC-DPP converter; (b) Hardware prototype of the MAC-DPP converter; (c) Communication links of the prototype.

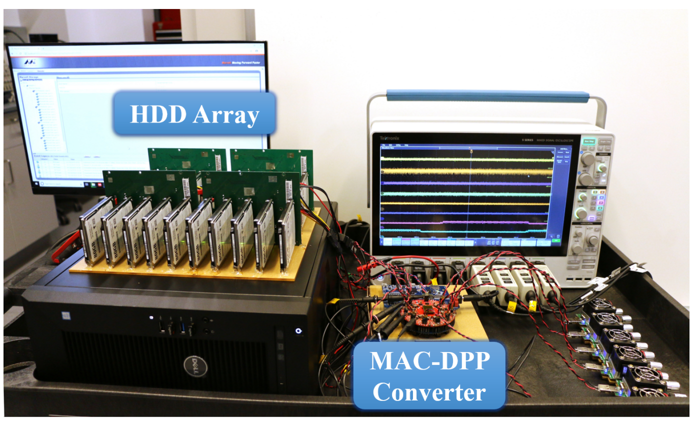

Fig. 4. The prototype HDD storage system testbench.

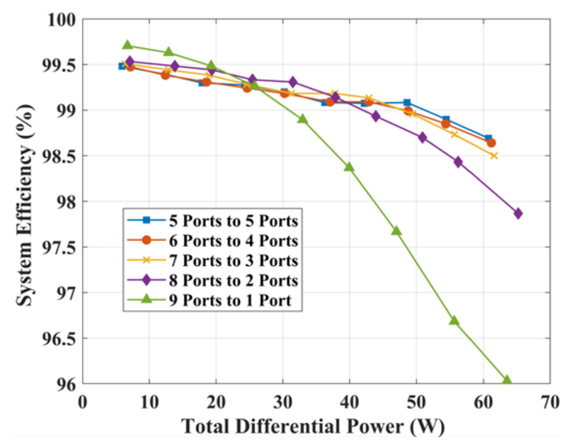

Fig. 5. Converter efficiency on rack.

are connected in series, supporting a dc voltage bus of 50 V. Advanced magnetics modeling and power flow control algorithms are needed to achive high efficiency and high power density [1]. This converter can achieve 10x reduction in loss, 3x improvement in density, and 3x reduction in cost than a typical 48 V to 5 V dc-dc converter.

Data transfer across different voltage domains are implemented as standard communication protocols with differential signals (e.g., SAS, SATA). In the testbench, we adopted a three-layer communication architecture as illstrated in Fig. 3c. There are ten backplanes boards (based on Marvell9235) in the testbench. Each backplane is placed on one voltage domain, hosting five HDDs. The backplane is connected to a SATA extension card with the grounding wire disconnected. The SATA card is inserted into the PCIE express sockets of the motherboard which manages all the data flow. The storage system functions normally with series-stacked power delivery architecture. Reading, writing and hot-swapping are tested.

In server racks, the storage tasks are usually well-distributed among many HDDs (e.g., RAID systems). Therefore, differential power between different voltage domains can be very small compared to the total load power. Series-stacked architecture that can step down voltage by nature only needs to process the differential power, resulting in extreme system conversion efficiency. Figure 5 shows the measured system efficiency in different power flow scenarios. The rack-level system efficiency can remain above 99% if the differential power only varies within 10% of the total load power.

III. Conclusion

This paper presented a novel series-stacked power architecture for HDD storage servers. A large number of HDDs are grouped into voltage domains and connected in series to eliminate the otherwise needed dc-dc converters. The proposed MAC-DPP converter offers low component count, high efficiency (99.7%) and high power density (100W/in^2). A prototype testbench has been tested with complete reading, writing, and hot-swapping to verify the proposed concept.

References

[1] P. Wang and M. Chen, "Towards Power FPGA: Architecture, Modeling and Control of Multiport Power Converters," *2018 IEEE 19th Workshop on Control and Modeling for Power Electronics (COMPEL)*, Padua, 2018, pp. 1-8.