A 99.7% Efficient 300 W Hard Disk Drive Storage Server with Multiport Ac-Coupled Differential Power Processing (MAC-DPP) Architecture

Ping Wang, Yenan Chen, Parker Kushima, Youssef Elasser, Ming Liu, and Minjie Chen
Princeton University, Princeton, NJ, United States
Email: {pwang2, yenanc, pkushima, yelasser, ml45, minjie}@princeton.edu

***Abstract*—High performance computing needs high performance power electronics. This paper presents the design of an ultra-efficient series-stacked hard-disk-drive (HDD) data storage server with a multiport ac-coupled differential power processing (MAC-DPP) architecture. A large number of HDDs are connected in series and ac-coupled through a multi-winding transformer with a single flux linkage. The MAC-DPP architecture offers very low power conversion stress, can achieve extremely high efficiency, and can reduce the magnetic size and the component count. A hybrid time-sharing and distributed phase-shift control strategy is developed to modulate the ac-coupled multi-input-multi-output (MIMO) power flow. A 10-port MAC-DPP prototype was designed to support a 300 W data storage system with 10 series-stacked voltage domains. The MAC-DPP converter was tested with a 50-HDD 12TB testbench, which can maintain normal operation of the server against the worst hot-swapping scenario. The 300 W MAC-DPP prototype can achieve 99.7% peak system efficiency and over 100 W/in^3 power density.**

***Index Terms*—multiport converter, series-stacked architecture, differential power processing, data center, distributed control**

I. Introduction

Due to the ever-increasing generation of data and computing demands in emerging applications (e.g., artificial intelligence and cloud computing), the scale of data centers and their power consumption increase dramatically. Data centers currently contribute about 2% of U.S. total electricity consumption and are still rapidly growing [1]. To maximize the computing power per unit volume and reduce the energy consumption, power converters for data centers need to be compact and efficient.

Data centers nowadays are still using a power delivery architecture that is based on the designs developed for single server scenarios - each server is connected to an ac voltage bus through an ac-dc PFC converter. Due to the low dc supply voltage need of IT equipments (e.g., 1.8 V$\sim$12 V for CPUs, RAMs, and HDDs), power supplies in servers usually employ many cascaded power conversion stages, only to limit the overall system performance. In a cascaded multi-stage architecture, the overall system efficiency tends to be low, as the full load power is processed at each stage. Moreover, due to the high voltage conversion ratio as needed, conventional power electronics designs are usually bulky and inefficient [2].

A recent trend in data center power architecture is to deliver dc power to each of the racks [3]. A dc voltage bus is created and the uninterruptible power supply (UPS) is placed on the rack level. The dc power distribution scheme contributes to reduced power conversion stages and improved power delivery efficiency. Compared to the 12 V intermediate bus architecture, delivering power at 48 V$\sim$54 V dc bus can reduce the power delivery loss and can leverage the existing telecom power ecosystem. To transfer power from the dc voltage bus to low voltage IT equipments, conventional power architecture employs numerous dc-dc converters with a variety of output voltage levels for different applications. In data storage servers, since the HDDs or SSDs are highly modular with uniform voltage ratings (5 V or 12 V), there are opportunities to connect a large number of HDDs or SSDs in series to realize inherent voltage step down. Since the vast majority of power is directly delivered to the loads and only a small amount of power difference is processed by the power converter (i.e., through differential power processing [4]), the power conversion stress can be significantly reduced, improving the system efficiency and power density [4].

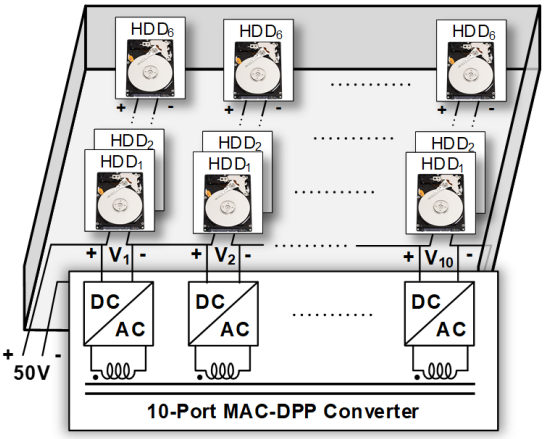

Fig. 1. A 300 W 60 HDD storage server with a 10-port MAC-DPP converter. Ten 5 V voltage domains are connected in series to the 50 V dc bus.

This paper presents an ultra-efficient series-stacked data storage system with a *multiport ac-coupled differential power processing* (MAC-DPP) architecture. The key concept is illustrated in Fig. 1. A 300W 10-port MAC-DPP converter is designed to support 60 HDDs which are configured into 10 series-stacked voltage domains (6 HDDs$\times$10). High-speed data transfer across different voltage domains is achieved with standard data transfer solutions (e.g., SAS, SATA). A hybrid time-sharing and distributed phase-shift control strategy is

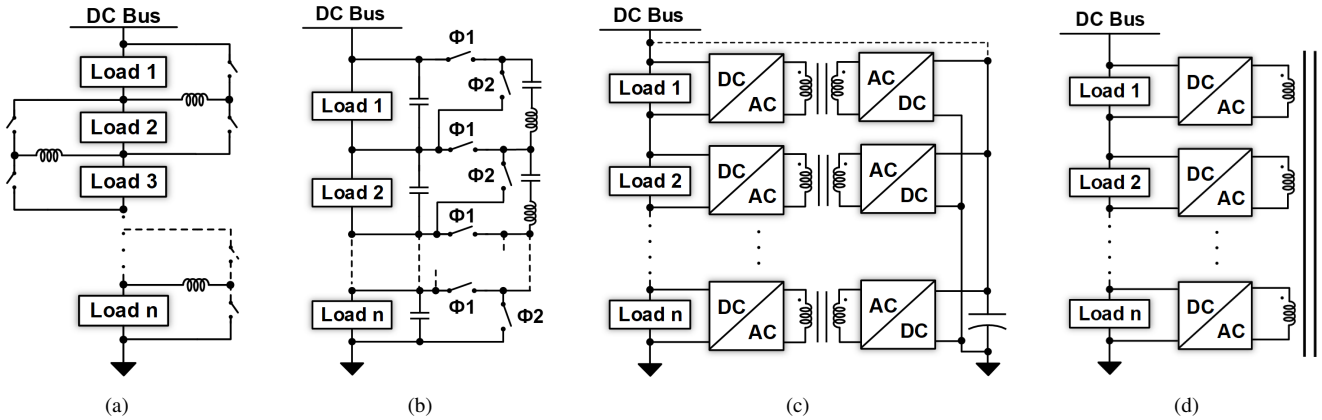

(a) (b) (c) (d)

Fig. 2. Block diagrams of a few DPP architectures: (a) Load-to-load DPP; (b) Switched-capacitor DPP; (c) Dc-coupled DPP; (d) Proposed MAC-DPP. The MAC-DPP architecture offers reduced power conversion stress, higher efficiency, smaller magnetic size, and lower component count.

TABLE I
COMPARISON OF A FEW DPP TOPOLOGIES

Topology	Load-to-load DPP	Switched-capacitor DPP	Dc-coupled DPP (half-bridge)	MAC-DPP (half-bridge)
Switch count	$2N-2$	$2N$	$4N$	$2N$
Switch voltage stress	$2V_{load}$	V_{load}	V_{load} or V_{bus}	V_{load}
Magnetic components	$N-1$ inductors	$N-1$ inductors	N two-winding transformers	one multi-winding transformer
Power conversion stages	multiple stages	multiple stages	two "dc-ac-dc" stages	one"dc-ac-dc" stage
Port-to-port isolation	non-isolated	non-isolated	galvanic-isolated	galvanic-isolated
Publication	[5]	[6], [7]	[8]	This work

utilized to modulate the differential power flow and regulate the voltage of each domain, which can maintain the normal operation of the storage server against the worst-case hot-swapping scenario. The remainder of this paper is structured as follows: Section II compares a few different DPP topologies and clarifies the advantages as well as design challenges of the MAC-DPP architecture. Section III analyzed the minimum cross-section area of the multi-winding transformer and Section IV discusses two typical power flow control strategies for the MAC-DPP converter. Detailed experimental results are provided in Section V, including the design of a 10-port MAC-DPP prototype and the configuration of a 12TB HDD server testbench. Finally, Section VI concludes this paper.

II. MULTIPORT-AC-COUPLED DPP ARCHITECTURE

Many DPP converter topologies have been proposed [5]–[8]. Fig. 2 compares the proposed MAC-DPP architecture against many existing DPP solutions. As shown in Fig. 2a, a load-to-load DPP architecture uses a bidirectional buck-boost circuit to process the differential power between two neighboring loads [5]. Compared to a traditional load-to-bus structure [4], the load-to-load DPP converter has lower switch voltage stress ($2V_{load}$), enabling it to utilize switches with lower voltage rating and operate at higher frequencies with smaller passive component size. However, the differential power that needs to be delivered between two non-adjacent loads has to go through multiple power conversion stages. This creates higher power conversion losses and limits the system dynamic performance, making it hard to be extended to a large number of ports or applied to the cases of highly-dynamic loads. Fig. 2b shows a resonant ladder switched-capacitor (SC) DPP topology [6], [7]. The ladder SC-DPP converter can achieve higher efficiency and higher power density. However, this solution is another embodiment of the load-to-load DPP architecture, and it can only process the differential power between two neighboring loads in one switching cycle.

An alternative DPP approach is to employ multiple isolated dc-dc converters and connect each series-stacked voltage domain to a common dc bus [8], as depicted in Fig. 2c. The dc-coupled DPP architecture can transfer power directly between two arbitrary loads. Compared to ladder-structure based DPP options, this architecture is more scalable and can offer better dynamic performance. However, the dc-coupled DPP topology requires multiple magnetics as well as high component count, which will increase the cost and total converter size. Moreover, the differential power needs to go through at least two "dc-ac-dc" stages from one port to another, resulting in additional power conversion stress and losses [9].

As shown in Fig. 2d, the proposed MAC-DPP architecture connects each voltage domain to a multi-winding transformer through a dc-ac unit. The differential power flow is coupled by the centralized transformer, and each dc-ac unit can be simply implemented as a half-bridge stage with a dc blocking capacitor, or other options such as full-bridge-based topologies or Class-E-based topologies. The power transferred between two different loads is galvanic isolated and is bidirectional. Table I lists the detailed comparison of different DPP architectures

assuming half-bridge implementation for all dc-ac units. The advantages of the proposed MAC-DPP architecture include:

- **Fewer "dc-ac-dc" power conversion stages:** The MAC-DPP architecture directly transfers power between two arbitrary ports with one single "dc-ac-dc" conversion stage. Other DPP solutions usually need two or more "dc-ac-dc" stages when delivering power between two arbitrary loads. The reduced power conversion stress improves the system dynamic performance and reduces the losses.
- **Reduced component count:** In a MAC-DPP architecture, one voltage domain is connected to one dc-ac unit, and n voltage domains only need n dc-ac units. The MAC-DPP architecture is highly modular. It has almost the lowest component count among existing DPP options, leading to reduced cost and improved power density.
- **Smaller magnetic size:** Compared to the dc-coupled DPP converter that needs multiple transformers, the MAC-DPP architecture has only one magnetic core. In principle, the magnetic core area of a multi-winding transformer is determined by the highest volt-second-per-turn of all windings instead of the winding count. In a MAC-DPP architecture with fully symmetric configuration, each dc-ac unit have identical voltage rating, and all windings have identical volt-second-per-turn, which can stay the same as winding count increases. Therefore, the core area of a multi-winding transformer in the MAC-DPP is roughly the same as that of a two-winding transformer in other isolated DPP options. Only the window area increases as the winding count increases. Theoretically, the MAC-DPP architecture can reduce the magnetic core area by n times compared to other isolated DPP implementations (n is the number of series-stacked voltage domains).

III. Minimum Core Area of the Multiwinding Transformer

One challenge of designing the MAC-DPP converter is to build a high performance low volume multi-winding transformer. The basic requirement is to maintain normal operation of all the windings without saturating the magnetic core. Generally, the rule that the magnetic core area is determined by the maximum volt-second-per-turn in a two-winding transformer can be extended to the multi-winding cases, and the minimum core area can be obtained by analyzing the highest magnetizing current in the worst-case scenario. Fig. 3 shows the normalized equivalent lumped circuit model of the multi-winding transformer in a multiport ac coupled converter. The dc-ac units at each port (e.g., half-bridge stages with dc blocking capacitors) are modeled as square wave voltage sources driving an n port inductor network. V_i is the square wave voltage amplitude and N_i is the winding turn. L_i includes the leakage inductance as well as external inductance of port $\#i$. The current that flows through the equivalent magnetizing inductance in one switching period can be obtained by applying superposition:

$$I_m = \sum_{i=1}^{n} \frac{V_i}{N_i L_{eqi}} \left(\frac{T}{4} - \left| t - \frac{T}{2}\left(1 - \frac{\phi_i}{\pi}\right) \right| \right) \times K_i, \quad (1)$$

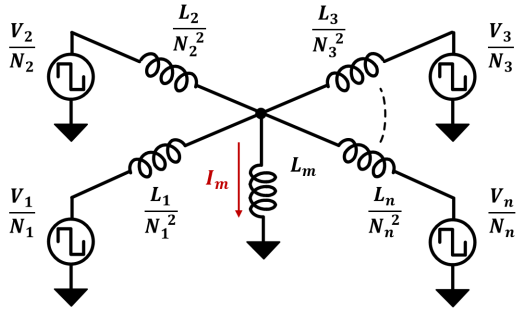

Fig. 3. Lumped circuit model of the multi-winding transformer in the multiport ac-coupled converter. L_m is the magnetizing inductance and L_i is the branch inductance of each port.

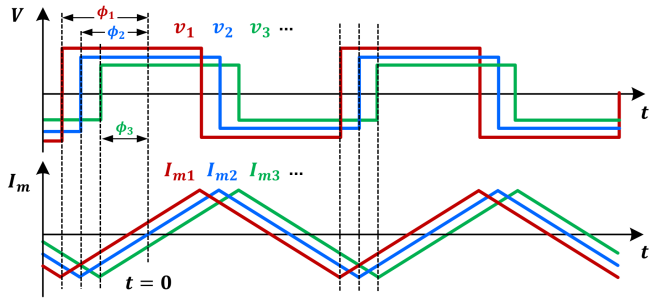

Fig. 4. Voltage and current waveforms when applying superposition to the n dc-ac units. The magnetizing current reaches maximum when all current waveforms are in phase.

in which T is the switching period and ϕ_i is the phase-shift. L_{eqi} is the equivalent inductance seen from the i_{th} port and the K_i is the current shunting ratio of port $\#i$, which can be expressed as functions of the inductance values:

$$\begin{cases} L_{eqi} = \left(\sum_{k \neq i} \frac{N_k^2}{L_k} + \frac{1}{L_m} \right)^{-1} + \frac{L_i}{N_i^2}, \\ K_i = \frac{\left(\sum_{k \neq i} \frac{N_k^2}{L_k} \right)^{-1}}{\left(\sum_{k \neq i} \frac{N_k^2}{L_k} \cdot \right)^{-1} + L_m} \end{cases} \quad (2)$$

Each square wave voltage source generates a triangular magnetizing current in L_m as illustrated in (1). Fig. 4 demonstrates the magnetic currents induced by different voltage sources considering the phase-shifts. For each voltage source, the induced magnetic current will reach the peak at the time of $\frac{T}{2}\left(1 - \frac{\phi_i}{\pi}\right)$, and the total magnetic current will reach the maximum when all the sources are in phase. In a symmetric multiport ac-coupled (MAC) converter design, each winding has identical V_i/N_i and L_i/N_i^2. Denoting them as V_0 (volt-per-turn) and L_0 separately, the maximum magnetic current is:

$$I_{m\_\max} = \sum_{i=1}^{n} \frac{K_i V_i}{N_i L_{eqi}} \times \frac{T}{4} = \frac{V_0}{\frac{L_0}{n} + L_m} \times \frac{T}{4} \approx \frac{V_0 T}{4 L_m}. \quad (3)$$

The maximum magnetic current is almost irrelevant with the winding count n, and the maximum magnetic flux ($\Phi_{m\_\max} =$

$L_m I_{m\_\max} = V_0 T/4$) is only determined by the volt-second-per-turn ($V_0 T$). Accordingly, the minimum core area ($A_{\min}$) of the multiwinding transformer in a symmetric MAC converter configuration is:

$$A_{\min} = \frac{\Phi_{m\_\max}}{B_{\max}} \approx \frac{V_0 T}{4 B_{\max}}. \tag{4}$$

As a result, the minimum core area of a multi-winding transformer is roughly the same as that of a two-winding transformer, if their volt-second-per-turn are the same. This is the fundamental reason why the proposed MAC-DPP architecture can achieve much higher power density and better magnetic utilization than other isolated DPP implementations. Compared to other non-isolated DPP options, the MAC-DPP architecture offers reduced power conversion stress (fewer "dc-ac-dc" stages) and lower component count.

IV. Power Flow Control of MAC-DPP Converter

The MAC-DPP converter is a multi-input-multi-output system. To maintain stable port voltage with unbalanced load power, accurate differential power flow control is needed. Fig. 5 shows the equivalent delta model modified from the star-model in Fig. 3. The power flow that transfers through each link inductor (L_{ij}) can be calculated in the same way as that in a dual active bridge (DAB) converter [10]. The total power feeds into the passive network from the i_{th} port is:

$$P_i = \sum_{j=1}^{n} \frac{V_i V_j}{2\pi f_s N_i N_j L_{ij}} \phi_{ij} \left(1 - \frac{|\phi_{ij}|}{\pi}\right). \tag{5}$$

As indicated in (5), the input differential power of one port (P_i) is related with the phase-shifts of all ports $\{\phi_1, \phi_2, ..., \phi_n\}$. The closely-coupled power flow brings challenges to the port voltage regulation, especially in the case where a large number of loads are stacked in series. [11] discusses a few different control methods for multi-active-bridge (MAB) converters. Two control strategies, the time-sharing control and the phase-shift control are adopted in this MAC-DPP architecture.

Fig. 6 shows the principles of the time-sharing control. Each port of the MAC-DPP converter is enabled for a specific proportion of time during one modulation cycle (T_m), which usually contains numerous switching periods. Only two ports are enabled at one time instance. They are denoted as active ports. Other ports are inactive and do not transfer power. In this case, the MAC-DPP converter delivers power in the same way as a dual-active-bridge converter. The switching frequency and the relative phase-shift between two active ports (ϕ_{ij}) can be adjusted to achieve the maximum port to port power delivering efficiency. The amount of the input differential power at the i_{th} port is controlled by modulating its active duty ratio in one modulation cycle (D_i). With time-sharing control, the MAC-DPP converter can always operate at the phase-shift of the optimal efficiency and reduce losses during the idling period (i.e. when no port is enabled), which will benefit the overall system efficiency. Time-sharing control applies to the case where the load power is lightly unbalanced with small differential power.

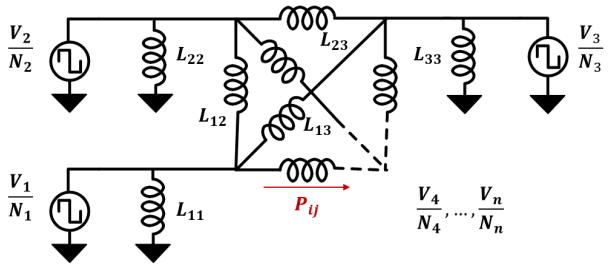

Fig. 5. Modeling the multi-winding transformer in the MAC-DPP converter.

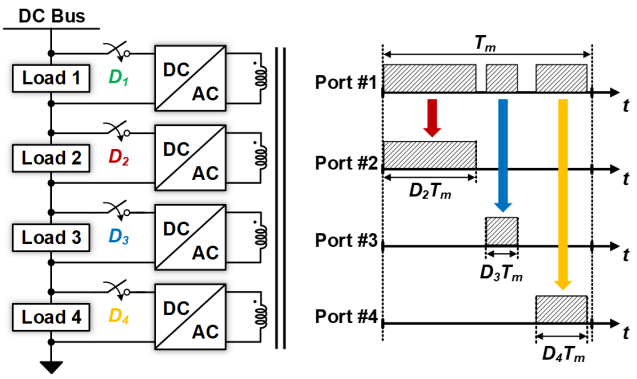

Fig. 6. Time-sharing control of an example four-port MAC-DPP design.

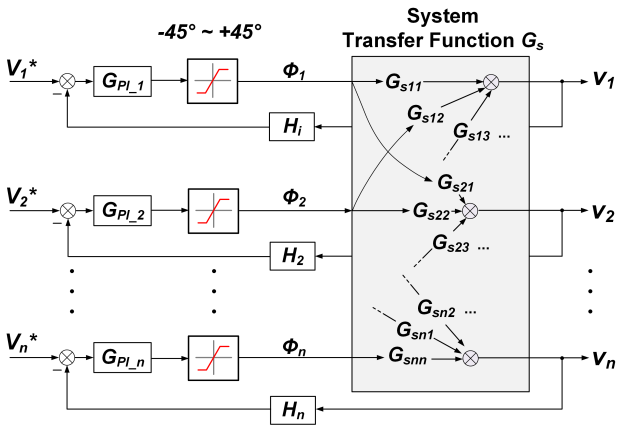

Fig. 7. Block diagrams of the distributed phase-shift control strategy.

Different from time-sharing control, phase-shift control has all the ports enabled simultaneously, and the power flow can be manipulated by changing the phase-shift of each port. Phase-shift control is capable of delivering large differential power at heavily unbalanced load, but it might run into voltage oscillation because of the closely-coupled power flow. [12] presents a systematic approach to identify the small signal model for the MAC-DPP converter with an arbitrary number of ports. The modeling approach accurately captures the impacts of power losses, and the system transfer function matrix ($\boldsymbol{G_s}$) that describes the dynamics from any control phase-shift (ϕ_i) to port voltage (v_j) is derived. To regulate the load voltage, a feedback loop can be implemented at each port to adjust its own phase-shift based on the locally measured port voltage, as depicted in Fig. 7. However, the non-diagonal

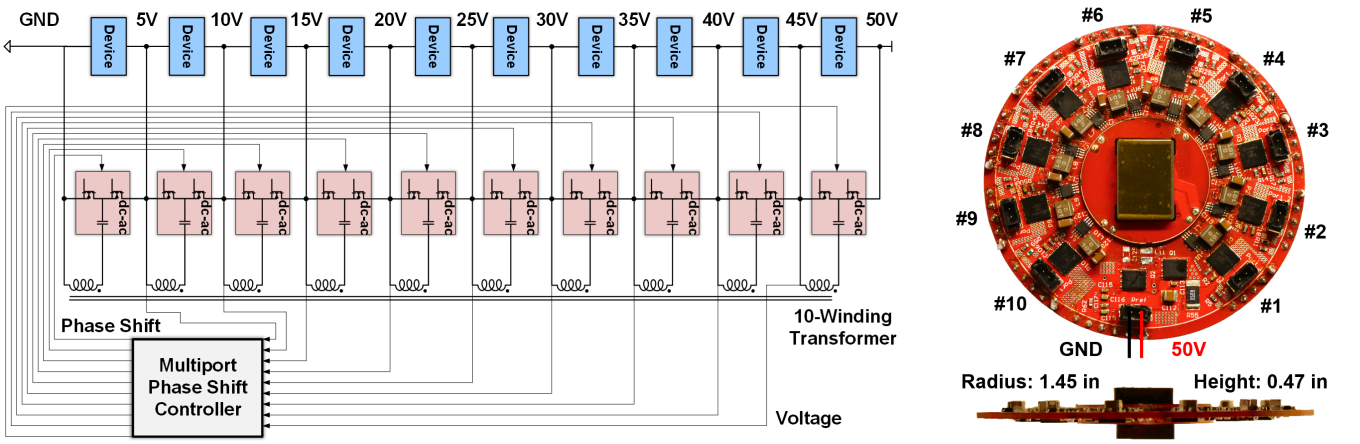

(a) (b)

Fig. 8. (a) Topology of a 10-port MAC-DPP converter with dc-ac units implemented as half-bridge circuits. (b) Prototype of a 300W 10-port MAC-DPP converter. DrMOS (FDMF6833C) switches are utilized as the half-bridge units, operating at 100kHz. Each port is connected with $300\mu F$ blocking capacitance (X5R, $100\mu F \times 3$) and $100nH$ external inductance (Coilcraft XEL4030).

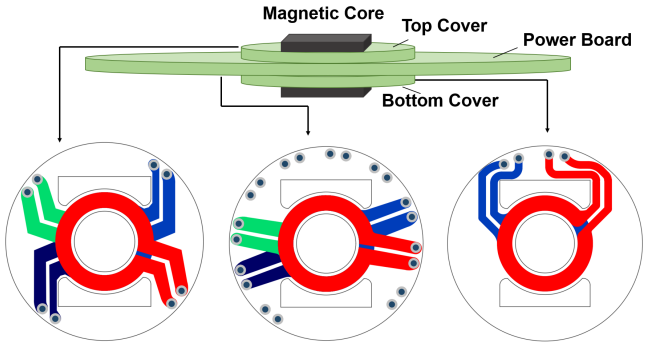

Fig. 9. Structure of the 10-winding planar transformer. Each winding has one single turn on a four-layer PCB with 1 oz copper. Magnetic core: EQ20-N97.

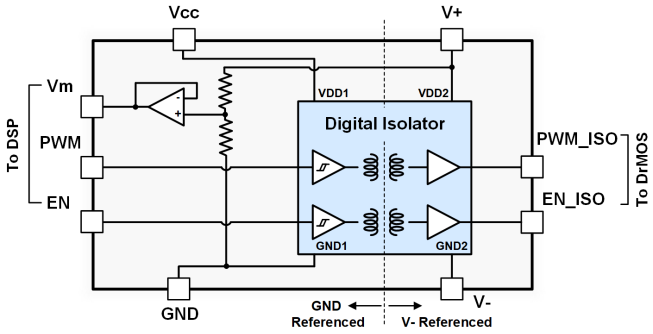

Fig. 10. Modular implementation of the isolated control circuitry (PWM & enable) and voltage sampling circuit at each port.

elements ($G_{sij(i\neq j)}$) of the transfer function matrix indicates the interaction between different control loops. To decouple the control loops, an inverse matrix of the system transfer function could be applied [13], but this method places heavy computation burden on the controller, making it difficult to be extended to the cases with a large number of ports. This paper adopted a distributed phase-shift control strategy which considers the interactions between different feedback loops as disturbances. Based on the derived system transfer function, the loop gain of each individual port voltage regulator is:

$$G_{Li}(s) = G_{PI\_i}(s) \times G_{sii}(s) \times H_i(s) \tag{6}$$

Here $G_{PI\_i}(s)$ is the PI controller parameters. $G_{sii}(s)$ is the diagonal elements of the system transfer function matrix. $H_i(s)$ is the transfer function of the sampling circuitry. The explicitly derived loop gain can be used to analyze the dynamic performance of the system. Through carefully designing the phase margin of each loop gain, the oscillation caused by the interactions between different feedback loops can be avoided.

A well-designed storage server usually has equally-allocated storage tasks among many HDDs (e.g., RAID systems). Each HDD has similar reading/writing load. If many HDDs are connected in parallel in one series-stacked voltage domain, the power consumptions of multiple voltage domains are usually very close to each other with very low differential power. Time-sharing control can be applied in this scenario to improve the system efficiency. Under the circumstances where there is large differential power (e.g., hot-swapping, startup, etc.), the MAC-DPP converter can shift to phase-shift control to ensure rapid and stable port voltage regulation.

V. Experimental Prototype and Results

To experimentally validate the MAC-DPP architecture in data center rack-level power delivery, a 10-port MAC-DPP converter of 300 W power rating is built and tested. A 50-HDD data storage server is set up for testing the operation of MAC-DPP converter in normal operation and hot-swapping. This section introduces the hardware setup of both the MAC-DPP converter and the HDD server testbench. The effectiveness of the MAC-DPP architecture and the control strategy are experimentally validated.

A. A 10-Port 300 W MAC-DPP Converter Power Stage

Fig. 8 shows the circuit topology and hardware prototype of a 10-port MAC-DPP converter. The dc-ac units are imple-

mented as half-bridges with dc blocking capacitors, connected to a 10-winding transformer. The 50 V dc bus is split into 10 series-stacked 5 V voltage domains to support numerous 2.5 inch HDDs. The MAC-DPP converter is designed to process over 30 W differential power for each port. Therefore, each voltage domain can support 6 HDDs of 30 W peak power (5 W$\times$6 during spinning up) in case of the worst hot-swapping scenario where the HDDs of an entire voltage domain are removed, and the entire 10-port MAC-DPP architecture can support 60 HDDs with over 300 W peak power rating. The MAC-DPP prototype is 1.45" in radius and 0.47" in height, so the power density of the prototype is over 100 W/in^3.

Fig. 9 shows the structure of the 10-winding planar transformer. Three 4-layer PCBs and two EQ20-N97 magnetic cores are stacked together. Ten windings are distributed on the three PCBs, and each winding has one single turn. The windings on the top cover and bottom cover are connected to the power board through vias. The cross-section area of the central magnetic core is only 0.17" in radius, which is the similar size as that in a two winding transformer of the same volt-second-per-turn (5V at 100kHz). A conventional isolated DPP design would require ten similar magnetic cores, while the MAC-DPP architecture only needs one.

The voltage sampling circuits and the isolated control signals including PWM signals and the enable signals are implemented modularly for each port as depicted in Fig. 10. In each module, a two-channel isolator is applied, and its secondary side is powered by the corresponding 5 V voltage domain directly. No additional isolated supply is needed. The voltage sampling circuit utilizes a resistive divider with a buffer to scale down the upper node voltage (V_+) and send it back to the DSP. Distributed phase-shift control is implemented in the DSP, generating the control signals with different phase-shifts, as indicated in Fig. 8a. The centralized multiport phase-shift controller (DSP) can also be implemented as separate control modules, which, together with the voltage sampling circuits, can be further integrated into the half-bridge power stage, enabling fully integrated modular building blocks for the MAC-DPP architecture.

B. HDD Server Testbench and Data Link Infrastructure

Fig. 11 shows the 50-HDD 12TB storage server testbench. A linux based operating system (Ubuntu) is installed to test the full reading, writing and hot-swapping capability of the HDD array. A dc voltage source (QPX-600D) is applied to mimic the 50 V dc bus on rack. Fig. 12 shows the detailed implementation of the high-speed data transfer infrastructure across series-stacked voltage domains. The communication infrastructure comprises three layers. The 50 HDDs are divided into 10 groups, and each group contains five HDDs in parallel on a SATA III port multiplier (i.e. backplane). Ten backplanes in different voltage domains transfer data to the SATA extension card through differential signal with dc blocking capacitors. Actually, standard SATA/SAS protocol is based on differential signal. By removing the common ground wires and adding blocking capacitor to SATA/SAS differential signal links, the

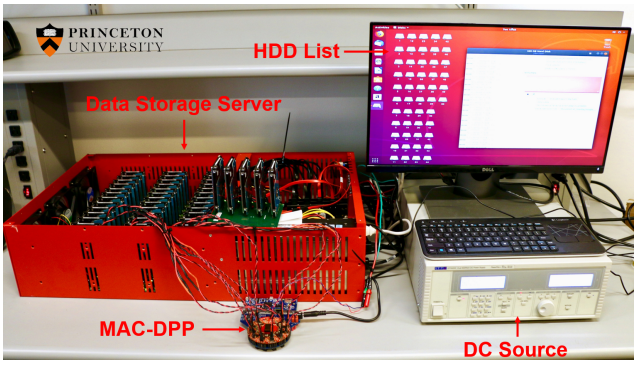

Fig. 11. Picture of the 12TB HDD testbench modified from a Backblaze server.

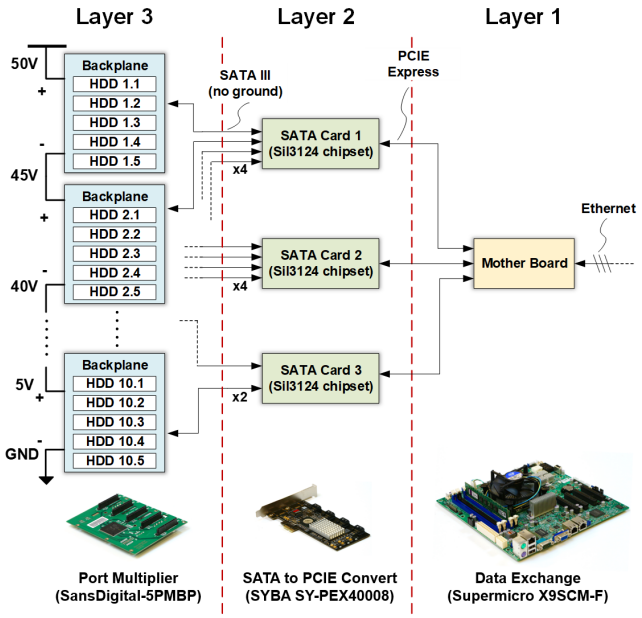

Fig. 12. Data link infrastructure of the series-stacked HDD testbench.

isolated signal transfer across voltage domains are achieved without major modification to standard protocols. Different SATA to PCIE extension cards are on the same voltage domain. They are directly connected to the motherboard through the PCIE Express slots. The 3-layer data link infrastructure is highly scalable, and it can be extended to support larger scale of HDD array with more series-stacked voltage domains. In applications where high data rate is needed, the isolated SATA transmission can also be replaced with optic fibers, which can offer higher communication bandwidth.

C. Experimental Results

In the experiment with heavily unbalanced loads, we adopted distributed phase-shift control to regulate the voltage for the proposed MAC-DPP architecture. In this case, multiple ports are inputting or outputting the differential power at the same time. We define the MAC-DPP *converter efficiency* as the total output differential power divided by the total input differential power, and define the *system efficiency* of

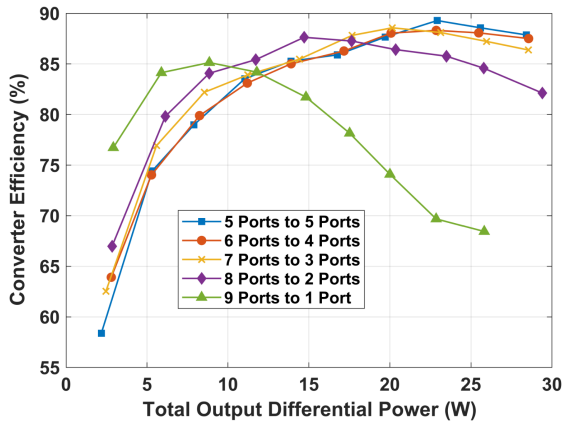

Fig. 13. Port to port power conversion efficiency.

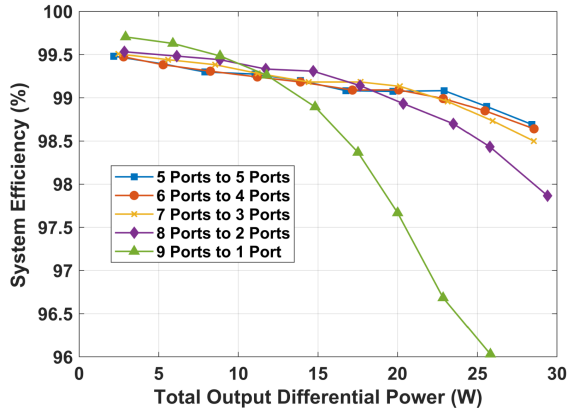

Fig. 14. System power conversion efficiency (total load power: 300 W).

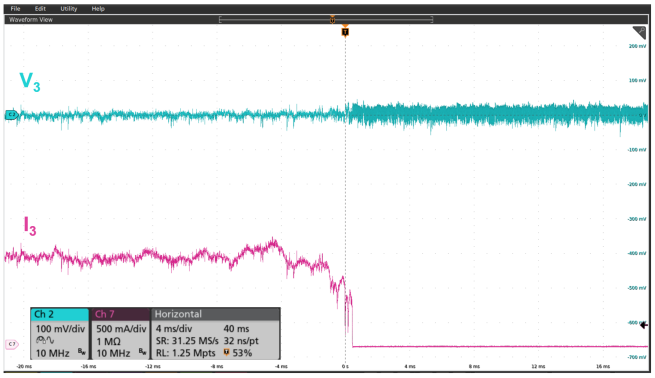

Fig. 15. Transient response when hot-swapping an entire voltage domain (removing 5 HDDs from port #3) of the HDD server testbench.

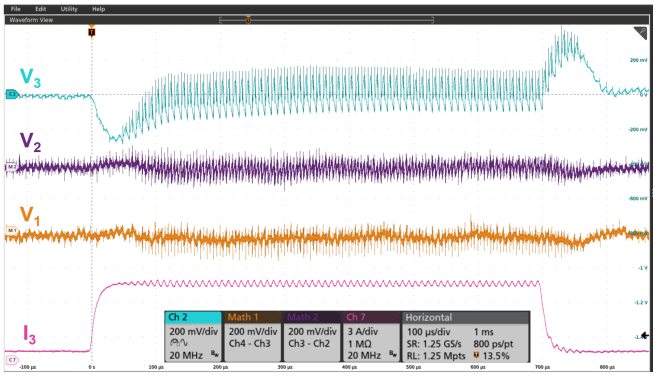

Fig. 16. Transient response of a 30 W step load change at port #3. The settling time is 100 μs. The voltage overshoot is less than 250 mV.

a MAC-DPP system as the total power consumed by the load (e.g., HDDs), over the total power sourced from the dc bus. Both the converter efficiency and system efficiency were tested using ten electronic loads with 300 W total load power, and the overall output differential power varies from 0 W to 30 W. Since the HDDs on each voltage domain usually have similar power consumption, the tested data points well cover a majority of the operating conditions (20% power mismatch among different voltage domains) as well as the worst hot-swapping scenario (removing 30 W full load HDDs of an entire voltage domain). Fig. 13 shows the measured converter efficiency in different power transfer scenarios. The maximum port-to-port converter efficiency is about 90% when it is delivering 23 W differential power from 5 ports to 5 ports. The converter efficiency drops very fast when delivering power from 9 ports to 1 port because of the dramatically increased current conduction loss at one port. As indicated in Fig. 14, the MAC-DPP system can achieve 99.7% peak system efficiency when processing 7 W differential power, and the system efficiency maintains over 99% across a majority of the operation scenarios when the power mismatch among different voltage domains is within 10% (i.e. 15 W total differential power). Compared to existing solutions, the proposed MAC-DPP converter can realize extremely high system power efficiency with very small converter size, and can significant improve the computing power per unit volume in data center.

Voltage stability is of great importance for HDD's robust operation. A typical requirement for a 2.5" HDD is to regulate the voltage within 5% of the nominal value (5 V). Fig. 15 shows the port voltage and load current waveforms at the 3^{rd} port during a hot-swapping test on the HDD server testbench. A 2.5 mF output capacitor was included at each port. In the test, all five HDDs of the same voltage domain were removed abruptly when they were performing reading/writing tasks. The voltage transition during the hot-swapping was very smooth. After hot-swapping, since the MAC-DPP converter processes more differential power, the voltage ripple at port #3 increases (still within the allowed range for 2.5" HDDs).

Since the MAC-DPP converter is designed to support 30 W peak power at each voltage domain, we tested the transient response of the prototype in an extreme case where 30 W full load was plugged in and removed at one port (i.e. worst hot-swapping scenario). In the test, each series-stacked voltage domain was connected to an electronic load. All the load currents kept at 1A except for the current at port #3, which jumped up from 1A to 7A and then returned back to 1A, as shown in Fig. 16. With the distributed phase-shift control strategy, the MAC-DPP converter can successfully limit the overshoot of the "hot-swapping" port voltage to 250 mV with only 100 μs settling time, fulfilling the voltage requirements

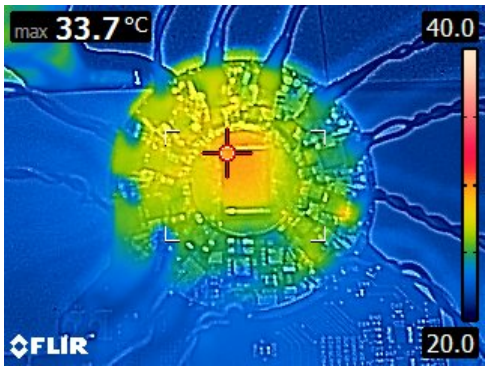

(a)

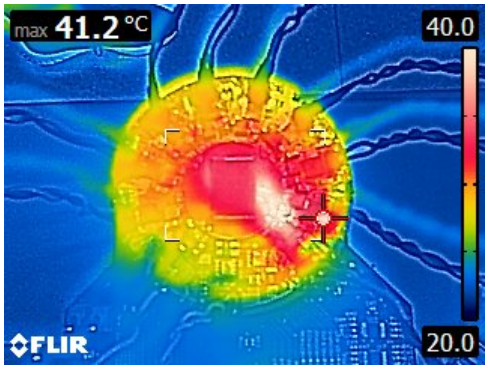

(b)

Fig. 17. Thermal images of the MAC-DPP prototype in (a) balanced load and (b) hot-swapping with 25°C ambient temperature.

for typical 2.5" HDDs. Fig. 16 also indicates that the interactions between ports will lead to voltage fluctuation on the other ports (e.g., V_1&V_2), but they can be effectively controlled by the proposed control strategy as well. These hot-swapping experiments verified that the designed MAC-DPP prototype is capable of maintaining smooth operation of a 300 W data storage server against the worst-case hot-swapping scenarios.

Fig. 17 shows the thermal images of the MAC-DPP converter operating at different operating conditions on the HDD server testbench. Both the thermal images are taken under the natural air cooling of 25°C ambient temperature without the air flow. If all the HDDs are doing similar reading/writing tasks, very little differential power needs to be processed by the MAC-DPP converter. Almost no temperature rise was observed. When all five HDDs of an entire voltage domain are removed, the hot-swapping port will deliver about 10 W differential power to the other 9 ports. Since the current at the hot-swapping port is roughly the summation of currents at the other 9 ports, its conduction loss is significantly higher than others. Therefore, a significant temperature rise was observed at one port (port #1 in this case) as shown in Fig. 17b.

VI. Conclusion

This paper presents a MAC-DPP architecture for series-stacked storage servers in data centers. The MAC-DPP architecture offers low component count, single "dc-ac-dc" conversion stage and "almost" the smallest magnetic size. The multi-winding transformer is implemented as a PCB transformer. Hybrid time-sharing control and phase-shift control is adopted in the MAC-DPP design for the HDD storage systems. A 300 W 10-port MAC-DPP prototype is designed and tested in a 50-HDD data storage server testbench. The HDD server can maintain normal reading/writing operation when hot-swapping the HDDs of an entire voltage domain. The MAC-DPP prototype was also tested in an extreme case where 30 W full load was hot-swapped at one port. The transient response of the MAC-DPP system meets the requirements of typical HDDs, and the system efficiency for a 300 W storage server can remain above 99% for a majority of operating conditions. The MAC-DPP architecture is a very promising solution for powering highly modular energy systems, such as series-stacked HDD arrays, battery cells, and PV modules.

Acknowledgment

The authors would like to thank the DOE ARPA-E CIRCUIT program and the National Science Foundation (# 1847365) for supporting this work.

References

[1] A. Shehabi, S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, and W. Lintner, "United States Data Center Energy Usage Repor," 2016.

[2] M. H. Ahmed, C. Fei, F. Lee and Q. Li, "Single-Stage High-Efficiency 48/1V Sigma Converter with Integrated Magnetics," *IEEE Transactions on Industrial Electronics*, to appear.

[3] S. Jiang, S. Saggini, C. Nan, X. Li, C. Chung and M. Yazdani, "Switched Tank Converters," *IEEE Transactions on Power Electronics*, vol. 34, no. 6, pp. 5048-5062, June 2019.

[4] P. S. Shenoy and P. T. Krein, "Differential Power Processing for DC Systems," *IEEE Transactions on Power Electronics*, vol. 28, no. 4, pp. 1795-1806, April 2013.

[5] G. L. Brainard, Non-dissipative battery charger equalizer, U.S. Patent 5479083, Dec. 26, 1995

[6] J. T. Stauth, M. D. Seeman and K. Kesarwani, "Resonant Switched-Capacitor Converters for Sub-module Distributed Photovoltaic Power Management," *IEEE Transactions on Power Electronics*, vol. 28, no. 3, pp. 1189-1198, March 2013.

[7] A. Stillwell and R. C. N. Pilawa-Podgurski, "A Resonant Switched-Capacitor Converter with GaN Transistors for Series-Stacked Processors with 99.8% Power Delivery Efficiency," *2015 IEEE Energy Conversion Congress and Exposition (ECCE)*, Montreal, QC, 2015, pp. 563-570.

[8] E. Candan, P. S. Shenoy and R. C. N. Pilawa-Podgurski, "A Series-Stacked Power Delivery Architecture with Isolated Differential Power Conversion for Data Centers," *IEEE Transactions on Power Electronics*, vol. 31, no. 5, pp. 3690-3703, May 2016.

[9] 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.

[10] R. W. A. A. De Doncker, D. M. Divan and M. H. Kheraluwala, "A Three-Phase Soft-Switched High-Power-Density DC/DC Converter for High-Power Applications," *IEEE Transactions on Industry Applications*, vol. 27, no. 1, pp. 63-73, Jan.-Feb. 1991.

[11] Y. Chen, P. Wang, H. Li and M. Chen, "Power Flow Control in Multi-Active-Bridge Converters: Theories and Applications," *2019 IEEE Applied Power Electronics Conference and Exposition (APEC)*, Anaheim, CA, USA, 2019, pp. 1500-1507.

[12] P. Wang, Y. Chen, Y. Elasser and M. Chen, "Small Signal Model for Very-Large-Scale Multi-Active-Bridge Differential Power Processing (MAB-DPP) Architecture," *2019 IEEE 20th Workshop on Control and Modeling for Power Electronics (COMPEL)*, Toronto, ON, 2019, pp. 1-8.

[13] C. Zhao, S. D. Round and J. W. Kolar, "An Isolated Three-Port Bidirectional DC-DC Converter With Decoupled Power Flow Management," *IEEE Transactions on Power Electronics*, vol. 23, no. 5, pp. 2443-2453, Sept. 2008.