A Series-Stacked Power Delivery Architecture With Isolated Differential Power Conversion for Data Centers

Enver Candan, *Student Member, IEEE*, Pradeep S. Shenoy, *Member, IEEE*, and Robert C. N. Pilawa-Podgurski, *Member, IEEE*

***Abstract*—In this paper, an alternative method to achieve more efficient dc power distribution and voltage regulation for future data centers is presented. This paper describes a series-stacked power delivery architecture, where servers are connected electrically in series, thereby, providing inherent step down of voltage. Server voltage regulation is performed by differential power converters, which only process the mismatch power between servers. The bulk power flows with no power processing, yielding greatly increased system efficiency compared to conventional architectures. We demonstrate the series-connected architecture with an experimental proof of concept and compare the proposed architecture with a conventional dc power delivery architecture employing a best-in-class power supply unit for servers. The proposed power delivery architecture is validated with a series-stacked server rack consisting of four 12 V servers, powered from a 48 V dc bus, performing two different real-world operations: web traffic management and computation. Through experimental measurements, we demonstrate up to a 40x reduction in power conversion losses compared to state-of-the-art hardware, and an overall best-case system conversion efficiency of 99.89%.**

***Index Terms*—Bidirectional power flow, dc-dc power converters, differential power processing, high efficiency data center power delivery.**

I. Introduction

DATA centers are key facilities for processing large datasets, enabling online services such as search, video streaming, and scientific computing. With the rapid increase in both number and size of data centers, the energy efficiency of data centers has become an important challenge for power engineers [1]. Worldwide data centers consumed between 1.1% and 1.5% of the global electric load in 2010, while the ratio for the US was between 1.7% and 2.2% [2]. Moreover, a recent report has estimated that US data centers consumed 91 billion kWh of energy in 2013 [3]. The continued growth of the Internet and cloud services indicate that data centers are going to be significant energy consumers in the future, and therefore, their energy efficiency is a critical consideration for sustainable growth of this technology.

The most power-hungry Information Technology (IT) equipment in data centers are the servers, which store and process data. Servers are computational loads that operate from low dc voltage (e.g., 12 or 48 V). Consequently, power conversion is required to transfer the high-voltage ac utility input to the distributed low voltage loads. The concept of dc power delivery architectures has been explored for data centers recently, with the potential to offer higher reliability while employing fewer power conversion steps compared to conventional ac distribution [4], [5]. However, since no power conversion can be performed with 100% efficiency, regardless of the preferred power distribution architecture (ac or dc), large power conversion losses are inevitable in both architectures.

Due to the high power consumption of data centers, several solutions have been proposed in the recent literature to improve their energy efficiency. In [6], a traffic flow scheduling approach is designed for energy consumption optimization for data centers that have cloud computing systems. Another scheduling algorithm is proposed to minimize the amount of workload and avoid hot spots in data center networks by optimizing the job consolidation and traffic distribution to the servers [7]. In [8], the input voltages of series-connected web servers are regulated by scheduling incoming network load in software only. In addition to the software solutions mentioned in the recent literature, several hardware solutions have also been proposed. For example, in [9], thermal- and mechanical-oriented design of a server chassis, and in [10], a reduction of the number of power conversion stages are shown to reduce the power consumption and increase the energy efficiency of the servers. Today, mid-90% efficient power equipment is commercially available for data center applications [11], [12]. Certainly, improved conversion efficiency at each power stage can also be achieved, but this typically results in higher power electronics cost and/or volume. Power electronics already take up considerable space in modern servers, which makes it desirable to pursue alternative architectures that offer high efficiency and small overall volume. The work presented in this paper aims to improve the energy efficiency of data centers by proposing a new system-level solution, rather than focusing on efficiency improvement of each power element in the system. The main focus of the work presented here is that of energy conversion efficiency. While our proposed architecture achieves extreme efficiency, there are implications with regards to reliability and practical challenges such as

Manuscript received December 22, 2014; revised April 8, 2015; accepted July 20, 2015. Date of publication August 5, 2015; date of current version December 10, 2015. This work was supported in part by Texas Instruments and Google. This paper was presented in part at the INTELEC on October 2, 2014 in Vancouver, BC, Canada. Recommended for publication by Associate Editor M. Vitelli.

E. Candan and R. C. N. Pilawa-Podgurski are with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: candan2@illinois.edu; pilawa@illinois.edu).

P. S. Shenoy is with the Kilby Labs, Texas Instruments, Dallas, TX 75243 USA (e-mail: pshenoy@ieee.org).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TPEL.2015.2464805

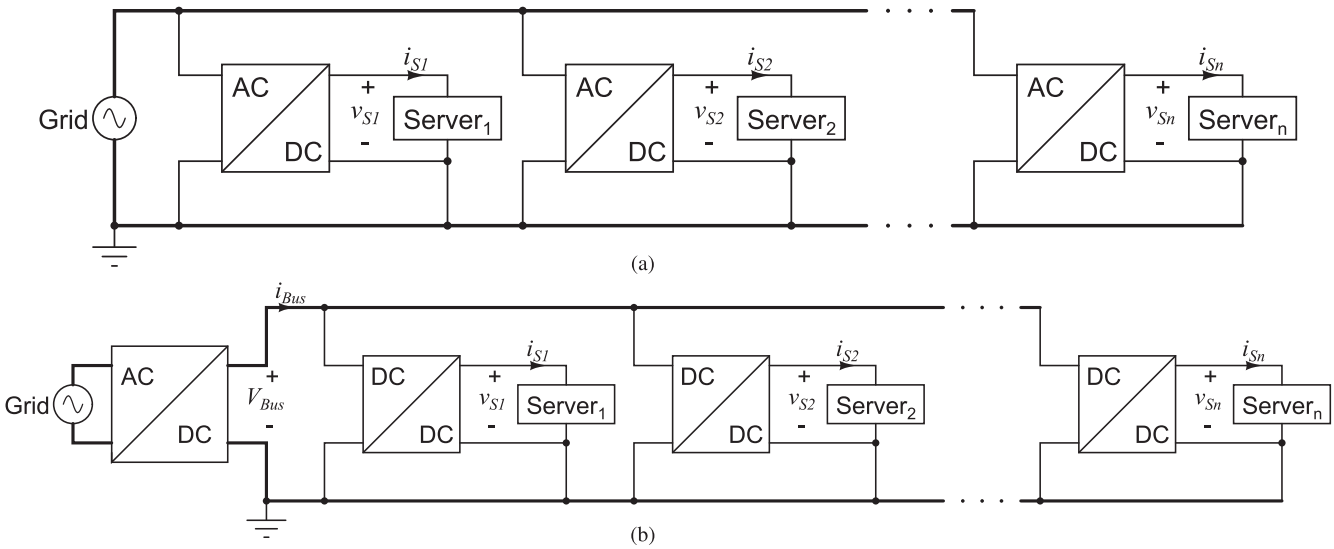

Fig. 1. Conventional power delivery stages in data centers. (a) AC power delivery stage for server racks. A single power converter provides rectification and step down to the 12 V server input. (b) DC power delivery stage for server racks. A central rectifying stage provides and intermediate dc voltage, followed by separate dc–dc converters, which step down the voltage to the 12 V server input.

hot-swapping and hardware protection, which are currently under investigation. Finally, we note that owing to the relatively fast replacement time of modern data centers (i.e., every 3–5 years), the architecture proposed here can be widely deployed in relatively short time as data centers are replaced.

The remainder of this paper is organized as follows: Section II states the research problem and provides background information. Section III introduces the concept of series stacking and differential power processing for server power delivery, and Section IV discusses the series-stacked system properties and the rack level control technique. Details of the differential power converter prototype designed for the proposed architecture are provided in Section V. Section VI provides the experimental results of both the proposed architecture, and as a comparison, a conventional architecture with a commercial power supply unit (PSU). Finally, Section VII concludes this paper.

II. Problem Statement and Background

The fundamental challenge, which is addressed in this paper is the coupling between the delivered power to the server and the conversion losses that occur at the PSU of the server. As the power ratings of modern servers are approaching kW ranges [13], even high-efficiency power converters can waste tens of watts at the power conversion stage.

A. Power Conversion Losses Throughout the System

In data centers, the utility power has to go through several power conversion stages before it reaches the servers. Shown in Fig. 1 are representative schematics of conventional power delivery architectures for data centers. Not shown in these figures are the uninterrupted power supply systems, power factor correction circuits, transformers, etc., typically present in these systems. Fig. 1(a) depicts a typical ac power distribution architecture, where each server has an ac–dc converter that provides the requisite dc power for the server through rectification and down conversion. In the dc power distribution architecture shown in Fig. 1(b), a central converter regulates a 380 V rectified grid dc bus ($V_{\rm Bus}$) voltage. The dc bus is fed to each server rack, and then, one dc–dc converter is installed on each server to convert the high voltage ($V_{\rm Bus}$) to a lower voltage for the servers, which is typically 12 V [13] or 48 V [14]. Although these architectures are well accepted in present data center designs, they have several drawbacks, such as limited system-level efficiency and large converter size. This is because each server's power converter has to perform a large voltage step down and process the full server power. Since each converter has to process the full server power during server operation, the direct coupling between delivered power and the associated power losses limits the system-level efficiency to that of individual power converter efficiency.

B. Previous Work

Large operating voltages are commonly encountered in dc systems, such as photovoltaic sources and battery systems, where the desired operation voltage of each element is far below the system dc bus voltage. By serially connecting these elements (sources or loads), high voltage step-up or step-down power conversion can be avoided [15], [16].

Recently, series-stacked systems have been explored in various fields, such as photovoltaics, where it has been shown that local, submodule maximum power point tracking can be achieved with the use of differential power processing (DPP) [17]–[20] with greatly improved system efficiency. In another application area, series stacking of two voltage regulators has been used to provide power to low-voltage biomedical implants in an efficient manner [21], [22]. Moreover, in order to deliver power to CMOS circuits, series stacking of the circuitry [23] and active voltage regulation with linear regulators to individual loads have been proposed [24], [25]. In such designs, the power delivery stage can also be implemented on chip, together with

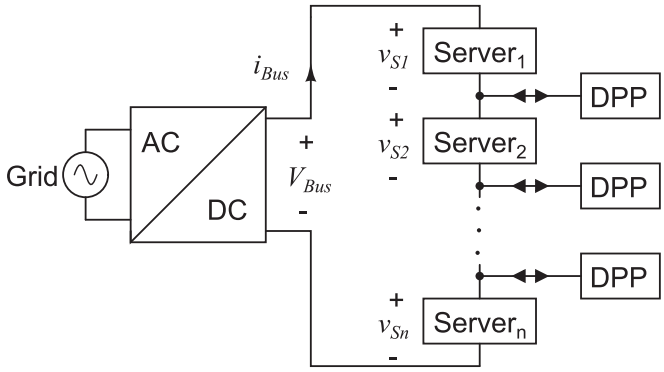

Fig. 2. Series stacking and DPP concept for server power delivery.

the series-connected logic loads [24]–[26]. Multilevel converter topologies have also been investigated for power delivery of vertically stacked cores [27].

In this paper, we seek to achieve greatly increased power delivery efficiency through the series connection of servers in data center applications. To this end, we present a hardware architecture that achieves significantly higher efficiency than state-of-the-art solutions. Through the use of DPP and a virtual bus architecture, we simultaneously achieve high conversion efficiency, small size, and excellent voltage regulation.

III. Series Stacking and DPP for Server Racks

In a series-stacked architecture, as shown in Fig. 2, all servers in a rack are connected electrically in series, rather than in parallel as in the conventional systems. Instead of employing a separate dc–dc converter for each server to convert the bus voltage ($V_{\rm Bus}$) to the server voltage, N servers are connected in series to equally share the bus voltage. For a suitable choice of N, the series-stacked architecture can provide inherent step down, where each server's input voltage is $V_{\rm Bus}/N$. However, all the series-connected servers conduct the same current ($i_{\rm Bus}$). This leads to a variation of server voltages even if there is only a small mismatch in power consumption between servers. Since the servers need to conduct the same amount of current, their input voltages have to vary to satisfy the instantaneously different load powers.

Regulated voltage of all servers in the series stack can be achieved through the use of bidirectional differential power converters, as illustrated in Fig. 2. Here, differential power converters are employed to maintain the allowed voltage band for the servers by injecting or rejecting the *difference* between the server current and the bus current. Providing instantaneous mismatch current between servers with differential converters can ensure that all servers in the rack operate in an allowed input voltage band.

Advantages of the configuration in Fig. 2 can be summarized as follows. Since the input voltage of the series-connected servers divide the bus voltage, the need for large voltage step down through external means (i.e., a step-down converter) is eliminated. Moreover, in the proposed architecture, the bulk power to the servers is delivered by the dc bus current that passes through all the servers in the series stack. Through series stacking and DPP, only the *difference* in power between servers needs to be processed. The amount of processed power, as well as the power lost during the power conversion can thus be reduced compared to conventional systems, where each server's dc–dc converter has to process the full power needed by the server. When care is taken to ensure a relatively uniform load distribution between servers in a stack, the proposed architecture yields a significant reduction in power conversion losses.

A. DPP Topology Selection

There are several possible techniques to implement DPP for series-connected elements. Common to all techniques is the series connection of all elements and the shared dc bus current for bulk power transfer. Differential power converters are connected to the intermediate nodes to create alternative current flow paths to inject or reject the mismatch current so that fine voltage regulation can be performed. In [16], different ways of performing DPP are explained in a systematic way for general dc systems. Here, DPP techniques are briefly explained considering server loads as series-connected elements.

Two of the main DPP techniques for series-connected servers, namely *server-to-server* and *server-to-bus*, are depicted in Fig. 3(a) and (b).

In the server-to-server DPP technique, alternative current flow paths are created between two neighbor servers by the differential converters, while in the server-to-bus DPP technique, an alternative current flow path for each server is created between the server and the dc bus.

In the server-to-server DPP architecture, the input and the output of the differential converters see a maximum of twice the server voltages, which enables the use of lower voltage switches in the differential converters. This leads to smaller converter size and better voltage regulation by increasing the switching frequency and control bandwidth in this architecture. However, when there is a mismatch in one of the servers in the series stack, all differential converters may need to work to deliver the required difference in current to the corresponding server. Therefore, the amount of processed differential power is heavily dependent on the load mismatch order of the series-stacked servers in the server-to-server DPP architecture.

In the server-to-bus DPP architecture of Fig 3(b), alternative current flow paths from the dc bus to an intermediate node in the series stack are provided by isolated converters, where the dc bus side of the differential converters need to be rated to the full dc bus voltage. In this architecture, the direct current flow path from the dc bus to each of the servers not only decouples the amount of processed differential power and the load mismatch order of the series-stacked servers, but also results in processing the lowest amount of power among the DPP architectures.

It is possible to achieve low-voltage-rated switches *and* reduced power processing with a slight modification to the server-to-bus architecture. For a given series-connected stack of servers, the secondary side of the differential power converters can be connected to a *virtual bus*. The virtual bus is essentially a capacitor bank that is isolated from the dc bus and connected in parallel with the secondary sides of the differential power

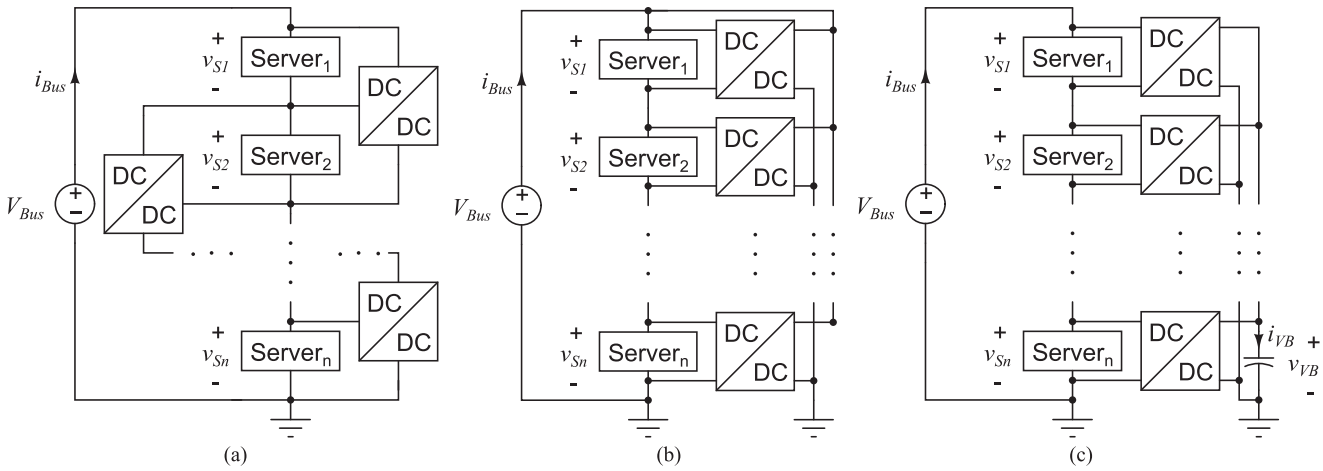

Fig. 3. DPP techniques for series-stacked server power delivery. (a) Server-to-server DPP. (b) Server-to-bus DPP. (c) Server-to-virtual bus DPP.

converters. This architecture is depicted in Fig. 3(c). Note that while the virtual bus capacitor is shown as a discrete element in Fig. 3(c), it can also be distributed among the secondary side capacitors of the differential power converters. A key benefit of the proposed virtual bus architecture is the scalability of the approach. Here, each DPP converter must only be rated for the server voltage and the virtual bus voltage, while the number of servers can be increased to accommodate a higher bus voltage if so desired. It should be noted that the virtual bus capacitor topology employed here is quite similar to active battery balancing [28], [29] using isolated power converters [30], [31]. Moreover, it has recently been successfully employed in DPP architectures for photovoltaic applications [19], [32].

In the server-to-virtual bus DPP architecture, the bulk power that is consumed by the servers is still provided by the dc bus, passing through the series connection in the stack. Each server in the stack has an isolated differential converter, whose secondary side is connected to a shared capacitor bank (the *virtual bus*). This capacitor bank acts as an energy reservoir to compensate for instantaneous power mismatch. This architecture thus enables any differential power converter in the series stack to exchange power with the others by injecting current to or rejecting current from the virtual bus.

Since there is no common ground between the servers in the series stack and the virtual bus, the differential converters still need to be isolated; however, the virtual bus voltage is an unrestricted design parameter, and can be chosen to be the same as the nominal server voltage. This enables a symmetric bidirectional isolated converter with the same high frequency, low-voltage switches in both the primary and secondary side of the differential power processors.

In addition, the power mismatch order of the series stack has no effect on the total amount of processed power as long as the virtual bus voltage is regulated within a limit. It should be noted that the virtual bus voltage regulation slightly increases the total processed power and the control complexity compared to the server-to-bus DPP architecture. However, as shown in this study, a suitable control implementation can achieve excellent regulation and efficient power transfer. A comparison of the three DPP architectures is summarized in Table I. In this study, the server-to-virtual bus DPP topology is explored. A hardware implementation of the server-to-server DPP can be found in [33].

B. Case Study

In a well-utilized data center (where each server is operating near its full capacity), the relative differences in power consumption between servers is expected to be small, yielding low overall power conversion in the proposed architecture. Shown in Table II is a comparison between a conventional power delivery architecture and the proposed approach. In this example, a 10 kW rack consisting of 32 servers is analyzed. It is assumed that the rack is powered from a 380 V dc bus, and that 96% efficient power converters are employed in the conventional dc architecture. This compares the proposed architecture with some of the highest efficiency values for existing data centers [5]. It is furthermore assumed an efficiency of 96% for the differential converters.

For an average computational load of 95% of full power, the rack consumes 9.5 kW for both architectures. In a conventional power delivery system, all of this power must be processed by the power converters, leading to a power conversion loss of 380 W. In contrast, the power that is processed in the server-to-virtual bus DPP architecture is only slightly higher than the power mismatch between servers ($\pm 5\%$, or ± 250 W). Owing to the Gaussian nature of the mismatch modeled here, the total power converted is 384 W. Thus, the resulting power loss is only 15 W in this scenario, leading to a substantially increased overall system efficiency. It should be noted that this power saving is further multiplied by the resulting reduction in cooling requirements, since less energy is wasted as heat.

Figure 4 illustrates in graphical form the server power consumption and corresponding differential power processed for the scenario indicated in Table II. The red bars indicate the server power consumption, the blue bars indicate the magnitude of the differential power processed between each server and the virtual bus, and the green line indicates average server power utilization. In the proposed series-stacked power delivery architecture, the bulk power flows with no power conversion, and no attendant loss. Only the mismatch power between servers

TABLE I
COMPARISON OF DPP ARCHITECTURES

	Server-to-Server	Server-to-Bus	Server-to-Virtual Bus
Topology	Nonisolated	Isolated	Isolated
Switch voltage rating	$2 \times V_{\text{Server}}$	V_{Server} and V_{Bus}	V_{Server}
Relative amount of processed power	High	Lowest	Low
Processed power mismatch order dependent	Yes	No	No

TABLE II
EFFICIENCY IMPROVEMENT WITH SERIES-STACKED ARCHITECTURE FOR A 10-KW RACK.

	Conventional Architecture	Series-Stacked Architecture
Average computational load	95%	95%
Computational load range	(90–100%)	(90–100%)
Rack power consumption [W]	9500	9500
Total power converted [W]	9500	384
Power converter efficiency	96%	96%
Power electronics loss [W]	380	15
Total system efficiency	96%	99.84%

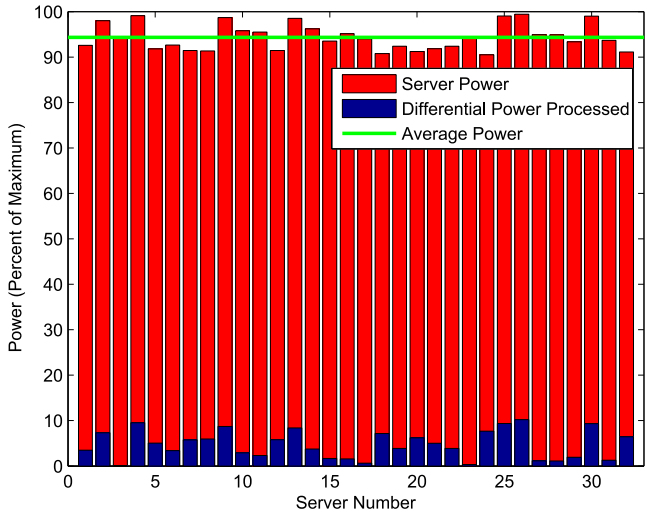

Fig. 4. Plot of server power consumption and differential power processed in a 10 kW, 32 server series-stack with ±5% average computational load imbalance.

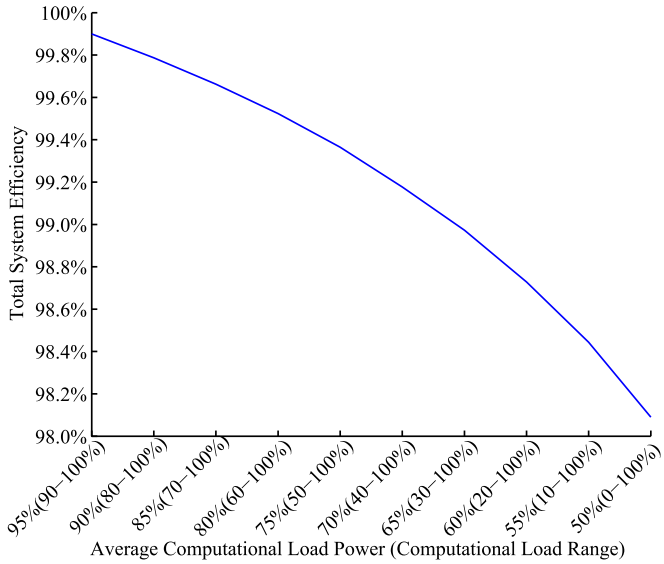

Fig. 5. Average computational load power versus calculated total system efficiency on the case study.

is processed through the differential power converters, yielding greatly reduced power conversion losses, and extreme conversion efficiencies.

The scenario illustrated in Fig. 4 is expanded to show the total system efficiency of the proposed architecture under different power mismatch scenarios. For this statistical example, the average computational load is swept between 95% and 50% of the full server power in 5% increments. At each average computational load point, the computational load range is varied with a Gaussian distribution for 1000 iterations, causing different power mismatch scenarios. The efficiency of each iterated scenario is calculated assuming 96% efficient differential converters and the total system efficiency is calculated by averaging the efficiency of each iterated scenario. Fig. 5 shows the total system efficiency for each average computational load power. Since the proposed architecture processes only the mismatch powers, the total system efficiency decreases as the computational load range (i.e., amount of mismatch power) increases. The worst-case mismatch scenario of this statistical study, which is 50% average computational load power with ±50% load range, results in a 98.09% theoretical efficiency.

IV. RACK LEVEL POWER CONTROL FOR SERIES-CONNECTED SERVERS AND VIRTUAL BUS

The objective of series stacking and DPP is to reduce power conversion losses by processing only the power mismatch between series-connected servers. In order to achieve the extreme efficiency desired, the power conversion needs to be avoided (i.e., the differential converters need to be OFF) whenever possible, as long as the server voltages are within a safe limit. A

system-level control method for the server-to-virtual bus DPP is briefly summarized in this section. The objective of the control is to regulate both the server and the virtual bus voltages simultaneously. A detailed description of the proposed control algorithm, as well as additional example scenarios, can be found in [34] and [35].

A. Light-Load Operation

When scheduling algorithms are used with the objective of spreading the computational load evenly among servers, the mismatch power between the series-connected servers as well as the processed power in DPP converters are typically far lower than the rated server power. On the other hand, to be able to maintain conventional server operation capabilities such as hot-swapping, the differential power converters must also be rated for the full server power. As a result, the differential power converters must be capable of working at light load with high efficiency (during normal operation) while also being rated for full server power (during abnormal operation).

Hysteresis control is a well-known control technique to increase the light-load efficiency [36]. To date, hysteresis control is well applied to unidirectional converters [37], [38]. However, relatively little attention has been given to bidirectional power flow applications. In this study, a bidirectional hysteresis control method is developed and used to regulate both individual server voltages and the virtual bus voltage simultaneously within a desired band.

B. Bidirectional Hysteresis Control

The server-to-virtual bus DPP architecture depicted in Fig. 3(c), where a stack of N servers are connected in series, employs N differential converters. However, there are $N+1$ independent voltages to regulate since the voltage of the virtual bus is also free to vary. Controlling $N+1$ voltage domains with N controllers is challenging; however, it should be noted that the $N+1$ voltage domains are not independent, but coupled through their series connection. Our proposed control technique exploits this fact. The control objective is to regulate all individual server voltages within a predefined hysteresis band, while ensuring that there is always energy in the virtual bus capacitor. Each differential converter is thus tasked with regulating both its input (i.e., virtual bus) and output (i.e., server) voltage. This requires the hysteresis control to be bidirectional, where the direction of power flow needs to be determined dynamically, depending on both the input and output voltage of the converters. In the implementation of this control objective, only voltage measurements are desired for high efficiency and simplicity. Local converter control is also desired since a central controller requires communication between differential converters in the series stack, leading to an increase in the cost and complexity of the system.

Figure 6 shows a schematic drawing of a four-server system, which will be used as an example system in this control discussion. Each isolated differential power converter is modeled as an isolated current amplifier with two current sources connected with opposite polarity. In Fig. 6, v_{Si} represents the input voltage

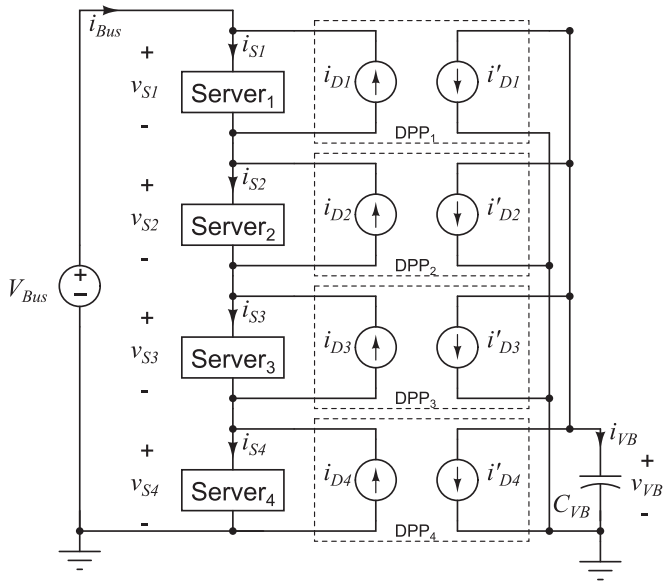

Fig. 6. Server-to-virtual bus DPP model with current sources.

of the ith server, where $i = 1, 2, 3, 4$. Moreover, i_{Si} represents the current consumed by the ith server, i_{Di} represents the injected server current, i'_{Di} represents the current rejected from the virtual bus, v_{VB} and i_{VB} represent the virtual bus voltage and current, respectively, and V_{Bus} and i_{Bus} represent the bus voltage and current, respectively. For this four-server rack model, V_{Bus} is assumed to be a well-regulated dc bus, and i_{Bus} is the current that provides the bulk power to the series-connected servers in the rack. Depending on the load of the ith server, i_{Si} will change, which in turn causes the input voltages of the servers to vary to satisfy the instantaneous power requirements.

In order to achieve the control objectives, the properties and constraints of the server-to-virtual bus DPP architecture need to be well understood. The first and most basic constraint of the system is the voltage constraint of the series stack, which is given for the server stack in Fig. 6 by

$$V_{\text{Bus}} = \sum_{i=1}^{4} v_{Si}(t). \tag{1}$$

As in Fig. 1(b), the dc bus is regulated by a central converter, as well as supplied with capacitor banks in order to provide a steady bus voltage. In practice, this means that its time constant is larger than the other dynamics in the series-stacked system, and its value can be assumed to be constant with respect to the instantaneous server voltages (v_{si}) and virtual bus voltage (v_{VB}) in this discussion.

In Fig. 6, the sum of KCLs at each server node results in

$$\begin{aligned} i_{\text{Bus}}(t) &= i_{S1}(t) - i_{D1}(t) \\ &= i_{S2}(t) - i_{D2}(t) \\ &= i_{S3}(t) - i_{D3}(t) \\ &= i_{S4}(t) - i_{D4}(t). \end{aligned} \tag{2}$$

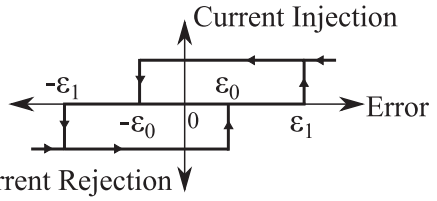

Fig. 7. Proposed hysteresis shape for bidirectional control.

TABLE III
SUMMARY TABLE FOR BIDIRECTIONAL HYSTERESIS CONTROL

Previous Decision	Present ε Range	Present Decision
Current rejection	$< \varepsilon < -\varepsilon_1$	Current rejection
Current rejection	$-\varepsilon_1 < \varepsilon < \varepsilon_0$	Current rejection
Current rejection	$\varepsilon_0 < \varepsilon <$	No action
No action	$\varepsilon_1 < \varepsilon <$	Current injection
Current injection	$\varepsilon_1 < \varepsilon <$	Current injection
Current injection	$-\varepsilon_0 < \varepsilon < \varepsilon_1$	Current injection
Current injection	$< \varepsilon < -\varepsilon_0$	No action
No action	$< \varepsilon < -\varepsilon_1$	Current rejection
No action	$-\varepsilon_1 < \varepsilon < \varepsilon_1$	No action

The second constraint is the virtual bus voltage $v_{\rm VB}(t)$, given by

$$v_{\rm VB}(t) = v_{\rm VB}(t_0) + \frac{1}{C_{\rm VB}} \int_{t_0}^{t} i_{\rm VB}(\tau) d\tau \tag{3}$$

where $i_{\rm VB}$ is the instantaneous sum of all differential currents, given by

$$i_{\rm VB}(t) = -\sum_{i=1}^{4} i'_{Di}. \tag{4}$$

In order to ensure that the virtual bus is well regulated, (3) and (4) imply that the net sum of the differential currents flowing into the virtual bus should be within a limit in any arbitrary time period. Let $I_{\rm VB} =< i_{\rm VB} >$ in this arbitrary time period; then, having $I_{\rm VB} > 0$ results in an increase in virtual bus voltage, while $I_{\rm VB} < 0$ results in a decrease in virtual bus voltage.

In order to determine whether any of the voltage domains (servers and virtual bus) in the system requires current injection or current rejection, at every sampling period, each DPP converter samples both the corresponding server voltage (v_{Si} in Fig. 6) and the virtual bus voltage ($v_{\rm VB}$ in Fig. 6). Then, the sampled voltages are compared with a reference voltage and errors are calculated by (5).

$$\varepsilon = V_{\rm ref} - v_{\rm sampled} \tag{5}$$

where $v_{\rm sampled}$ corresponds to the sampled voltages (v_{Si} or $v_{\rm VB}$) and $V_{\rm ref}$ corresponds to the nominal server voltage.

A proposed bidirectional hysteresis control algorithm is shown visually in Fig. 7, which comprises two hysteresis voltages (ε_0 and ε_1). A description of the corresponding controller decisions is provided in Table III. The control method of Fig. 7 and Table III is general in nature, and applies to both the converter input and output.

TABLE IV
DECISION TABLE FOR BIDIRECTIONAL HYSTERESIS CONTROL

		Virtual Bus		
		No Action	Injection	Rejection
Server	No Action	OFF	$-I_{D\,i}$	$+I_{D\,i}$
	Injection	$+I_{D\,i}$	OFF	$+I_{D\,i}$
	Rejection	$-I_{D\,i}$	$-I_{D\,i}$	OFF

The proposed hysteresis control method, thus, requires knowledge of the previous decision, so the decision made during the previous sampling period must be saved in memory in order to make the correct decision in the present sampling period. After the error is calculated for both the server and the virtual bus voltage, the present decision about current requirement and direction is made by both considering the previous decision and referring to the hysteresis shape in Fig. 7. For example, if the calculated error for a voltage domain is between ε_1 and $-\varepsilon_0$ while the previous decision for that voltage domain was not to inject or reject any current to the voltage domain, it means that this voltage domain is still within the predefined limits. Therefore, it does not require current injection or rejection. However, if the previous decision was current injection, then present decision must be current injection as well for the same error range.

Once the current decisions for all voltage nodes in the system have been made for that sampling time, the controller should determine if the corresponding differential power converter needs to be ON or OFF. If it must be ON, the power flow direction should be determined. Note that each differential power converter is responsible for both its input and output voltage, so the decision made by the controller must meet the needs of both voltage domains. Since for each differential converter, there are two voltage domains (i.e., the server and the virtual bus) and three possible requirements (i.e., current injection, current rejection, and no action), there are nine possible states for each differential power converter at any sampling period. These nine possible states and the decisions for each state are summarized in Table IV, with $+I_{Di}$ corresponding to average current being injected to the server ($+ < i_{Di} >$) in that sampling time, and $-I_{Di}$ corresponding to the current being removed from the server ($- < i_{Di} >$) in that sampling time (i.e., injected to the virtual bus). It should be noted that the decision based on the rows of Table IV may be different for each differential converter, since each server voltage is unique. However, the decision based on the columns of Table IV will be the same for all differential converters, since they share the same information and requirement on the virtual bus.

As an example, consider one of the differential power converters in the stack ($\rm DPP_i$ in Fig. 6). Both the server and the virtual bus voltage are sampled and the errors are calculated according to (5). Then, the needs of both voltage domains are determined by using the hysteresis shape in Fig. 7 (or by referring to Table III). If, for example, neither the server nor virtual bus requires current injection or rejection (no action), the corresponding control decision, according to Table IV, is to turn

DPP$_i$ OFF. For that sampling period, 100% power conversion efficiency is achieved for the ith server since DPP$_i$ is OFF and all the power consumed by the ith server is delivered through the series connection by the dc bus.

As seen in Table IV, as long as the ith server and the virtual bus do not require current injection (or rejection) at the same time, an appropriate current injection (or rejection) decision can be made in the DPP$_i$, in order to satisfy the need of one of the voltage domains. In this case, this decision will not adversely affect the other voltage domain. On the other hand, when both the ith server and the virtual bus require current injection (or rejection) at the same time, the appropriate decision is to keep the DPP$_i$ OFF, as shown in Table IV. This control decision relies on the series-stacked system properties to achieve the control action, as explained earlier.

As an example, consider the case where the ith server and the virtual bus require current injection at the same sampling time, which means v_{Si} and $v_{\rm VB}$ are both lower than their nominal values. In this case, since (1) enforces the sum of all series-stacked server voltages to be fixed, there must be at least one other server in the series stack that has a voltage higher than its nominal value. We denote this server the jth server. According to Table IV, the decision made by DPP$_j$ is then to inject current to the virtual bus, which satisfies the requirement of the virtual bus at that sampling time. Also, since (2) is valid at every server node, the current injection to the virtual bus by the DPP$_j$ increases the bus current ($i_{\rm Bus}$) by the same amount that the DPP$_i$ would provide if it was ON and injecting current to the ith server. This increase in $i_{\rm Bus}$ thus satisfies the current injection requirement of the ith server and increases v_{Si}. The same approach is valid when the ith server and the virtual bus require current rejection at the same sampling time. Consequently, when both the ith server and the virtual bus have the same demands, keeping the DPP$_i$ OFF and relying on series-stacked system properties, not only satisfy the need of both voltage domains but also achieves 100% power conversion efficiency for the ith server during that sampling period since all the current needed by the ith server is provided by $i_{\rm Bus}$ without being processed.

C. Design Considerations

A four-server system will be used for hardware verification of the server-to-virtual bus DPP architecture. The implementation of the system requires consideration of the size of the virtual bus capacitor ($C_{\rm VB}$) along with the hysteresis bands (ε_0 and ε_1) since these parameters determine how often the differential converters need to be ON. Also, the amount of power injected or rejected by the differential converter (i_{D1}) and the sampling time of the control algorithm are important parameters since they affect the amount of energy lost in the differential converters. The main focus of this paper is the proof of concept server-to-virtual bus DPP demonstration. The detailed characterization and optimization of these parameters for given computational load profiles will be more thoroughly explored in future work. Here, the design and parameters for the virtual bus capacitor size, the hysteresis bands, the magnitude of the differential current, and the sampling time are briefly explained.

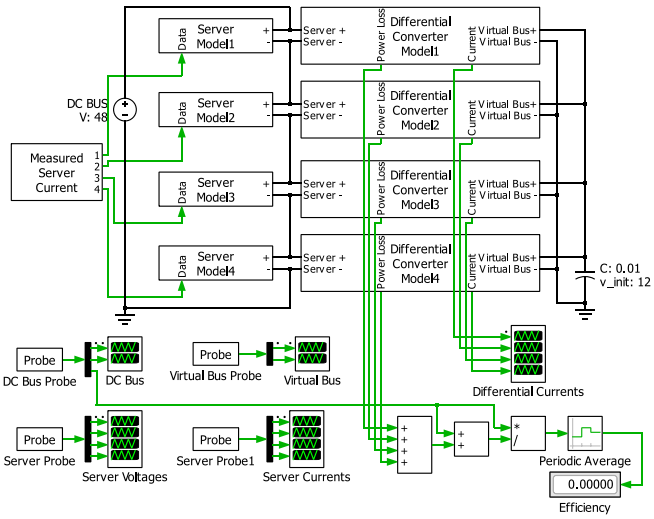

Fig. 8. Simulation schematic in PLECS.

The bidirectional hysteresis control used in this paper injects or rejects current to the servers when needed. In order to reduce the power loss of the differential power converters, the differential current magnitude is chosen to correspond to the value where the differential converter has its peak efficiency. Then, the sampling time for the control algorithm is chosen approximately 100 times the switching period, allowing enough time for the differential current to reach its steady-state value and to have a measurable effect on the system. The input voltage specification of the servers that will be used in the experiment is 12 V $\pm$5%($\pm$0.6 V). With an additional safety margin, the hysteresis bands are selected as $\varepsilon_1 = 0.4$ V and $\varepsilon_0 = 0.2$ V for servers. The voltage swing on the virtual bus capacitor can be more relaxed; therefore, the virtual bus hysteresis bands are selected as $\varepsilon_1 = 0.6$ V and $\varepsilon_0 = 0.3$ V. After these parameters are fixed, the effect of increased virtual bus capacitor size is investigated in simulation.

Shown in Fig. 8 is the simulation model of the four-server system. Each server model in Fig. 8 uses recorded server currents while the server is performing web traffic management, providing a realistic simulation scenario. A mathematical model for 95% efficient bidirectional converters is used to execute the proposed control algorithm, while also capturing the instantaneous power loss. The system-level efficiency is calculated by $\sum p_{\rm server}/\left(\sum p_{\rm server} + \sum p_{\rm loss}\right)$ and its average is displayed at the end of the simulation. In the simulation study, the hysteresis bands for the servers are $\varepsilon_1 = 0.4$ V and $\varepsilon_0 = 0.2$ V, for the virtual bus are $\varepsilon_1 = 0.6$ V and $\varepsilon_0 = 0.3$ V. Also, the sampling time for the control algorithm is chosen as 500 μs and the magnitude of the differential current is 2 A.

A plot illustrating the effect of increased virtual bus capacitor size on averaged system-level efficiency is given in Fig. 9. As the virtual bus capacitance size increases, the series-stacked system requires less differential effort for regulation, yielding higher system-level efficiency. However, as shown in Fig. 9, the system-level efficiency saturates around $C_{\rm VB} = 80$ mF and does not increase dramatically for the simulated load mismatch scenario. Considering the possible nonidealities in the real test-bed

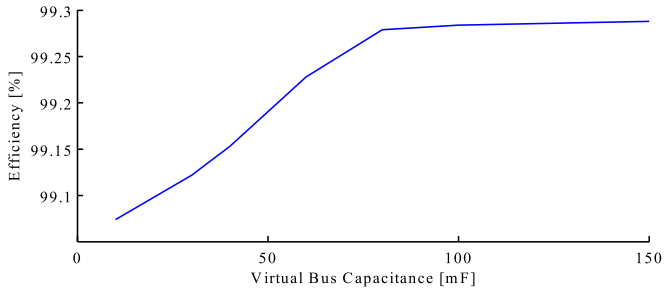

Fig. 9. Virtual bus capacitor size versus the system-level efficiency.

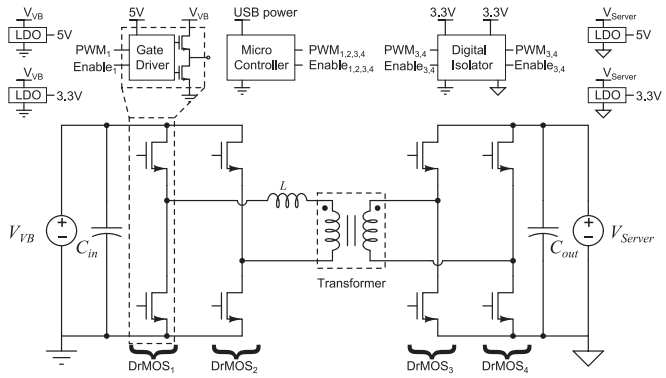

Fig. 10. Annotated schematic of DPP converter.

and testing the proposed architecture under different load mismatch scenarios, the following experimental studies use a 100 mF virtual bus capacitor. Further optimization of the hysteresis bands and the virtual bus capacitor sizing is possible, and will be explored in follow-up work.

V. Prototype Converter for Server-to-Virtual Bus DPP Architecture

In order to validate the series-stacked system with the proposed server-to-virtual bus DPP control method, isolated bidirectional differential power converters are needed. In the server-to-virtual bus DPP architecture, the differential power converters are rated for the same voltage at both sides of the transformer. The dual active bridge (DAB) converter topology offers bidirectional power flow and symmetrical switch voltage ratings [39], [40]. In this study, phase-shift modulation is used to control the power flow. Four prototypes of 120 W 12 V–12 V DAB converters have been designed. A schematic diagram and an annotated photograph of the prototype converter is shown in Figs. 10 and 12, respectively. Key specifications and selected components of the prototype converter are given in Tables V and VI, respectively.

The efficiency of the prototype converter is plotted in Fig. 11 for both power flow directions. As expected, the symmetric design of the DAB offers almost identical efficiency curves for both directions of power flow. The peak efficiency of 95% is achieved for both directions of power flow around 35 W.

In the experimental testing performed here, the series-stack was initialized manually, using an external, isolated power supply to ensure a stable initial condition for the servers, and to reduce any in-rush current at startup.

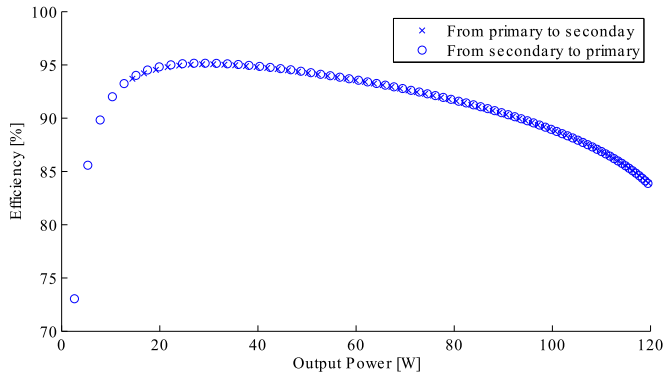

Fig. 11. Efficiency versus output power of the DPP converters.

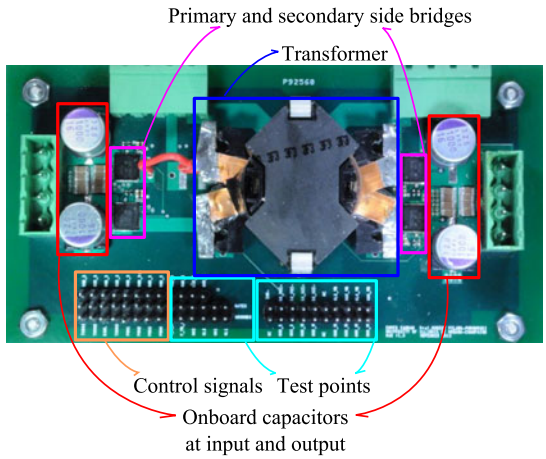

Fig. 12. 120 W 12 V–12 V DAB converter.

TABLE V
DAB Converter Specifications

Rated Power	120 W
Peak Efficiency	95%
Switching Frequency	175 kHz
Modulation Technique	Phase shift
Control Mode	Bidirectional Hysteresis

VI. Experimental Setup and Results

In order to experimentally validate the server-to-virtual bus DPP architecture for data center power delivery, a four-server rack was created in the laboratory. In addition, a conventional dc power delivery architecture with a best-in-class PSU was tested to demonstrate the improvement in the energy efficiency with the proposed architecture compared to the state-of-the-art solution. This section explains the experimental setup for both architectures, gives the results for both, and compares the energy efficiencies.

A. Test Bed

A flexible and modular laboratory test bed was developed as shown in Fig. 13. Four Dell Optiplex SX775 Core 2 Duo

TABLE VI
DAB CONVERTER KEY COMPONENTS

Switches	DrMOS - Vishay SiC780ACD
Digital Isolator	TI - ISO7241C
LDOs	TI - LP2985
Microcontroller	TI - C2000 Piccolo
Transformer	1:1 custom design; 3F3, RM12 core
Windings	0.500 in $\times$ 0.005 in copper foil, 6 turns
Inductor	Leakage inductance of the transformer
C_{in} and C_{out} Ceramics	$6 \times 10\mu$F 16 V X5R, Samsung
C_{in} and C_{out} Aluminum	2×1 mF 16 V SMD, Panasonic

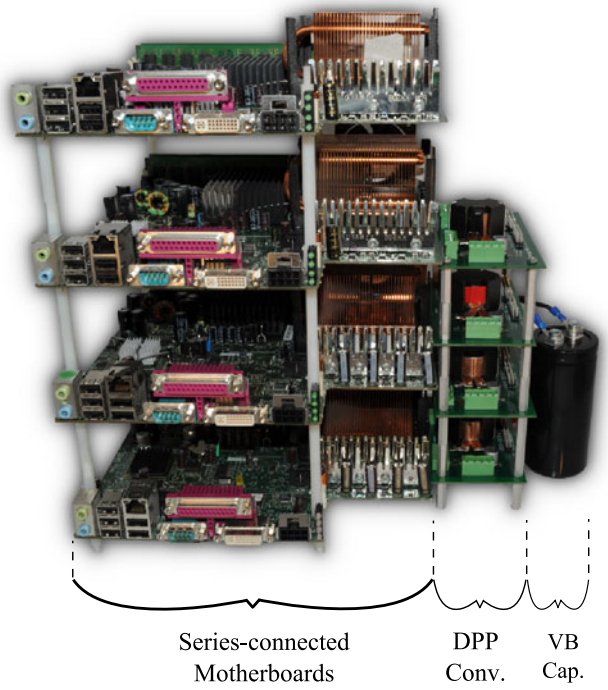

Fig. 13. Photograph of the modular laboratory test-bed.

computers are used as servers, chosen specifically for their single 12 V motherboard input, similar to what is used in modern data centers. While the servers used in the experimental evaluation here are of lower power rating than what is typically encountered in data centers, they behave similarly to high-power server systems in terms of response to computational loads, albeit at reduced absolute current values. In this regard, they serve as scaled down versions of the costly servers typically encountered in data centers, and enable us to explore the series-stacked concept without expensive custom server hardware. All servers run Ubuntu server images, initialized from a common network image. A more detailed description of the test bed can be found in [33].

To demonstrate the applicability of the proposed architecture on real-world workloads, the servers were operated to represent two separate scenarios: web traffic management and computation. In the web traffic test, a python script is used to generate high rate web traffic over the Ethernet network. It is generated by seven computers at a total rate of 700 requests per second and sent to all four servers for approximately 140 s. This test represents a more bursty power consumption profile with correspondingly large instantaneous power mismatch between servers. In addition to the web traffic test, a computational load test is executed. In this test, the standard Linux "stress" utility is used as a computational load [41], which loads the CPU, disk, and memory by performing repeated numerical computations and memory writes and reads. This test is run on the servers for approximately 100 s.

A data acquisition unit from National Instruments is used to simultaneously sample both individual server voltages and currents, as well as bus and virtual bus voltages and currents at 5000 samples per second (200 μs sampling period). Current sensors are implemented with sense resistors and operational amplifiers, and the voltage across the current sensors are sampled with the same data acquisition unit. The output voltage of the current sensors are calibrated by an Agilent 34410A 6 1/2 digit digital multimeter. Identical tests and measurements are performed on the same servers, using the proposed series-stacked and conventional dc power delivery architecture in order to have a fair comparison. The schematics for both setups are shown in Figs. 14 and 15. In both setups, the bus voltage is 48 V, corresponding to standard telecommunication infrastructure. An advantage of the server-to-virtual bus architecture is that it can be applied to a system with higher bus voltage without modification to the differential converters.

B. Results

1) Proposed Series-Stacked and Server-to-Virtual Bus DPP Architecture: Four DAB converters are designed as explained in the previous section and a 100 mF capacitor is employed as the virtual bus capacitor. For both the web traffic and the computational loads, all server voltages and the virtual bus voltage are regulated to predefined hysteresis bands for the duration of the tests. Instantaneous server currents as well as the server and the virtual bus voltages for a representative 200 ms snapshot are given in Figs. 16 and 17 for web traffic load, and in Figs. 18 and 19 for computational load.

Despite the large instantaneous variations in server currents in both tests as shown in Figs. 16 and 18, the server voltages as well as the virtual bus voltage can be regulated by using the proposed bidirectional hysteresis control. Moreover, the active voltage regulation is performed without processing of the full server power during both tests. The key advantage of DPP is reducing the amount of processed power by processing only the differences between individual server power consumptions.

During the web traffic test, the average input power to the series stack is 241.09 W, and the average output power is 237.98 W. This corresponds to a 98.71% efficiency in power conversion. Likewise, during the computational load test, the average input power to the series stack is 426.60 W, and the average output power is 426.11 W. This corresponds to a 99.89% efficiency in power conversion. As expected, the computational load test, which represents a uniform power consumption for all servers (i.e., small power mismatch), yields the highest efficiency.

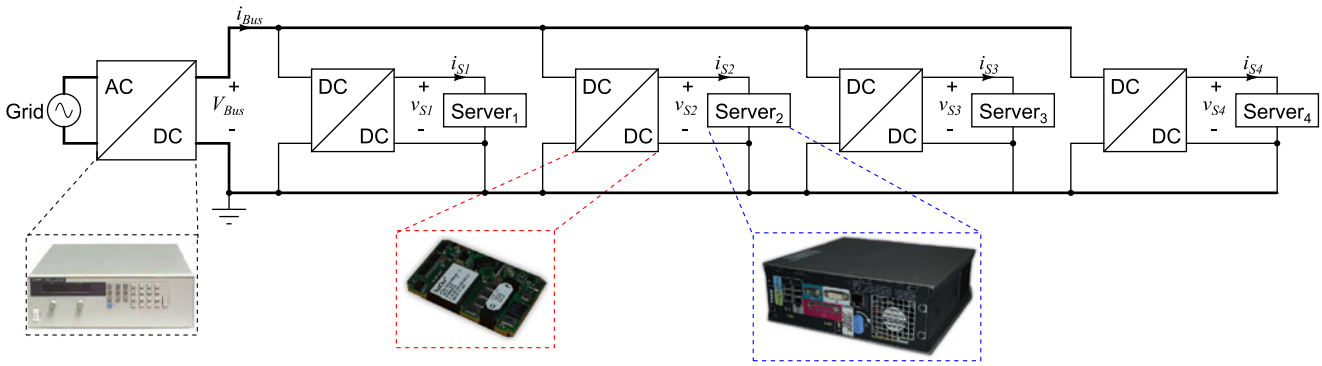

Fig. 14. Schematic of the experimental setup to compare the conventional architecture.

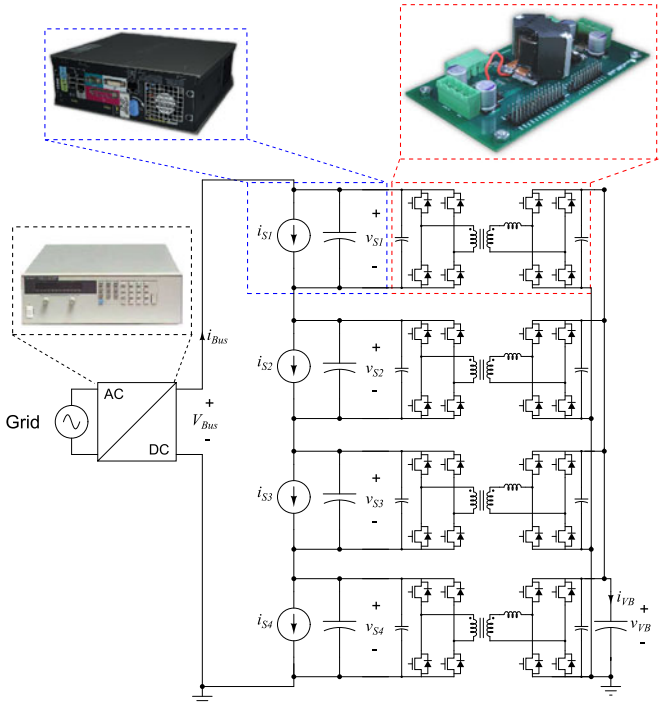

Fig. 15. Schematic of the experimental setup to test the proposed architecture.

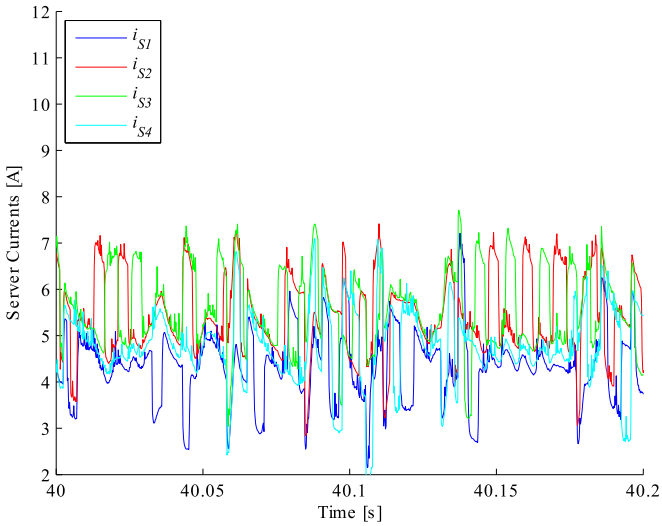

Fig. 16. Typical server current waveforms during the web traffic test.

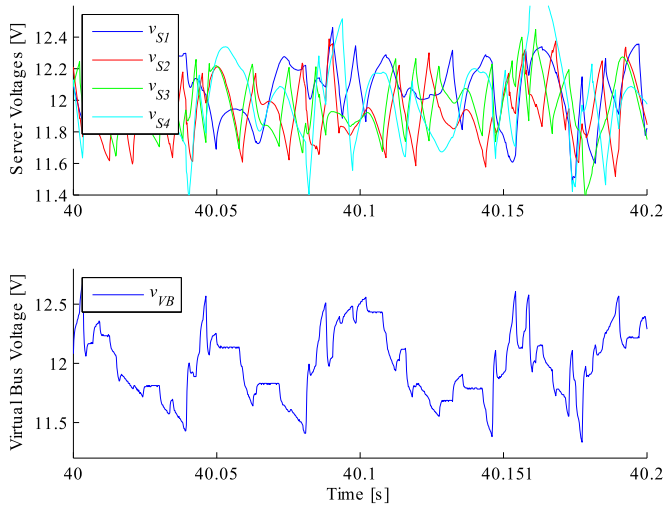

Fig. 17. Typical server and virtual bus voltage waveforms during the web traffic test.

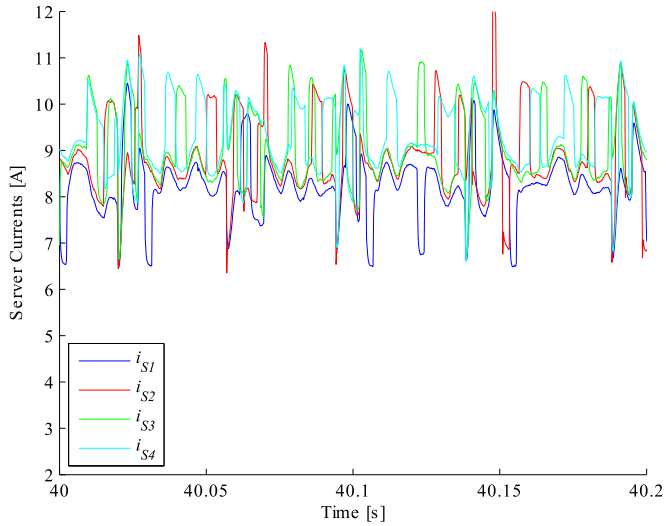

Fig. 18. Typical server current waveforms during the computation test.

2) Conventional DC Power Delivery Architecture With a Best-in-Class PSU: In order to illustrate the effectiveness of series stacking and DPP as an alternative power conversion and delivery architecture for server racks, identical tests are

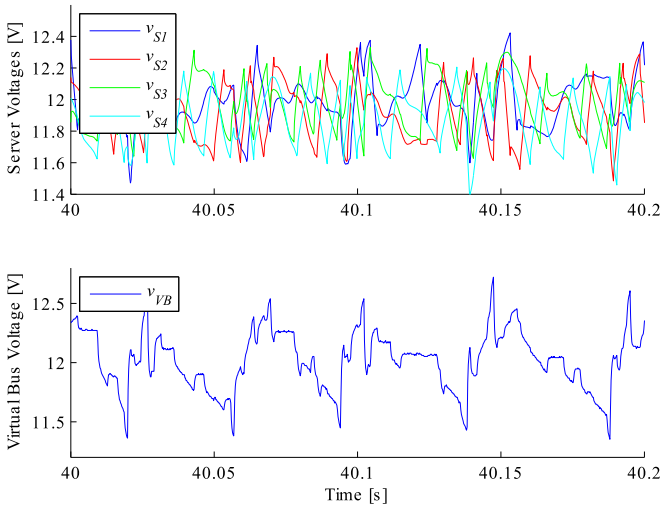

Fig. 19. Typical server and virtual bus voltage waveforms during the computation test.

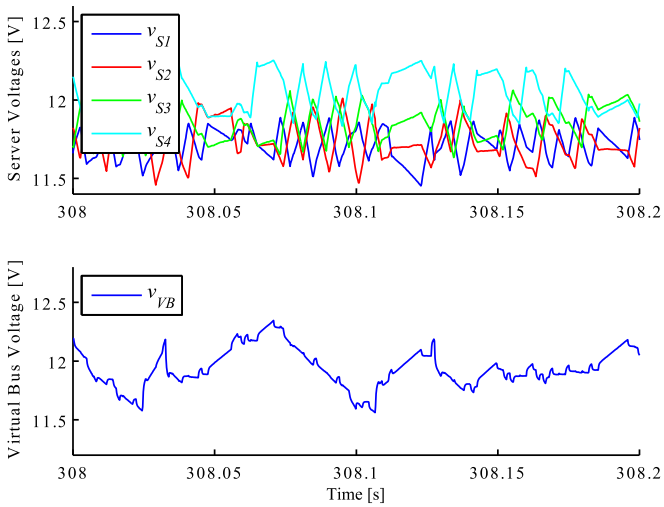

Fig. 20. Typical server and virtual bus voltage waveforms under severely unbalanced series-stacked loads.

TABLE VII
COMPARISON OF PROPOSED ARCHITECTURE WITH CONVENTIONAL ARCHITECTURE

	Web Traffic Test		Computation Test	
	Proposed	Conventional	Proposed	Conventional
$< P_{\text{in}} >$ [W]	241.09	252.87	426.60	447.59
$< P_{\text{out}} >$ [W]	237.98	238.58	426.11	426.51
$< P_{\text{loss}} >$ [W]	3.11	14.29	0.49	21.08
Efficiency [%]	98.71	94.35	99.89	95.29

performed while each server is powered with a 48–12 V step-down dc–dc converter as shown in Fig. 14. The dc–dc converter shown in Fig. 14 and used in these tests (PQ60120QEx25 from SynQor) is considered as a best-in-class PSU for telecommunication equipments with 96% peak efficiency [11]. In addition to being the highest efficiency commercial solution found by the authors, the chosen PSU is also the one used in Bluewaters, a recent supercomputer with modern hardware [42].

During the web traffic test, the average input power to the server stack in this setup is 252.87 W, and the average output power is 238.58 W. This corresponds to a 94.35% efficiency in power conversion. Likewise, during the computational load test, the average input power to the server stack is 447.59 W, and the average output power is 426.51 W. This corresponds to a 95.29% efficiency in power conversion. It should be noted that all measurements are performed with the same measurement unit regardless of the power delivery architecture.

C. Comparison of Conventional and Proposed Architectures

The experimental comparison of the proposed server-to-virtual bus DPP architecture with the conventional dc power distribution architecture for a four-server system is summarized in Table VII. The efficiency of the proposed architecture greatly overcomes the efficiency of the conventional architecture with a state-of-the-art solution in both tests. In the proposed architecture the converters are kept OFF whenever possible in order to maximize the power delivery efficiency. For the same web traffic load to a four-server rack, a more than four times reduction in power loss is achieved. Moreover, for the same computational load in four servers, more than 40 times reduction in power loss is achieved in the power delivery stage, when the server-to-virtual bus DPP architecture is compared to the conventional dc power distribution architecture with a best-in-class PSU.

D. Severe Mismatch Conditions

In Section VI-B, the proposed series-stacked servers have been loaded with example scenarios where the same web traffic and computational loads are used, thus, the series-stacked servers have somewhat similar power consumption. This example scenario not only helped the differential converters to process significantly less power, (which resulted in 99.89% overall efficiency) but also demonstrated how minimizing the load difference between the series-stacked servers can boost the system level efficiency. This shows the importance of the load-balancing algorithms for series-stacked systems; nevertheless, identical loads are not always encountered by the natural behavior of the system. In order to show that the proposed architecture is capable of providing power to unbalanced series-stacked loads with high efficiency, another experiment is conducted while the series-stacked servers are operating under unbalanced loads and the results are presented here. Experimental results and implementation details of a hot-swapped operation (i.e., one of the servers is completely isolated from the series-stack while the others are operational) can be found in [43].

Shown in Figs. 21 and 20 are the instantaneous server currents as well as the server and the virtual bus voltages, where one of the servers is in idle, while the others are operational. Despite the unbalanced series-stacked loads, the server voltages as well as the virtual bus voltage can still be regulated without processing the full server power. During the unbalanced load conditions, the average input and output power of the server

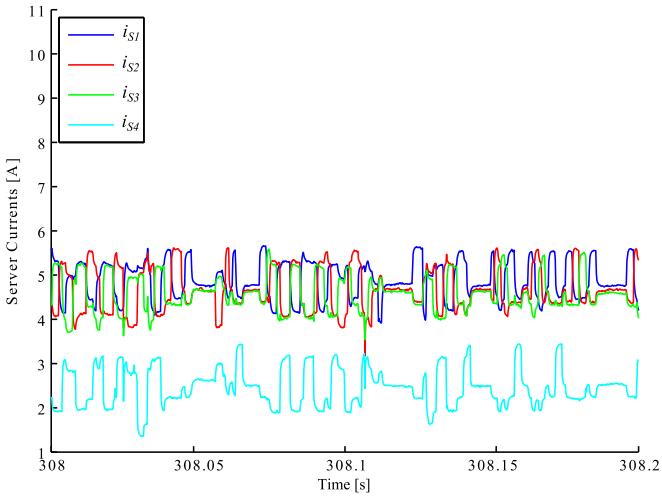

Fig. 21. Typical server current waveforms under severely unbalanced series-stacked loads.

stack are 119.73 and 118.20 W, respectively. This corresponds to a 98.72% efficiency in power conversion stage.

VII. Conclusion

The work presented in this paper is a successful hardware demonstration of a dc power distribution and conversion architecture for high-efficiency data centers. In the presented architecture, servers are connected in series, rather than in parallel as in the conventional architecture, and server-to-virtual bus DPP technique is used to actively regulate individual server voltages. In this technique, the amount of processed power is greatly reduced in comparison to the conventional architecture with a correspondingly large reduction in power loss in the power distribution and conversion stage.

Isolated bidirectional differential converters are designed as DAB dc–dc converters and rated for full server power. In order to regulate the series-stacked server voltages and the virtual bus voltage with high efficiency, bidirectional hysteresis control is developed. The proposed architecture is experimentally validated and compared with the conventional dc power distribution architecture. A modular experimental test bed is created with four 12 V servers, which operate two different real-world applications: web traffic management and computation. The same four-server rack is powered with both the proposed architecture and the conventional architecture with a best-in-class PSU for servers.

Identical tests and measurements are performed with both architectures. The four server rack system achieves up to 99.89% efficiency with the proposed architecture, which corresponds to 40x reduction in average power conversion losses in comparison to the conventional architecture with state-of-the-art hardware.

Acknowledgment

The authors would like to thank J. McClurg for his assistance with some experimental measurement design.

References

[1] E. Masanet, R. Brown, A. Shehabi, J. Koomey, and B. Nordman, "Estimating the energy use and efficiency potential of U.S. data centers," *Proc. IEEE*, vol. 99, no. 8, pp. 1440–1453, Aug. 2011.

[2] J. Koomey. (2011 Aug). Growth in data center electricity use 2005 to 2010. Analytics Press. Oakland. CA, USA. [Online]. Available: http://www.analyticspress.com/datacenters.html

[3] J. Whitney and P. Delforge, "Data center efficiency assessment," Natural Resources Defense Council, Tech. Rep. 14-08-A, Aug. 2014.

[4] P. Krein, "A discussion of data center power challenges across the system," in *Proc. Int. Conf. Energy Aware Comput.*, Dec 2010, pp. 1–3.

[5] M. Ton, B. Fortenbery, and W. Tschudi, "DC power for improved data center efficiency," Lawrence Berkeley National Laboratory, Berkley, CA, USA, Tech. Rep., Mar. 2008.

[6] D. Kliazovich, S. Arzo, F. Granelli, P. Bouvry, and S. Khan, "E-stab: Energy-efficient scheduling for cloud computing applications with traffic load balancing," in *Proc. IEEE Int. Conf. Green Comput. Commun./ Int. Conf. Cyber, Phys. Soc. Comput.*, Aug. 2013, pp. 7–13.

[7] D. Kliazovich, P. Bouvry, and S. Khan, "Dens: Data center energy-efficient network-aware scheduling," in *Proc. 2010 IEEE/ACM Int. Conf. Green Comput. Commun./Int. Conf. Cyber, Phys. Soc. Comput.*, Dec. 2010, pp. 69–75.

[8] J. McClurg, Y. Zhang, J. Wheeler, and R. Pilawa-Podgurski, "Re-thinking data center power delivery: Regulating series-connected voltage domains in software," in *Proc. IEEE Power Energy Conf. Illinois*, Feb. 2013, pp. 147–154.

[9] E. Frachtenberg, A. Heydari, H. Li, A. Michael, J. Na, A. Nisbet, and P. Sarti, "High-efficiency server design," in *Proc. Int. Conf. High Perform. Comput. Netw. Storage Anal.*, Nov. 2011, pp. 1–11.

[10] Efficiency: How we do it. (2014). [Online]. Available: http://www.google.com/about/datacenters/efficiency/internal/#servers

[11] *Technical Specification PQ60120QEx25*. (2013, Sep.). SynQor. [Online]. Available: www.synqor.com

[12] P. Eve, F. R. Gomez, and E. Gao, "Trial of super high efficiency rectifiers (98%) in telecom networks," in *Proc. IEEE Int. Commun. Energy Conf.*, 2014, pp. 1–5.

[13] Intel. (2014, Feb.). *Intel Server Board S2600GZ and GL Technical Product Specification*, 2nd ed. [Online]. Available: http://download.intel.com/support/motherboards/server/sb/s2600gzgl_tps_r2_1.pdf

[14] Sun Microsystems. (2001, Dec.). *Sun Fire V120 and Netra 120 Server User's Guide*, A ed. [Online]. Available: https://docs.oracle.com/cd/E19088-01/v120.srvr/816-2090-10/816-2090-10.pdf

[15] P. Shenoy, I. Fedorov, T. Neyens, and P. Krein, "Power delivery for series connected voltage domains in digital circuits," in *Proc. Int. Conf. Energy Aware Comput.*, Nov. 2011, pp. 1–6.

[16] P. S. Shenoy and P. T. Krein, "Differential power processing for dc systems," *IEEE Trans. Power Electron.*, vol. 28, no. 4, pp. 1795–1806, Apr. 2013.

[17] P. Shenoy, K. Kim, B. Johnson, and P. Krein, "Differential power processing for increased energy production and reliability of photovoltaic systems," *IEEE Trans. Power Electron.*, vol. 28, no. 6, pp. 2968–2979, Jun. 2013.

[18] S. Qin and R. C. Pilawa-Podgurski, "Sub-module differential power processing for photovoltaic applications," in *Proc. IEEE 28th Annu. Appl. Power Electron. Conf. Expo.*, Mar. 2013, pp. 101–108.

[19] C. Olalla, D. Clement, M. Rodriguez, and D. Maksimovic, "Architectures and control of submodule integrated dc-dc converters for photovoltaic applications," *IEEE Trans. Power Electron.*, vol. 28, no. 6, pp. 2980–2997, Jun. 2013.

[20] Y. Levron, D. Clement, D. Maksimovic, and C. Olalla, "Nonlinear control design for the photovoltaic isolated-port architecture with submodule integrated converters," in *Proc. IEEE Energy Convers. Congr. Expo.*, Sep. 2013, pp. 2398–2405.

[21] Y. Yang and T. Lehmann, "Current recycling in linear regulators for biomedical implants," in *Proc. IEEE 53rd Int. Midwest Symp. Circuits Syst.*, Aug. 2010, pp. 545–548.

[22] Y. Yang, H. Chun, and T. Lehmann, "Dual-stacked current recycling linear regulators with 48% power saving for biomedical implants," *IEEE Trans. Circuits and Syst. I, Reg. Papers,*, vol. 60, no. 7, pp. 1946–1958, Jul. 2013.

[23] S. K. Lee, D. Brooks, and G.-Y. Wei, "Evaluation of voltage stacking for near-threshold multicore computing," in *Proc. ACM/IEEE Int. Symp. Low Power Electronics and Design*, 2012, pp. 373–378.

[24] S. Rajapandian, Z. Xu, and K. Shepard, "Energy-efficient low-voltage operation of digital CMOS circuits through charge-recycling," in *Proc. Symp. VLSI Circuits*, Jun. 2004, pp. 330–333.

[25] S. Rajapandian, K. L. Shepard, P. Hazucha, and T. Karnik, "High-voltage power delivery through charge recycling," *IEEE J. Solid-State Circuits.*, vol. 41, no. 6, pp. 1400–1410, Jun. 2006.
[26] Y. Byung-Do, "A high-efficiency on-chip dc-dc down-conversion using selectable supply-voltage charge-recycling," *IEICE Trans. Fundamentals Electron., Commun. Comput. Sci.*, vol. 94, no. 12, pp. 2676–2684, 2011.
[27] K. Kesarwani, C. Schaef, C. Sullivan, and J. Stauth, "A multi-level ladder converter supporting vertically-stacked digital voltage domains," in *Proc. IEEE 28th Annu. IEEE Appl. Power Electron. Conf. Expo.*, Mar. 2013, pp. 429–434.
[28] N. Kutkut and D. Divan, "Dynamic equalization techniques for series battery stacks," in *Proc. 18th IEEE Int. Telecommun. Energy Conf.*, Oct 1996, pp. 514–521.
[29] C. Pascual and P. Krein, "Switched capacitor system for automatic series battery equalization," in *Proc. IEEE 12th Annu. Appl. Power Electron. Conf. Expo.*, vol. 2, Feb 1997, pp. 848–854.
[30] C. Karnjanapiboon, K. Jirasereeamornkul, and V. Monyakul, "High efficiency battery management system for serially connected battery string," in *Proc. IEEE Int. Symp. Ind. Electron.*, Jul. 2009, pp. 1504–1509.
[31] S.-H. Park, K.-B. Park, H.-S. Kim, G.-W. Moon, and M.-J. Youn, "Single-magnetic cell-to-cell charge equalization converter with reduced number of transformer windings," *IEEE Trans. Power Electron.*, vol. 27, no. 6, pp. 2900–2911, Jun. 2012.
[32] R. Bell and R. Pilawa-Podgurski, "Asynchronous and distributed maximum power point tracking of series-connected photovoltaic sub-modules using differential power processing," in *Proc. IEEE 15th Workshop Control Model. Power Electron.*, Jun. 2014, pp. 1–8.
[33] J. McClurg, P. S. Shenoy, and R. C. N. Pilawa-Podgurski, "A series-stacked architecture for high-efficiency data center power delivery," in *Proc. IEEE Energy Conver. Congr. Expo.*, Sep. 2014, pp. 170–177.
[34] E. Candan, "A series-stacked power delivery architecture with isolated converters for energy efficient data centers," Master's thesis, University of Illinois at Urbana-Champaign, Champaign, IL, USA, 2014.
[35] E. Candan, P. S. Shenoy, and R. C. N. Pilawa-Podgurski, "A distributed bi-directional hysteresis control algorithm for server-to-virtual bus differential power processing," in *Proc. IEEE 16th Workshop Control Model. Power Electron.*, Jul. 2015, pp. 1–8.
[36] R. Pilawa-Podgurski, A. Sagneri, J. Rivas, D. Anderson, and D. Perreault, "Very high frequency resonant boost converters," in *Proc. IEEE Power Electron. Spec. Conf.*, Jun. 2007, pp. 2718–2724.
[37] Y. S. Lee and Y. C. Cheng, "A 580 kHz switching regulator using on-off control," *J. Inst. Electron. Radio Eng.*, vol. 57, no. 5, pp. 221–226, Sep. 1987.
[38] K. Ka-Sing Leung and H.-H. Chung, "Dynamic hysteresis band control of the buck converter with fast transient response," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 52, no. 7, pp. 398–402, Jul. 2005.
[39] M. Kheraluwala, R. Gascoigne, D. Divan, and E. Baumann, "Performance characterization of a high-power dual active bridge dc-to-dc converter," *IEEE Trans. Ind. Appl.*, vol. 28, no. 6, pp. 1294–1301, Nov. 1992.
[40] F. Krismer and J. Kolar, "Accurate power loss model derivation of a high-current dual active bridge converter for an automotive application," *IEEE Trans. Ind. Electron.*, vol. 57, no. 3, pp. 881–891, Mar. 2010.
[41] A. Waterland *Stress POSIX workload generator*. (2014). [Online]. Available: http://people.seas.harvard.edu/∼apw/stress
[42] Blue waters sustained petascale computing. (2014). [Online]. Available: https://bluewaters.ncsa.illinois.edu/
[43] E. Candan, D. Heeger, P. S. Shenoy, and R. C. N. Pilawa-Podgurski, "A series-stacked power delivery architecture with hot-swapping for high-efficiency data centers," presented at the IEEE Energy Conversion Congr. Expo., Montreal, Canada, Sep. 2015.

Enver Candan (S'12–M'14) received the dual B.S. degrees in control engineering and in electrical engineering from Istanbul Technical University, Istanbul, Turkey, both in 2012, and the M.S. degree in electrical engineering from the University of Illinois, Urbana-Champaign, Champaign, IL, USA, in 2014, where he is currently working toward the Ph.D. degree.

In the summer of 2015, he was a Hardware Engineering Intern at Google Inc. His research interests include data center power delivery systems and isolated dc–dc converters.

Mr. Candan received a Fulbright Scholarship in order to pursue his M.S. degree at the University of Illinois at Urbana Champaign, in 2012–2014.

Pradeep S. Shenoy (S'06–M'12) received the B.S. degree in electrical engineering from the Illinois Institute of Technology, Chicago, IL, USA, in 2007, and the M.S. and Ph.D. degrees in electrical engineering from the University of Illinois, Urbana-Champaign, Champaign, IL, in 2010 and 2012, respectively.

In 2008, he participated in the National Science Foundations East Asia and Pacific Summer Institutes program conducting research at Tsinghua University, Beijing, China. He interned with Caterpillars Electric Power Division in 2005 and with Texas Instruments Systems and Applications R&D Lab in 2011. He joined Kilby Labs, Texas Instruments, Dallas, TX, USA, in 2012.

Dr. Shenoy received the Camras scholarship in 2003–2007 and a Foreign Language and Area Studies fellowship in 2009–2010. He also received the Illinois International Graduate Achievement Award in 2010. He was a finalist for the Lemelson-MIT Illinois Student Prize for innovation in 2012 and the Jack Kilby Award for innovation in 2014. He was the Vice Chair of the IEEE Power Engineering Society and Power Electronics Society Joint Student Chapter at the University of Illinois in 2009–2010 and the Codirector of the 2011 IEEE Power and Energy Conference at Illinois. He currently serves in the IEEE Power Electronics Society as the Young Professionals Chair and as an Ad Com Member-at-Large.

Robert C. N. Pilawa-Podgurski (S'06–M'11) received the dual B.S. degrees in physics, electrical engineering and computer science in 2005, the M.Eng. degree in electrical engineering and computer science in 2007, and the Ph.D. degree in electrical engineering in 2012, all from the Massachusetts Institute of Technology, Cambridge, MA, USA.

He is currently an Assistant Professor in the Electrical and Computer Engineering Department, University of Illinois at Urbana-Champaign, Champaign, IL, USA and is affiliated with the Power and Energy Systems group. He performs research in the area of power electronics. His research interests include renewable energy applications, high-density and high-efficiency power electronics, CMOS power management, and advanced control of power converters.

Dr. Pilawa-Podgurski received the Chorafas Award for outstanding MIT EECS Master's thesis, the Google Faculty Research Award in 2013, and the 2014 Richard M. Bass Outstanding Young Power Electronics Engineer Award of the IEEE Power Electronics Society, for outstanding contributions to the field of power electronics before the age of 35. In 2015, he received the Air Force Office of Scientific Research Young Investigator Award. He is an Associate Editor for theIEEE TRANSACTIONS ON POWER ELECTRONICS, and for the IEEE JOURNAL OF EMERGING AND SELECTED TOPICS IN POWER ELECTRONICS. He is a coauthor of four IEEE prize papers.