

AC vs. Hybrid AC/DC Powered Data Centers: A Workload Based Perspective

Anuroop Desu, Udaya Puvvadi, Tyler Stachecki, Shane Case and Kanad Ghose
State University of New York, Binghamton, NY 13902
{adesu1, upuvvad1, tstachel, scase1, ghose}@binghamton.edu

Abstract— Proponents of AC-powered data centers have implicitly assumed that the electrical load presented to all three phases of an AC data center are balanced. To assure this, servers are connected to the AC power phases to present identical loads, assuming an uniform expected utilization level for each server. We present an experimental study that demonstrates that with the inevitable temporal changes in server workloads or with dynamic server capacity management based on known daily load patterns, balanced electrical loading across all power phases cannot be maintained. Such imbalances introduce a reactive power component that represents an effective power loss and brings down the overall energy efficiency of the data center, thereby resulting in a handicap against DC-powered data centers where such a loss is absent.

I. INTRODUCTION

US data centers consume about 71 billion kilowatt-hours of electricity annually [1]. The growing need for the diverse services offered by data centers has prompted a steady growth in the number of deployed data centers world-wide. The emergence of the Internet-of-things (IoT) will further fuel the demand for data center services worldwide. A large fraction of these data centers rely on coal-fired plants to provide their electric supply, affecting both the energy independence and the environment. Data centers deployed today primarily use Alternating Current (AC) utility power which is converted in multiple steps within the data center to different voltage levels. About 10% to 20% of the data center's total energy consumption is wasted in these conversions before delivering Direct Current (DC) power to the various components of the server and the cooling system [2, 3, 4]. It is thus imperative to investigate, develop and deploy techniques for reducing the energy demands of data centers by reducing losses in the wasteful power conversion and delivery/distribution steps.

The inherent inefficiencies in AC power distribution, conversions and use within data centers have not been widely questioned. Recently, segments of the data center industry have started taking a hard look at the advantages of distributing and using 380 volts DC within a data center. This approach requires a single AC-to-DC conversion from the utility line feed and retains only the DC-to-DC converters that have to be used. The technology challenges associated with first generation DC-powered data centers have been largely addressed on the electrical side. These challenges have been mitigated by the development of high-efficiency modular DC-to-DC power converters [2], new high-voltage switching devices [2, 5, 6], the development of arc-suppressing high voltage DC power connectors [3], efficient DC power distribution systems [3, 7, 8]. The foreseeable need to integrate micro-grid power from

renewable and/or locally generated sources to reduce the peak power draw of data centers makes another case for providing DC power directly to the servers.

One of the key sources of energy losses in a three phase AC power distribution system is the presence of reactive power. If the electrical loads on the three phases are not identical, a reactive power component is introduced, which represents an effective energy loss and a decrease in the overall energy efficiency of data centers [9].

DC-powered data centers offer inherently lower power conversion losses, higher reliability, smaller equipment footprints for power conversion and ease the integration of micro-grid power.

The pros and cons of using DC power instead of AC power in data centers have been widely discussed and arguments have been made to favor one over the other in [10, 11, 12]. Arguments made in favor of using AC power, which is typically distributed via three power phases, tacitly assume that servers present an equal electrical load to all three AC phases, assuming a specific utilization level of each server. Most of the existing literature (Section III) also make this assumption. Equal electrical loading across the AC power phases is facilitated with the use of three-phase Power Distribution Units (PDUs) to each rack. These PDUs bring in power from all three phases to a single rack to realize identical phase loads at the level of individual racks [13]. The power distributed by the PDUs is converted to regulated board level voltages by the power supply unit(s), PSU (s) inside individual servers.

Modern PSUs for AC-powered servers feature power factor correction to avoid presenting a reactive component to the power distribution system [14]. In reality, even with the use of these PSUs, server utilizations can vary temporally within a server and across servers, resulting in an imbalance across all three power phases. A similar power phase imbalance occurs when dynamic server capacity management techniques are in place. With dynamic server capacity management, servers are powered on or powered off dynamically to match the inevitable variation of the data center workload during the day in an attempt to improve the overall energy efficiency of the data center. The energy efficiency improvement comes by operating just the right number of servers to avoiding idling or low utilization servers. This paper experimentally demonstrates the impact of these phase imbalances to make another case for the use of DC power in data centers.

Power phase imbalance in a data center introduces two other problems. First, phase imbalances, resulting in reactive power components, introduce harmonics (higher frequency components) into the power line that can damage the PSUs when they cross the stipulated power quality standards for data




Fig. 1 (a). A typical AC-powered data center. AC power is delivered directly to the servers. The blue boxes represent bypass or distribution switches




Fig. 1 (b). A typical Hybrid AC-DC powered data center: AC power to the aisle, AC-DC with in the aisle followed by DC power distribution to the servers. Fewer conversion equipment reduces losses and footprint while enhancing reliability

centers (typically, a total harmonic distortion level of 5% [15]). Second, phase imbalances cause the currents in the three phase lines to differ from each other and may cause overload tripping when the current in one phase line exceeds the tripping threshold. All of the issues associated with power phase imbalance are obviously absent in DC-powered data centers.

This paper focuses on evaluating both AC and hybrid AC-DC power distribution systems using realistic workloads and by looking at inefficiencies that results in AC-powered data centers due to workload variations. Temporal workload variations during the day demand the use of dynamic server capacity provisioning. Such provisioning can result in imbalance across the AC power phases, representing a loss of energy efficiency. This also leads to the possibility of failure through the introduction of power line harmonics or interruptions caused by circuit tripping. In this paper, we thus make a hitherto unstated argument in favor of DC power delivery to servers in data centers.

II. BACKGROUND

Although grid power is delivered in AC form to the data centers, there are many compelling reasons to consider a change to DC power. In data centers, the power conversions from high voltage AC to low voltage DC are convoluted, expensive, occupy a large footprint and are highly inefficient. AC power also introduces barriers in the integration of locally distributed power generation and storage. DC power distribution and usage in data centers requiring a single AC-to-DC conversion step

from the utility feed, holds the promise of delivering higher levels of efficiency, the ability to incorporate locally-generated and/or renewable power sources easily, a smaller installation cost, a significantly smaller footprint and enhanced reliability and availability [2, 4, 16].

A common DC powered data center configuration is called a **hybrid AC-DC powered data centers**. In these data centers, AC power from the utility is converted to DC power (at voltages ranging from 12 Volts DC to 400 Volts DC to PSUs) within individual servers or at the rack level.

Fig. 1 (a) and 1 (b) present typical AC and hybrid AC/DC-powered data center configurations in a simplified manner. As seen in Fig. 1 (a), the AC power distribution system involves AC/DC and DC/AC conversions inside the UPS. When the AC power reaches the PSU of the server *after* going through few more AC to AC step-down conversions, the PSU converts the AC power at line voltage to the DC level voltages needed by the various server components. Regulators, local to the server components, are typically used on the motherboard to provide a stable voltage and/or to reduce noise.

In contrast, when the AC delivered by the utility grid reaches the DC power distribution system of a hybrid AC/DC data center (Fig. 1 (b)), only one AC/DC conversion takes place at the aisle level and high voltage DC (typically, 400 Volts) is delivered directly to the servers via aisle-local bus(es). The use of a 400 VDC feed to the server PSUs also eliminates resistive losses in the conductor compared to systems that use lower

voltage DC feed. The server PSUs then convert the 400 VDC input directly to other DC voltage levels that are needed by the various server components. In many configurations, independent rectifiers can be used for a few racks, for an entire row, for a single aisle or multiple adjacent aisles of servers.

III. RELATED WORK

Most of the existing work focus on the virtues of using DC power is in the context of fixed server loads. Examples of these are presented below.

A detailed case study about the improvements and opportunity for DC power in data centers is carried out by [17] and [18]. In [17], the authors conclusively demonstrated that DC delivery systems are practical and potentially cost less in long run.

In [19], authors have evaluated different configurations of a 400V DC power delivery system for data centers against an AC powered one. With a small-scale setup, authors have demonstrated that significant input power savings can be achieved between 480VAC and 400VDC with fixed loads.

A novel approach for DC powered data centers is proposed in [20], which involves a solar power generation system. The new integrated system demonstrated a 17% energy reduction and overall improvement in Power Usage Effectiveness (PUE) of the data center.

A thorough analysis on the reliability of the highly efficient DC distribution for data centers is carried out in [19]. With an analytical reliability model and quantitative evaluation, it was concluded that DC system has higher steady-state availability and survival probability. A detailed analysis on power quality and causes of its disturbances in DC data centers is studied in [22].

In [23], the authors concluded that the power losses due to phase imbalance is negligible, based on the use of modeling and experimental data. In fact, contrary to the observation made in [23], the energy losses in the conductor of the neutral line can be quite significant [9] when high-current lines are used - this being a growing trend as rack power levels continue to go up.

The work presented here contrasts with existing work in evaluating the advantages of hybrid AC/DC-powered data centers with realistic workload and natural temporal changes in the workload level. In AC powered data centers such workload variations introduce power-phase imbalance that reduce the overall data center efficiency.

In the rest of this paper, where used for the sake of brevity, “DC-powered” stands for the use of a hybrid AC-DC powered system for providing DC power to servers.

IV. CONSEQUENCES OF CAPACITY PROVISIONING

It is well-known that servers are most energy-efficient when they are operated at high utilization levels. This is a consequence of the fact that idling servers often dissipate a

significant amount of energy (in performing no useful work towards servicing requests). In typical servers, the idling power can often be about 15% to 35% of the power dissipated at full utilization [24]. This makes it imperative to reduce the online server capacity available to serve requests when the workload directed at the data center drops below the maximum expected workload level for prolonged periods. Many data center operators adjust the number of online servers devoted to serving online requests based on observed daily temporal load profiles – that is based on the observed workload level as a function of time during the day. This practice deploys just enough servers to serve the current requests, so as to operate these servers at high utilization levels instead of keeping all available servers on and running them at low utilization levels [25]. Some data centers also adjust the server capacity dynamically using preset rules [26, 27] or other sophisticated techniques [25, 28].

Adjusting the deployed online server capacity to match the demands of the instantaneous workload requires servers to be dynamically switched on or switched off, both of which introduce the possibility of power phase imbalance in an AC-powered data center. Such power phase imbalances result in an overall energy loss in the power distribution network. Of course, such losses will be absent in DC-powered data center when server capacities are adjusted dynamically to match workload demands as they vary temporally.

For workload patterns that are known apriori, such as predictable temporal variations of the workload during the day, the server capacity adjustments can be made to avoid idling and/or poor server utilization. To do this, one can use an “oracle table” that lists the number of online servers needed to handle a specific range of workload or equivalently overall utilization of the full capacity of the data center. To construct this table, we use the capacity provisioner (CP) introduced in [29]. In its normal operating mode, CP uses workload prediction to match the online server capacity to the workload expected in the immediate. We ran CP on the known daily workload pattern to construct the table that gives us the number of servers needed to serve the instantaneous requests, which is represented by the instantaneous overall data center utilization. CP ensures that the minimum number of servers are maintained to have negligible impact on request latency by following the technique described below.

For every utilization range for the data center, U_i , the number of servers needed to handle the corresponding workload, N_i , is obtained from a series of runs made using CP. The peak number of servers needed for a specific utilization level, obtained from a series runs using CP is stored in the oracle table. In effect, the runs made using CP serves as the oracle data that gives the number of online servers needed to serve the requests. Specifically, this table provides the value N_i , the number of servers needed to limit the average request service latencies to acceptable levels when the workload results in an overall data center utilization within the range $U_i \pm \Delta$ centered around U_i . The acceptable service latency is ensured to be close to what is realized when all servers are kept on. This is because the algorithm behind CP ensures that the latencies realized

using CP, approach latencies observed when the server capacity is fully provisioned.

To dynamically adjust the server capacity to match the instantaneous (and known) workload level, the data center's server utilizations are sampled at regular intervals and capacity adjustment decisions are taken at the end of such intervals. As the overall data center utilization changes from one utilization range U_i to another, say, U_j , CP either powers on an additional $(N_j - N_i)$ servers when $U_i + \Delta < U_j - \Delta$ or power off $(N_i - N_j)$ servers when $U_i - \Delta > U_j + \Delta$. To prevent frequent server turn-ons and turn-offs, the capacity adjustment decisions are only taken when the absolute differences between U_i and U_j persist for at least 6 second, thus effectively adding a hysteresis to the control system to prevent oscillations.

The relevant results of AC power phase imbalance on known workload patterns with oracle based server provisioning as described above appears in Section V.

V. EXPERIMENTAL ASSESSMENT

A. Experimental Setup

To evaluate the viability of DC power in data centers, several experiments are conducted between AC powered servers and DC powered servers. In our experimental setup, 112 HP ProLiant BL 460c Generation 7 (G7, dual Xeon-5675) Server blades are used for the evaluation. To provide power, cooling and I/O infrastructure to the blade servers, eight HP BladeSystem c7000 Enclosures are used. For the purpose of processing the user workload, all blade servers are running Linux kernel 4.4.0-131 and Ubuntu 16.04.2 LTS as the operating system. The networking in each enclosure is provided by HP Virtual Connect. Multiple Cisco Nexus 2248 switches are used for networking. To balance the workload across all servers equally, an F5 Network BIG-IP 4000s LTM programmable load-balancer is used whose list of target servers can be modified dynamically.

The three-phase power distribution system used in our studies has a wye configuration. Phase currents and neutral line currents are measured in real-time using Siemens power line instrumentation. The DC power distribution equipment includes an Emerson NetSure 9000 400V DC power system [30], a Starline Track Busway [31] and 4 ECHOLA Smart PDU (SPDU). The required data for the analysis is collected in real time from the DC (and AC) PDUs using Simple Network Management Protocol (SNMP).

To make a fair comparison between the use of AC and DC power in data centers, 8 HP blade enclosures are equally divided between two racks: one rack powered by AC and another rack powered by DC. For powering AC enclosures, 2 Server Tech 24V2C415A1 rack-mounted PDUs are used.

B. Workloads

Five unique workloads that are representative of online requests are used from the SPECjvm2008 benchmark suite [32] for the evaluation. The description of these five workloads is

given in Table I. Workloads are simulated using a custom harness tool by generating numerous unique HTTP GET requests. Configuration parameters for the tool can be modified to generate varying types of workload traces. All generated HTTP GET requests are recorded along with their time of generation in order to reproduce the same exact trace on another set of servers to make a fair comparison. The SPECjvm benchmarks are designed to have longer duration with an average service time of 16-21 seconds.

TABLE I. DESCRIPTION OF WORKLOADS USED

Name of the Benchmark	Description
SPECjvm Compress	Uses Lempel-Ziv for compression
SPECjvm Crypto.rsa	Does encryption and decryption of provided data
SPECjvm Scimark.Monte_carlo	Approximates the value of Pi using integral of a quarter circle.
SPECjvm Serial	Does producer-consumer work where communication is done using sockets.
SPECjvm Xml.Validation	Validates the XML files by comparing against .xsd files

We present the evaluation results in two sets. The first set of results demonstrates the power savings between a fully provisioned AC and a hybrid AC/DC power distribution system, without the use of CP. This gives us an opportunity to look at the inherent disadvantages and power losses in AC power distribution system and how it contributes to poor energy efficiency. The second set of results illustrates how phase-agnostic server capacity provisioning can lead to resource leakage even though overall IT equipment energy consumption is reduced. Table II describes different variants of CP used for the evaluation between AC-DC.

TABLE II. VARIANTS USED FOR EVALUATION

Name of the Variant	Description
Baseline-AC	Fully provisioned system using AC power for processing the workload.
Baseline-DC	Fully provisioned system using hybrid AC-DC power for processing the workload.
CP-AC	Capacity Provisioner for AC configuration.
CP-DC	Capacity Provisioner for hybrid AC-DC configuration.

C. AC vs. Hybrid AC-DC Powered Data Centers using Load Variations and without any Capacity Provisioning

Fig. 2 presents how the overall system utilization varies for a SPECjvm benchmark during the experiment. As shown, the average system utilization varies from 13% to 68%. Fig. 3 (a) shows the power consumption of the AC powered servers and the DC powered servers in Baseline configuration. As expected, the power consumed by DC powered servers is consistently less than the AC powered servers. The average power savings across all five benchmarks for DC powered servers is 5.40%. Compress has least power savings of 2.76% and Serial has highest power savings of 9.04%. The power




Fig. 2. Average System Utilization for a SPECjvm benchmark.

savings are dependent on type of workload and also on the workload trace. The mean request service latencies for both AC-powered, and the hybrid AC-DC powered systems are shown in Fig. 3 (b).

The average increase in the mean request service latency across all five benchmarks is 0.68% for requests processed by DC powered servers. As mentioned earlier, any increase or decrease in either the energy savings or the request latency is dependent on the nature of the workload. The increase in mean request latency is highest for Compress, measured at 4.8% and lowest for the Xml.Validation benchmark, a decrease of 3.2%, when compared to requests processed on AC powered systems.

D. Phase-Agnostic Server Capacity Provisioning

To evaluate how dynamic server capacity provisioning can lead to reduced energy efficiency in AC-powered data centers, we used a real-world workload trace, obtained from a data center operator (who prefers to remain anonymous). Fig. 4,




Fig. 3. AC and DC power consumption for different SPECjvm benchmarks.

compares different metrics for this trace (**Operator trace**) for the AC-powered system.

Fig. 4 (a) represents how the utilization changes over a period of time for Operator workload trace. As the system utilization changes from a high utilization phase to the low utilization phase, CP-based oracle provisions the required number of servers as described in Sec. IV and powers off excess number of servers.

In the 3-phase wye configuration used in our studies, with equal loading on the three phases, the neutral line current is (ideally) absent, that is, zero. A non-zero current in the neutral line implies a power phase imbalance and a reactive power component that represents an energy loss. Fig. 4 (b) shows how the neutral line current varies as the relative difference between the current flowing through each phase either increases or decreases. Even though power supply connections to each blade chassis are manually (that is, statically) configured to be phase balanced, workload variations result in different power draw from the three phases by each of the blade server and, in turn, by each chassis of blade servers. These variations




Fig. 4. Utilization and Power consumption and Neutral line Current in AC powered three phase power distribution system

introduce phase imbalances in a three phase AC power distribution system. By comparing Figures 4 (a) and 4(b), we can observe that, as the utilization drops to less than 25% and when power drawn decreases from each individual phase, the neutral line current does not drop commensurately as a consequence of the phase imbalance.

From Fig. 4 (b), the average current drawn across each phase during the high utilization period is 50.83 Amps. During this period, the average neutral line current is 19.57 Amps, which is 38.5% of the average current drawn from each phase. When the trace enters the low utilization period, the average current drawn from each phase reduces to 20.2 Amps but the average neutral current is 7.26 Amps, which is 35% of the average current draw from each phase. Furthermore, when workload trace transitions into the low utilization period, the average power drawn from each phase is reduced by 60% whereas the neutral line current is reduced by only 3.5%. This shows that the increase/decrease in neutral line current is not proportional to the current drawn from each phase, as expected. This also demonstrates that by using a CP oracle, which is phase oblivious, for dynamic server capacity provisioning during the low utilization period, will not reduce the neutral line current or the reactive power proportionately as the overall power consumption decreases.

The current wastage through neutral line in a three phase AC power distribution system is shown in Fig. 4 (c). As a consequence of powering on/off servers in the course of transitions between different utilization ranges, phase imbalance may occur. The increase in phase imbalances will cause energy wastage represented by the neutral line current as seen in Fig. 4 (c). The average increase in neutral line current between Baseline-AC and CP-AC is 9.23%.

Fig. 5 (a), presents the power consumption of Baseline and CP with respect to both AC and DC power distribution systems. From Fig. 5 (a), it is observed that because of the presence of reactive power and AC/DC conversion losses in the system, the power consumed by Baseline-AC is higher than Baseline-DC for the same IT workload. The average increase in power consumption between two Baseline variants is 5.47%, which revalidates the savings achieved in Fig. 3 (a) and reiterates that DC-powered data centers are much more energy efficient than AC-powered data centers. From Fig. 5 (a), the energy savings between Baseline and CP in either AC or DC powered system is significant. The average energy savings in AC configuration between Baseline-AC and CP-AC is 8.51%. For the corresponding DC configurations, the energy savings is 9.43%. Thus, the energy savings of CP itself in DC configuration are higher than its counterpart in AC configuration.

From Fig. 5 (a), the energy savings for CP-DC when compared against Baseline-AC is 12.06%, 17.34%, and 13.80% for Compress, Serial and Rsa benchmarks, respectively. Hence, by switching from AC power to DC power and with simultaneous use of a dynamic server capacity provisioner, an average energy savings of 14.40% can be realized, which is significant.





Fig. 5. Power Consumption and Mean request latency comparison between AC powered and hybrid AC-DC powered systems.

Fig. 5 (b), presents the mean request latency for all variants used in the second set of evaluations. The average increase in the mean request latency for CP against their respective Baseline in AC and DC powered systems is 0.60% and 0.57%, respectively. This demonstrates how effective, oracle based on CP is in limiting the request latency increase when compared to the Baseline system. This is achieved by provisioning enough online servers to guarantee SLA while achieving a significant saving in energy by powering down excess servers.


E. Phase Aware Server Capacity Provisioning

It is worth noting again that the amount of current flowing through the neutral line will stay high as long as the relative difference of the current flow among the three phases of an AC power distribution system stays high - even during the period of low utilization. Fig. 6 (a) shows how a power phase imbalance oblivious oracle based on CP is not sufficient in reducing the phase imbalances. Furthermore, the relative difference in currents among the three phases stays fairly constant even when current drawn from each phase decreases.

An obvious solution to reduce the neutral line current is to use a phase imbalance aware dynamic capacity provisioner. The oracle based on the CP discussed of Sec. IV was modified to do provisioning in a phase aware manner. Even though AC power is provided to two adjacent aisles in our research data center, the modified CP's oracle performs phase aware load-balancing by considering each aisle individually. This provides the modified CP's oracle with more fine grained control in provisioning servers to mitigate power phase imbalances. As seen from the Fig. 6 (b), the modified CP's oracle is able to reduce the relative difference among the three phases of an AC powered distribution system. During the period when modified



(a) Neutral Line current when phase oblivious CP is used.



(b) Neutral Line current when phase aware CP is used.

Fig. 6. Current Draw from each phase and Neutral line of a 3 phase AC power distribution system using different CPs.

CP is able to reduce the relative difference of the currents among the three phases, (i.e. when load is balanced across all three phases almost equally), the neutral line current momentarily reduces to almost zero. But, as the overall system utilization increases, the difference among the currents in the individual phases increases and, as a consequence, the overall neutral line current increases.

From Fig. 6 (b), it is evident that even with the use of phase aware provisioner, the power loss in the neutral line can never be eliminated completely and, at best, can only be reduced to some extent. Even though the results for Fig. 6 are derived from a different set of servers, *different from the ones in our studies*, a complete elimination of the neutral line current is impossible in practice.

In summary, the results establish that with the inevitable temporal variations in workload and/or with dynamic capacity provisioning, energy losses arising from a reactive power component are unavoidable in AC-powered data centers. DC-powered data centers have no such issues and permit the use of dynamic server capacity provisioning quite freely (without the constraints imposed by phase balancing considerations), to realize a higher energy savings.

VI. CONCLUSIONS

We presented an experimental study that demonstrates that electrical load imbalances are inevitable in data centers due to temporal workload variations or due to dynamic server capacity

management based on data center workload patterns that are known apriori. These imbalances are present even when three phase power is delivered to individual racks using 3-phase PDUs or when server loads are balanced statically across all three AC phases assuming a fixed utilization level for the servers.

We showed that, as workload level varies across servers on the three power phases, power phase imbalances are introduced. With phase-aware dynamic server capacity management, the degree of power phase imbalance can be reduced but cannot be completely eliminated. Even with the use of server power supply units with phase correction, reactive power and power phase imbalances need to be avoided. These imbalances represent a loss in its overall energy efficiency, introduce damaging harmonics and can cause a phase line circuit to trip because of overloading of a phase line with a high degree of imbalance. This study thus establishes another unquantified handicap that AC-powered data centers have against their DC-powered counterparts.

ACKNOWLEDGEMENT

This work is supported in part through NSF awards 1040666, 1134867, 1738793 and a NSF Grand Challenge award from the NSF IUCRC program.

REFERENCES

- [1] SHEHABI, A., S. SMITH, D. SARTOR et al, "United States Data Center Energy Usage Report", Lawrence Berkeley National Lab Report No. LBNL-1005775, June 2016.
- [2] D. Tuite, "400 V DC Power Distribution in the Data Center Gets Real", Electronics Design, March 2014.
- [3] H. A. Mantooth, M. D. Glover, and P. Shepherd, "Wide bandgap technologies and their implications on miniaturizing power electronic systems," in IEEE Journal on Emerging and Selected Topics in Power Electronics, Vol. 2, no. 3, pp. 374-385, Sept. 2014.
- [4] O. Saadeh, H. A. Mantooth, and J. C. Balda, "The modeling and characterization of silicon carbide gate turn off thyristors," in Proc. IEEE Energy Conversion Congress and Exposition (ECCE), 2012.
- [5] D. Geary, "380V DC Power for Data Centers" an Engineering Perspective", presentation at the National Energy Efficiency Technology Summit, Portland, OR, 2012.
- [6] K. Hirose, T. Tanaka, T. Babasaki, S. Person, O. Foucault, B. J. Sonnenberg, "Grounding Concept Considerations and Recommendations for 400 Vdc Distribution System", in Proc. IEEE/Intelec 2011.
- [7] W. Tschudi, M. Ton, B. Fortenberry, "DC Power for Improved Data Center Efficiency" Lawrence Berkeley National Lab Tech. Report, 2008.
- [8] Server Technology, Inc. Phase Balancing: The Last Few Inches Of a High-Efficiency Power System, White Paper STI-100-009, June 2010.
- [9] L. Boettger, "The Battle of the Currents, Take II – AC versus DC in the data center", ABB Conversations Blog, April 2013, available online at: <http://www.abb-conversations.com/2013/04/the-battle-of-the-currents-take-ii-ac-versus-dc-in-the-data-center>.
- [10] C. Garling, "AC/DC Battle Returns to Rock Data-Center World", Wired Magazine, Dec. 2, 2011. Available online at: <http://www.wired.com/2011/12/ac-dc-power-data-center>.
- [11] N. Rasmussen, and J. Spitaels, "A Quantitative Comparison of High Efficiency AC vs. DC Power Distribution for Data Centers", White Paper No. 127, Revision 3, Schneider Electric, 2014.
- [12] R. McFarlane, "Take data center efficiency to the max on three-phase power", 2015, available at: <https://searchdatacenter.techtarget.com/tip/Take-data-center-efficiency-to-the-max-on-three-phase-power>.

- [13] Data Center Power System Harmonics: An overview of effects on data center efficiency and reliability. In *The Green Grid*, August 2014.
- [14] Harmonic Currents in the Data Center: A Case Study. Available at: https://www.l-3.com/private/pacific_crest/articles/Harmonic%20Currents%20in%20the%20Data%20Center%20-%20A%20Case%20Study.pdf
- [15] S. Davis, "400 Vdc Power Distribution for Data Centers Emerges", *Power Electronics*, Oct. 2014, available online at: <http://powerelectronics.com/power-electronics-systems/400-vdc-power-distribution-data-centers-emerges>.
- [16] D. Geary, T. Martinson, "380v DC Power for Data Centers", in *Proc. National Low Voltage Direct Current Workshop*, Arlington, Virginia, Pennsylvania, April 8, 2011.
- [17] M. Ton, B. Fortenbery, and W. Tschudi, "DC power for improved data center efficiency," *Lawrence Berkeley National Lab*, report, Mar. 2008
- [18] M. Murrill and B. J. Sonnenberg, "Evaluating the opportunity for DC power in the data center" 2010. Available: <http://www.emersonnetworkpower.com/documentation/en-us/brands/liebert/documents/white%20papers/124w-dcdata-web.pdf>
- [19] A. Pratt, P. Kumar, and T. V Aldridge, "Evaluation of 400V DC distribution in telco and data centers to improve energy efficiency," in *Proc. Int. Telecommun. Energy Conf.*, 2007, pp. 32–39.
- [20] K. Hirose, "DC powered data center with 200 kW PV panels", in *Proc. Int. Power Electronic Conference*, 2018.
- [21] V. Sithimolada and P. W. Sauer, "Facility-level DC vs. typical ac distribution for data centers: A comparative reliability study," in *Proc. TENCON*, Nov. 2010, pp. 2102–2107.
- [22] S. Rajagopalan, B. Fortenbery and D. Symanski, "Power quality disturbances within DC data centers" 32nd International Telecommunications Energy Conference (INTELEC), pp. 1-7, 6-10 June 2010.
- [23] S. Gupta, A. Banerjee, Z. Abbasi, and S. K. S. Gupta. Effects of phase imbalance on data center energy management. In *Proceedings of 20th International Conference on High Performance Computing (HiPC)*, (2013), pp. 415-424.
- [24] FAN, X., WEBER, W.-D., AND BARROSO, L. A. "Power provisioning for a warehouse-sized computer". In *Proc. Intl. Sym. on Computer Architecture 2007*, ISCA, ACM, pp. 13–23.
- [25] WU, Q. "Making facebook's software infrastructure more energy efficient with autoscale", August 2014. [Online; posted 8-August- 2014].
- [26] IBM rules-based load balancing, available at: https://www.ibm.com/support/knowledgecenter/en/SSAW57_8.5.5/com.ibm.websphere.edge.doc/Ib/tprf_rules.html.
- [27] Solarwinds capacity planning, available at: <https://www.solarwinds.com/topics/server-capacity>.
- [28] D. LO, L. CHENG, R. GOVINDARAJU, L. A. BARROSO, AND C. KOZYRAKIS. Towards energy proportionality for largescale latency-critical workloads. *SIGARCH Comput. Archit. News* 42, 3 (June 2014), 301–312.
- [29] T. J. Stachecki and K. Ghose. "Short-term Load Prediction and Energy-Aware Load Balancing for Data Centers Serving Online Requests". In *proc. REC2: 1st Workshop on Resource-Efficient Cloud Computing*, held in conjunction with 42nd Intl. Sym. On Computer Architecture (ISCA), 2015.
- [30] Emerson Netsure 9000 Series, available at: <https://www.power-solutions.com/dc-power-systems/emerson/emerson-netsure-9000>.
- [31] StarLine Track Busway: <https://www.starlinepower.com/busway/>.
- [32] SPECjvm2008, <http://www.spec.org/jvm2008/>.