Effects of Phase Imbalance on Data Center Energy Management

Sushil Gupta, Ayan Banerjee, Zahra Abbasi, and Sandeep K.S. Gupta

***Abstract*—Phase imbalance has been considered as a source of inefficiency in the data center that causes energy loss due to line impedance and increases reactive power. Strategies assume high loss due to phase imbalance and propose sophisticated energy management algorithms including phase balance aware workload scheduling algorithm and dynamic power distribution unit assignment to servers. However, such attempts do not utilize an objective measure of the inefficiencies due to phase imbalance to evaluate the significance of their contributions. Excessive imbalance in a three phase load has various undesirable effects. This paper, first objectively characterizes the inefficiencies due to phase imbalance and then provides numerical measures of the losses in realistic data center deployments. Phase imbalanced load in a delta configuration results in reduced power factor, which is undesirable for several reasons. Also, an imbalanced load (both in delta or star configuration), results in higher line currents, leading to higher line loss. However, this increase in loss is a fraction of a percentage of the energy consumed. The paper also discusses effects of work load scheduling on phase imbalance, and how to minimize the same.**

Keywords*—*three phase imbalance, reactive power, data center

I. Introduction

Recently strategies to make data centers more power and energy efficient are highly governed by the electricity pricing models. It obviously makes more sense for energy management modules to control the particular parameter of the data center that is directly related to pricing. In most places around the world, the electricity price is proportional to the consumed power with extra penalties if the peak power exceeds a given value. Hence, a common strategy to save cost is to design data centers for low peak power. To this effect, the use of uninterrupted power supply (UPS) to store power during off-peak times for use during peak hours, or scheduling workload to avoid simultaneous peaks in different applications, or the use of renewable energy to reduce peak power draw from grid are considered effective. All these proposed solutions recognize a common evil - the inefficiency of the power distribution unit (PDU) due to imbalanced load. It is widely acknowledged that load imbalance in PDUs lead to several losses such as line loss, decrease in UPS efficiency and even workload allocation [1] and dynamic PDU allocation [2] algorithms are developed to avoid imbalance. However, there has been very limited analysis on the actual extent of these effects. This paper considers a theoretical analysis of load imbalance in three phase delta connected PDUs, and computes the line loss, effect on UPS capacity requirements and benefits of power imbalance aware workload scheduling in a realistic setting through simulation and experiments in an actual data center installation. Load imbalance also impacts the UPS efficiency, though a quantitative study for this is not included here.

In USA, PDUs are typically designed as three phase delta connected sources. In an ideal scenario all three phases of the PDU should be balanced, which means that the currents drawn from each line should be equal. If the load currents are not balanced, then the input power factor of the PDU deteriorates, implying the "apparent power" drawn from the UPS or the power grid is higher than the "true" or the useful power, also named as active power. "Power Factor" is the ratio of active power to the apparent power. In this case, input current drawn are higher than what it would be for the same load in balanced configuration, introducing a component called "Reactive Power", which does not deliver any useful power to the load, but contributed to higher currents, and higher I^2R losses. It would also require higher UPS capacity for the same true power of the load. Reactive power can also be induced in the circuit by the load if it has an inductive or capacitive component. However, power regulations in different countries such as the one in EU mandate power factor corrections in computing equipments. Hence, for load the power factor can be considered to be nearly unity. Thus, in this paper we consider that the power factor degradation is only possible through phase imbalance. A point to be noted is that regardless of the amount of active power or reactive power, the total power drawn from the power supply for both balanced and imbalanced configurations is same for the same load. This fact is discussed in detail in Section V.

One of the arguments against power imbalance in recent research is that it leads to increased power consumption. Researchers in [1] report a 14% energy saving by simply balancing load on the three phase of a PDU. In theory, the only power loss that can occur due to phase imbalance is due to heating effect of the supply wires. This is because although the average power consumptions for balanced or imbalanced phases under the same load are the same, an imbalanced phase leads to fluctuations in the power consumption. This leads to increased root mean square value of the current. Hence, there is an increase in line loss. However, the line resistance in contemporary PDU wires is very less of the order of a 10^{-3} ohms. Given such low resistance, it is expected that power loss due to heating will be insignificant. In this paper, we derive an objective measure of the line loss in Section VI and in Section IX we derive that the line loss is around 1% for a worst case phase imbalance.

Author Sushil Gupta is with HCL Infosystems Ltd, India. Ayan Banerjee, Sandeep K.S. Gupta, and Zahra Abbasi are with IMPACT Lab (http://impact.asu.edu) at Arizona State University. Their work is funded in part by the NSF grants CNS #1218505 and #0855277. The authors also thank Georgios Varsamopoulos and Joshua Ferguson for help with performing the experiments in Bluetool. This work was done when the first author was visiting IMPACT Lab. Email: skg001@gmail.com,{abanerj3,zabbasi1,sandeep.gupta}@asu.edu.

As discussed earlier, phase imbalance causes reduction in power factor of the PDU. This has serious consequences in the efficiency of power delivery by the UPS to a particular load. Typically, UPS capacity for a particular PDU or an entire data center is decided based on the peak power requirements. The UPS has its internal inefficiencies due to waveform distortion caused by conversion from AC to DC and then back to AC. The UPS provisioning for a data center also takes into account the load power factor however, it does not account for phase imbalance in the PDU. PDU phase imbalance may cause lesser useful power to flow through the load for the same apparent power drawn from the UPS. Thus, to support the given load at a poorer PDU power factor a higher capacity of the UPS is necessary, increasing the facility cost. In this paper, we quantify the power factor of PDU for a given phase imbalance and obtain numbers for realistic data center deployments.

Phase imbalance can be induced by oblivious workload allocation algorithm in the data center. Recent research [1] have focused on avoidance of phase imbalance while scheduling workload on servers. A lookup table based approach is proposed where a large number of experiments are carried out previously to determine the losses for different load configurations to the three phases. The lookup table is then used to search for an optimal configuration for a given load. Another approach is to use imbalance as a hard constraint in an optimization framework [2]. However, these approaches do not consider a quantitative measure of the inefficiencies due to phase imbalance and hence are inaccurate in estimating the benefits of performing phase aware workload allocation. In this paper, we consider an objective measure of the inefficiencies and report numbers on the energy benefits and power factor improvements of the PDU if workload is allocated such that there is always a nearly balanced PDU.

The paper makes the following contributions:

- characterizing the effects of phase imbalance on data center power consumption;
- a thorough analysis of the impact of line loss and power factor reduction on UPS capacity estimation; and
- evaluating the efficacy of phase aware workload scheduling through simulation and experimentation on a real data center deployment.

We use BlueCenter [15], a data center testbed deployed by us at Arizona State University, to perform the experiments and the BlueSim tool developed by us to perform simulations on the model of BlueCenter. From our findings we have the following conclusions:

- There is no significant energy savings in phase aware workload allocation since line loss due to phase imbalance is less than 1%.
- Power factor of the PDU degrades if there is phase imbalance and can result in very poor UPS efficiency. Thus, to support the same load the capacity of the UPS has to be increased.
- Incorporation of maximum phase imbalance as a hard constraint in an optimization problem for phase aware workload allocation leads to non-linear convex constraints, which are computationally infeasible to solve optimally.

The rest of the paper is organized as follows: Section II discusses the related work on data center energy efficiency and the treatment of phase imbalance by researchers, Section IV discusses the system model of BlueCenter, Section V discusses the theory of phase imbalance in detail, Section VI theoretically evaluates line loss due to phase imbalance, Section VII considers UPS power factor degradation, Section VIII discusses the benefits of phase imbalance aware workload scheduling, Section IX verifies the theoretical conclusions through simulations and experiments, Section X discusses some long term research implications of phase imbalance from our findings, and Section XI concludes the paper.

II. Related Work

Data center power management is increasingly receiving attention in academia and industry. Significant power management solutions in data centers rely on servers' power management knobs, those being DVFS, and servers' power state transitioning (e.g., c-states). These solutions affect the power consumption of the servers, the data center cabinets and the power phases i.e., can cause phase imbalance. However, little research work has been published about issues regarding phase imbalance in data centers. In the following, first we review the related work which motivate the necessity of research on the effect of phase imbalance on the data center power management, and then we give an overview on the existing works which study the effect of phase imbalance on the data center power management.

A. Power infrastructure

Power consumption and power infrastructure significantly contribute in the operational and capital expenditure of the data centers. Therefore, there are considerable amount of research which propose to increase the power efficiency and cost efficiency of the data centers' power infrastructure through workload scheduling [3]–[6], power management schemes [3], [4], [7], [8] and power efficient infrastructure designs [2], [9]. Particularly, the power infrastructure cost depends on the peak power consumption of the data centers. Therefore, some recent work propose schemes to under provision the power infrastructure without significantly affecting the performance using the existing energy storage devices (i.e., UPS) [9]–[12]. The basic idea comes from the assumption that simultaneous peak draw across all equipment will happen rarely, therefore it is more attractive for data center providers to overbook the power infrastructure for a high percentile of the needs rather than the occasional peaks. However, such an under-provisioning, results in non-zero probability of events where power needs exceed provisioned capacity. Those rare event can be addressed using combination of stored energy in UPSes and several power management knobs of servers (i.e., DVFS) which comes at the cost of negligible performance degradation of the applications (i.e., increasing response time) [9], [11]. Such schemes are shown to be effective, particularly for modern data centers which utilize distributed UPSes (deploying UPS per rack or per server rather than centralized UPS to remove double AC conversion) [11]. In other words, with increasing the availability of UPSes in data centers it is cost efficient to utilize them not only during utility power outages but also to shave data centers' peak power draw. However, the reported results are based on the assumption that the workload and power management do not cause phase imbalance and consequently do not deteriorate the power factor, which is not the case in practice. While the aforementioned related work motivate the importance of UPSes in energy and power management of the data centers, the effect of phase imbalance on the UPSes needs to be carefully studied. The reason is that the phase imbalance increases the current demand from energy storage, which may not be sufficiently supplied in the case of power factor deterioration.

III. Effect of Phase Imbalance

There are a small group of works in which the phase imbalance in data centers has drawn attention [1], [2]. However, these research focus either on the phase aware power allocation in data centers to avoid phase imbalance [1], [2], or on characterizing the power inefficacy resulting of phase imbalance [1]. Particularly, the existing work lack of theoretical, simulation or experimental study on the effect of phase imbalance on the reactive power and its potential impact on the data center power infrastructure.

Lama et al., propose pVOCL, a power and phase aware scheme to dynamically manage the placement and the migration of the virtulized GPUs across server clusters [1]. The proposed scheme dynamically consolidate GPUs across server clusters not only to satisfy the power budget and to remove unnecessary idle power, but also to increases the power efficiency of the cluster by avoiding phase imbalance. The authors frame their management scheme as an optimization problem and perform several experiments to show the power efficiency of their scheme. The experiments are performed in a Cabinet Power Distribution unit (CDU) with three computing nodes in each power-phase. There are 18 GPUs in total (2 GPU per node). The experiment is performed for a duration of 10 minutes and a time-varying workload where the number of required GPUs vary over time. The authors report an improvement of 14% in energy efficiency due to phase balance in this experiment (i.e., GPU placement in such a way that the phase balancing is preserved as much as possible). On the contrary, our experiments in BlueCenter shows a much lower improvement in energy efficiency (less than 1%). Further our simulation results for a range of workload each running for 24 hours also shows a maximum energy efficiency improvement of 1%. Furthermore, we find that workload scheduling, to reduce the power loss due to phase imbalance, in multiple PDUs is significantly different form scheduling in a single PDU. Despite the differences, we believe that phase balance is still important for data centers not only to improve the data center overall energy efficiency, but also to efficiently manage power draw from UPSes. We support our arguments through theoretical and simulation results.

Pelley et al., propose Power Routing which schedules IT load dynamically across redundant PDUs in order to reduce data center power facility size at peak demand and to balance power draw across AC phases [2]. The proposed algorithm is based on a constraint programming where the algorithm only allow the load on the three phases of any PDU to differ by no more than 20% of the per-phase capacity. The phase balance is performed in order to reduce heating and improve electrical stability. However, our study shows that such a heuristic scheme to create phase balance does not guarantee minimizing total power, power loss and consequently heating since (i) the percentage of power loss, out of total power consumption, due to phase imbalance is negligible (around 1%), and (ii) minimizing power loss in the power allocation of multiple PDU needs a complex nonlinear constraint. Consistent with the argument of the paper [2], we conclude that heuristic power allocation to induce phase balance improves power factor.

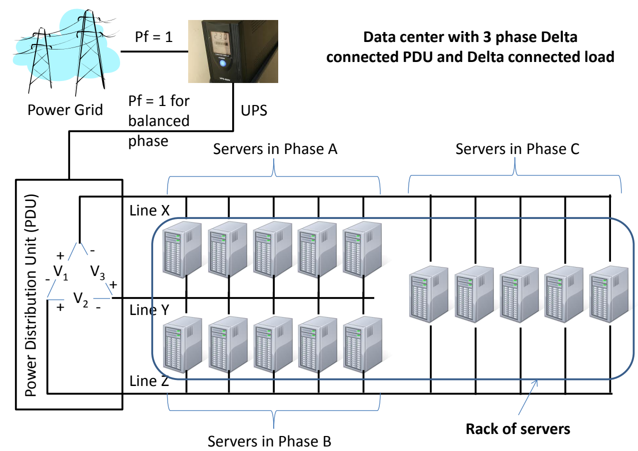

Fig. 1. **System model of data center power supply infrastructure.**

IV. Data Center System Model

Figure 1 shows the assumed system model of the power supply infrastructure of a data center. We assume that servers in a data center are arranged in chassis which in turn are stacked in racks. Each rack has a power distribution unit (PDU) that gets electricity from an uninterrupted power supply (UPS) unit. The UPS unit is connected directly to the power grid. The PDU can be considered as a balanced three phase delta connected source. Servers in the chassis are connected to different phases of the PDU such that the load is distributed as evenly as possible to the three phases. Each phase is a connection between two lines of the PDU e.g., phase A is a connection between lines X and Y as shown in Figure 1.

The even distribution of the load into the three phases of the PDU is necessary to reduce reactive power in the PDU and to increase its efficiency (discussed further in Section V). If the load in the three phases of the PDU are different then several inefficiencies are observed such as increase in the capacity of UPS to supply the same amount of power to the rack and increase in line loss. Further, server consolidation can lead to selection of active servers such that there is imbalance in the phase. This may cause power loss that negates the benefits of the consolidation algorithm.

V. Preliminaries

The PDU in a data center is a balanced three phase source as shown in Figure 2. The delta connected PDU is fed from

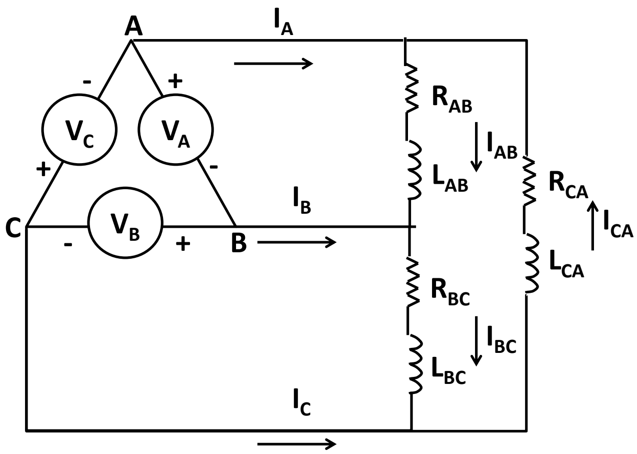

Fig. 2. **Circuit diagram for three phase source from the PDU and three phase load.**

the power supply grid through an UPS. The load in the data center i.e. the servers are connected to three lines of the delta connected source. In this configuration even if each of the loads is a unity power factor, the power factor of supply phases becomes below unity, worst case value being 0.87 ($\frac{\sqrt{3}}{2}$).

In a three phase source the voltage at the three lines V_A, V_B, and V_C have the same magnitude but are shifted in phase by 120°. If these are represented by $V_A = V_p sin(\omega t)$, $V_B = V_p sin(\omega t - 120°)$, and $V_C = V_p sin(\omega t - 240°)$, where V_p is peak voltage (line-Neutral) of each phase, ω is the angular velocity[1], t is time, the line to line voltage are given by Equation 1.

$$\begin{aligned} V_{AB} &= \sqrt{3}V_p \sin(\omega t + 30°), \\ V_{BC} &= \sqrt{3}V_p \sin(\omega t - 90°), \\ V_{CA} &= \sqrt{3}V_p \sin(\omega t - 210°). \end{aligned} \quad (1)$$

The currents in each line of the 3 phase delta connection are represented as I_A, I_B, and I_C. These line currents result in load currents I_{AB}, I_{BC}, and I_{CA} in directions as shown in Figure 2. The load in each line can have complex impedances Z_{AB}, Z_{BC}, and Z_{CA}. This causes the load currents I_{AB}, I_{BC}, and I_{CA} to have leading or lagging angles as shown in Equation 2. If the load is a capacitive load then the current is leading (ϕ is negative) and if it is an inductive load the current is lagging (ϕ is positive).

$$\begin{aligned} I_{AB} &= \sqrt{2}I_{abrms}(\sin(\omega t - \phi_A + 30°), \\ &= \sqrt{2}I_{abrms}(\sin(\omega t + 30°)\cos(\phi_A) - \cos(\omega t + 30°)\sin(\phi_A)), \\ I_{BC} &= \sqrt{2}I_{bcrms}(\sin(\omega t - \phi_B - 90,°\ ) \\ &= \sqrt{2}I_{bcrms}(\sin(\omega t - 90°)\cos(\phi_B) - \cos(\omega t - 90°)\sin(\phi_B)), \\ I_{CA} &= \sqrt{2}I_{carms}(\sin(\omega t - \phi_C - 210°), \\ &= \sqrt{2}I_{carms}(\sin(\omega t - 210°)\cos(\phi_C) - \cos(\omega t - 210°)\sin(\phi_C)), \end{aligned} \quad (2)$$

where, I_{abrms}, I_{bcrms}, and I_{carms} are the root mean square (rms) currents through each load. In phasor notation we have: $I_{AB} = I_{abrms}[\cos(\phi_A) - j\sin(\phi_A)]$. The phase angles ϕ_A, ϕ_B, and ϕ_C, depend on the characteristics of the load in each line. For a purely resistive load the angles are zero. If the load has an inductive component then the angles are positive and if the load has a capacitive component then the angles are negative. The real component of the current delivers useful power and is called *active power*. The imaginary component of the current is not useful and only results in loss of power due to the line impedance. The power through the load due to imaginary component of the current is called *reactive power*. The *apparent power* delivered to the load is the vector sum of the active power and the reactive power. The ratio of the active power to the apparent power is called *power factor*. The power factor can be quantified as the cosine of the angle ϕ_A, or ϕ_B, or ϕ_C. We assume each line has internal impedance equal to $Z_{line} = R_{line} + jwL_{line}$ or $R_{line} - j/(wC_{line})$ that may be complex.

The line currents can be expressed in terms of the load currents using Kirchoff's laws, as shown in Equation 3.

$$\begin{aligned} I_A &= I_{AB} - I_{CA}, \\ I_B &= I_{BC} - I_{AB}, \\ I_C &= I_{CA} - I_{BC}. \end{aligned} \quad (3)$$

Using Equation 2 to compute the line currents from Equation 3 we obtain the real and complex components of the line current as shown in Equation 4. Here for simplicity of representation we assume that the loads have equal power factors on each line. Even if the currents in the phases are unequal we assume that the power factor on each line remains the same, i.e., $\phi_A = \phi_B = \phi_C = \phi$.

$$\begin{aligned} I_A &= I^A_{line}\sin(wt - \phi - \theta_1), \\ I_B &= I^B_{line}\sin(wt - \phi - \theta_2 - 120°), \\ I_C &= I^C_{line}\sin(wt - \phi - \theta_3 - 240°), \end{aligned} \quad (4)$$

where,

$$\begin{aligned} I^A_{line} &= \sqrt{2}\sqrt{I^2_{abrms} + I^2_{carms} + I_{abrms}I_{carms}}, \\ I^B_{line} &= \sqrt{2}\sqrt{I^2_{abrms} + I^2_{bcrms} + I_{abrms}I_{bcrms}}, \\ I^C_{line} &= \sqrt{2}\sqrt{I^2_{bcrms} + I^2_{carms} + I_{bcrms}I_{carms}}, \\ \tan(\theta_1) &= \frac{1}{\sqrt{3}}\frac{I_{carms} - I_{abrms}}{I_{abrms} + I_{carms}} \\ \tan(\theta_2) &= \frac{1}{\sqrt{3}}\frac{I_{abrms} - I_{bcrms}}{I_{abrms} + I_{bcrms}} \\ \tan(\theta_3) &= -\frac{1}{\sqrt{3}}\frac{I_{bcrms} - I_{carms}}{I_{bcrms} + I_{carms}}. \end{aligned} \quad (5)$$

The total power loss due to line impedance is given by Equation 6.

$$P_{loss} = ((I^A_{line})^2 + (I^B_{line})^2 + (I^C_{line})^2)R_{line}. \quad (6)$$

Balancing a three phase delta connection means that the load in each phase are equal i.e., $I_{abrms} = I_{bcrms} = I_{carms}$. An imbalanced three phase delta connection can lead to several implications some of which are listed below:

1. Line loss due to line impedance: In a balanced three phase delta connection, the angle $\theta_1 = 0°$, $\theta_2 = 0°$ and $\theta_3 = 0°$ and the total imaginary or reactive current delivered to the load is zero if the power factor of the load is unity. However, if there is a load imbalance then the θ values deviate from the previous values and the total imaginary or reactive current drawn from UPS will be non-zero. This reactive current will induce more line loss for an imbalanced load (see Section VI).

2. Increase in required UPS capacity due to line imbalance: Phase imbalance leads to reactive power hence, to provide the

[1] ω is 2π frequency of the supply (2π60Hz).

same amount of power to the load the UPS has to supply more current. With a poorer imbalance the power factor of the PDU will get worse and hence the active power will reduce for the same amount of apparent power. To maintain the required amount of active power the UPS has to supply more apparent power. As a result, UPS with a larger capacity is needed to provide the required amount of active power (see Section VII).

3. Effect of workload scheduling on phase imbalance: Workload scheduling to servers may induce phase imbalance since most of the recent workload scheduling algorithms do not consider balancing the load on different phases. A workload scheduling that causes a phase imbalance can induce more line loss and can also cause increase in reactive power. More line loss can negate the energy savings from the workload scheduling algorithm. Also, increased reactive power means VA rating of UPS needs to be increased (see Section VIII).

In this paper, we will evaluate the extent of the three above-mentioned effects through simulations on a model of real data center with actual workload and also through experimentation in a real data center installation.

VI. Effects of Phase Imbalance on Total Power Draw

The average active power delivered to the servers through the three phase delta connection does not change with phase imbalance. However, due to phase imbalance the power factor deteriorates and hence there is reactive power. Thus, the apparent power from the source increases which results in increased heat losses through line impedance. The average active power delivered by the source to the load can be computed using Equation 7.

$$\begin{aligned} P_{active}^{avg} &= V_{prms} I_{line}^{A} \cos(\theta_1 + \phi) + V_{prms} I_{line}^{B} \cos(\theta_2 + \phi) \\ &+ V_{prms} I_{line}^{C} \cos(\theta_3 + \phi). \end{aligned} \quad (7)$$

where θ and the current values are obtained from Equation 5. When there is no phase imbalance the line currents are equal and the angles are exactly 120 °apart which results in a total active power of $3V_{prms}I_{linerms}$. However, with phase imbalance the power factor deviates from 1 in each line and the apparent power increases. The power factor due to phase imbalance can be expressed using Equation 8.

$$Pf_{imb} = \frac{P_{active}^{avg}}{P_{app}^{avg}}, \quad (8)$$

where P_{app}^{avg} denotes the average apparent power. Due to poorer power factor caused by phase imbalance the apparent power increases, which causes increase in the line loss P_{loss} as derived in Section V. However, given that the line impedance is minimal, the percentage of line loss due to phase imbalance is very small as observed in our experiments and simulations discussed more in Section IX.

In order to get a numerical estimate of the total power loss due to phase imbalance, let us consider Example 1.

Example 1: Estimate of line loss due to phase imbalance for delta connected PDU:. Let's say the line resistance is 0.04 ohms for each phase. The 3-phase PDU has a total load capacity of 6 kVA at 208 Volts, ie, 2 KVA per phase (delta connected, line to line) Active load on the PDU is only 2 KVA / 2 KW, consisting of three servers, unity power factor, each consuming 666.6 W.

Case - I If the three servers are distributed to three different branches (one on AB, one on BC, and one on CA). Each server draws 666.6/208 = 3.2048 A (say x). Each phase has a current $(\sqrt{3}) \times 3.2048 = 5.5509A$ ($(\sqrt{3})x$). The I^2R loss (on each phase) is $5.5509 \times 5.5509 \times 0.04 = 1.2325W$ ($3x^2R$). Thus the overall power loss for all three phases is 3.69749 W (0.18% of load) ($9x^2R$)

Case-II All three servers are connected on the same branch (say branch AB). The current on each of phase A and Phase B = 2000/208 = 9.6154 ($3x$). The current on phase C is zero. The total I^2R loss on each of phase A and B = $9.6154 \times 9.6154 \times 0.04$ = 3.6982 W ($9x^2R$). Loss on all (two) phases = 7.3964 W (0.37% of the load) ($18x^2R$)

As the above examples show, loss increases for unbalanced load, from 0.18% to 0.37%. Theoretically, the loss will be exactly double in case-II.

Increase in loss is because of (i) poor power factor in case-2, resulting in larger current (sum of currents in all three phases is 19.2308 A as opposed to 16.6527 A), and (ii) instead of three wires, the current is spread over two wires, thus increasing the heat loss.

Similar computation can be made for a star connected PDU as well. Unlike the delta connection, for an star connected set of loads, the line current will always equal the load current connected to that phase, and power factor of phase current will remain same as of the load. However, neutral current is a sum of all the phase currents. We show a numerical example (Example 2) for star connected PDU in this paper but we wont discuss the theory in detail.

Example 2: Line loss for imbalance in star connected PDU[2] - Let us take the similar case of a PDU, star connected, connected to a 230 V L-N supply, with total power capacity of 6 KVA, ie, 2 KVA per phase. Each of phases and neutral has a line resistance of 0.04 ohms. Each phase is connected to three servers, each server having a load of 666.6 W at unity power factor. We compare two cases, first with only three servers on one phase is on, and the rest are off. Second case with one server on each phase is on.

Case-I One server each on phases A, B and C. Current in each phase I_A = I_B = I_C = 666.6/230 = 2.8983 A (x) $I_{neutral}$ = zero, I^2R loss in each phase = 2.8983 × 2.8983 × 0.04 = 0.3360 (x^2R) Total in all three phases = 1.0080 W = 0.05 % ($3x^2R$)

Case - II All three servers on phase-A and rest of servers are off. I_A = 2000 / 230 = 8.6957 A (3x), $I_{neutral}$ = 8.6957 (3x).

I^2R loss in phase A = 8.6957 × 8.6957 × 0.04 = 3.0246 W ($9x^2R$) I^2R loss in Neutral = 3.0246 W ($9x^2R$) Total loss = 6.0491 W = 0.302 % ($18x^2R$)

So, a phase imbalance results in six times increase in line loss.

VII. Effects of Phase Imbalance on UPS

There are three common types of UPSes including the stand by UPS, line interactive UPS and double conversion UPS, each

[2]As opposed to delta, star configured UPSs are used in India and elsewhere where supply voltage is 230 V. At this point authors are looking into better understanding the pros and cons of using one over the other.

having their own inefficiencies [13]. The stand by UPS either solely uses the power grid or the battery to supply the load. The line interactive UPS uses the power grid to supply the load but simultaneously uses the battery to compensate for variations in the power supply. A double conversion UPS which is the most commonly used completely isolates the power grid and the load. The power from the grid charges a battery after passing through a AC-DC converter. The power from the battery is fed to an inverted which converts the DC power to AC power and serves the load. In this paper, we only consider double conversion UPSes and their inefficiencies. One of the main source of inefficiency in a double conversion UPS are the AC-DC converter and DC-AC inverter. These two conversions distort the waveshape of the delivered power and hence results in harmonics in the signal which result in power loss. The power loss due to each harmonic h_i is characterized by a parameter called total harmonic distortion (THD_i). THD is the ratio of the rms current of the component at the harmonic frequency to the rms value of the overall current. This power loss can be expressed using a distortion power factor (DPf), where $DPf_i = 1/(\sqrt{1 + THD_i^2})$. To avoid this distortion power loss delta conversion UPSes are proposed which can dynamically regulate the output power factor. However, most of the contemporary research on data center power consumption efficiency do not assume the existence of such cutting edge technology. We thus only consider a double conversion UPS in our research and quantify the effects of phase imbalance.

Commercially available UPSes have two ratings: a) the VA rating, VA_{ups}, which signifies how much apparent power can they deliver to the load, and b) the wattage rating, W_{ups}, which signifies the amount of active power that the UPS is capable of delivering. The UPS itself has a rated power factor, Pf_{ups} for delivering active power to the load, which takes into account the distortion caused by all harmonics. The load to the UPS also has a VA rating, VA_{load} and a power factor Pf_{load}. Several power supply regulations such as EN61000-3-2, require the load to provide for power factor correction such that $Pf_{load} \approx 1$, and $VA_{load} \approx W_{load}$. The UPS can supply power to a load if: a) the VA rating of the load is less than the UPS, i.e., $VA_{ups} > VA_{load}$ and b) the wattage rating of UPS is more than the load, i.e., $W_{ups} > W_{load}$.

UPSes are connected in several configurations in a data center as shown in Figure 3. With a higher infrastructure cost a data center operator can install a delta conversion UPS of very high capacity to serve the whole data center, or there can be several cheaper double conversion UPSes for groups of chassis or UPSes can be used for each server. In the first two cases the UPS will deliver power to the PDU which is a three phase delta connected source. In the third case, the UPS directly supplies current to the server [10], however, this may require special type of servers which support direct power draw from the UPS. The UPS in this case does not need the inverter and thus have lower power loss due to distortion. However, this calls for greater expenses in infrastructure. So for the purpose of this paper we consider the second configuration where there is an UPS for each PDU and the PDU servers a pre-decided number of chassis.

In the assumed configuration, the UPS supplies power to a delta connected three phase source. The UPS capacity is designed such that even with its distortion power factor it is

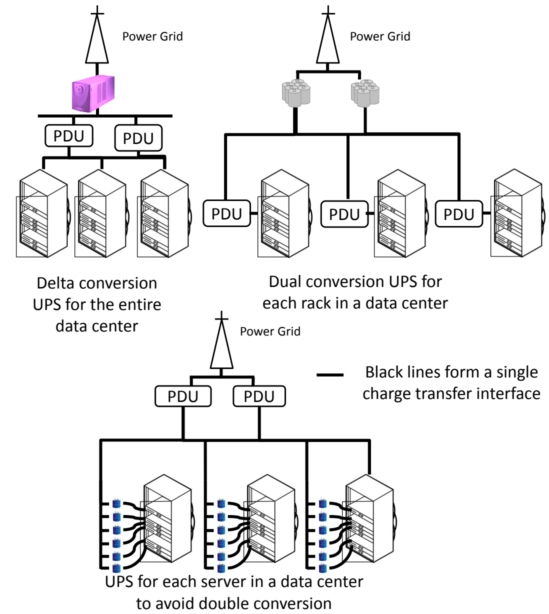

Fig. 3. **Different power supply architectures used in data centers.**

capable of supporting the maximum load that the PDU can handle. Thus, if the PDU is rated at VA_{pdu} and the load power factor is nearly 1 as mandated by regulations then the PDU is designed such that $VA_{ups} \times Pf_{ups} > VA_{pdu}$. Let us consider that the load has a maximum wattage requirement of $W_{load} \approx VA_{load} < VA_{pdu}/3$. The total power required by the load does not change even if there is imbalance in the 3 phases. Thus, the total power required by the load with the three phases combined is $3W_{load} < VA_{pdu}$.

VIII. Effects of Workload Scheduling on Phase Imbalance

This section studies how workload scheduling affects the phase imbalance. Workload scheduling and distribution involves splitting incoming workload among the allocated servers, always with respect to quality of service requirements (i.e., delay requirements).

Specifically, we study the phase imbalance due to power management by Thermal Aware Server Provisioning (TASP) [3], [4]. TASP resizes the active server set to the incoming workload, chooses the active server set such that the cooling power is minimized, and distributes the workload equally among active servers. The rest of the servers are assumed to be turned off which helps to remove unnecessary idle power and make a power proportional data center. In other words, TASP minimizes the sum of cooling power and computing power by dynamically resizing the active server set. However, server provisioning consolidates the workload to one "side" of the data center, where the servers are utilized at or near their peak level; potentially resulting in phase load imbalance if the active server selection is oblivious to the load balancing across phases. We compare TASP over no server provisioning where the workload is equally distributed among all servers, consequently resulting in load balanceing across phases. To study the effect of TASP on the phase load imbalnces in a data center we design several phase allocation algorithms to illustrate the highest and the lowest phase imbalance due to TASP:

- **No Server Provisioning (NoSP)**: All servers are assumed to be active and the incoming workload are equally distributed across all servers. This scheme causes a balanced load across different phases.
- **TASP-Worst:** The phase allocation is such that the active servers cause an imbalanced load across phases as much as possible. This is a reference algorithm and provides the upper bound on the maximum load imbalance that a server provisioning algorithm such as TASP can induce. The algorithm can not be implemented in practice as the phase allocation is performed according to the servers' incoming workload. The algorithm works as follows: (i) start allocating active servers to the phase one as much as possible, (ii) if there is no capacity left to load the phase one and there is unallocated active servers, start allocating the rest of the active servers to the phase two as much as possible, otherwise allocate enough inactive servers to the phase two, and (iii) allocate the rest of unallocated active server (if any) and unallocated inactive servers to the phase three.
- **TASP-Best:** The phase allocation to the active servers are such that the active servers cause a balanced load across phases as much as possible. Similar to the TASP-Worst, this algorithm can not be implemented in practice since the phase allocation is optimally performed based on the servers' incoming workload. The algorithm allocates the active servers to the different phases using "Round-Robin" scheme to provide the maximum possible phase load balancing.
- **TASP-Random:** Servers are randomly allocated to the phases, i.e., the phase allocation in this algorithm is oblivious to the active server set.
- **TASP-LRH:** Least Recirculated Heat (LRH) is a heuristic TASP algorithm to choose the active servers such that the cooling power is minimized [3], [4], [14]. LRH ranks servers based on their static characteristics (i.e., servers' peak power consumption and their contribution on the heat recirculation in the data center room that is assumed to be statically calculated). The active server set in TASP when using LRH, is chosen among servers that their associated LRH ranking is the least. The idea of TASP-LRH algorithm here is to perform the phase allocation using a similar algorithm of the active server selection to avoid load imbalance due to server provisioning (note that optimal TASP needs to be aware of the dynamic of the workload and can not be achieved using LRH). This algorithm can be implemented in practice since the phase allocation is based on servers' LRH ranking which can be statically calculated. The algorithm uses "Round-Robin" (similar to TASP-Best) to allocate the low LRH ranking servers to the phases.

IX. Verification and Validation

In this section, we validate our arguments in Sections VI, VII, and VIII through experiments conducted in BlueCenter and through simulations using GDCSim [15], a steady state data center simulator.

TABLE I. BlueCenter phase imbalance experiment.

Phase	Phase Balanced		Phase	Phase Imbalanced	
	Machine	Power		Machine	Power
1	3	182.4 $(\pm 2.5)W$	1	3	182.4 $(\pm 2.5)W$
2	2	176.4 $(\pm 3.2)W$	1	2	179.8 $(\pm 0.5)W$
3	4	192.2 $(\pm 1)W$	1	4	186.2 $(\pm 1)W$

A. Validation

We use BlueCenter, a small scale data center facility located in the ASU campus [16]. BlueCenter is a NSF-funded project, which offers experimentation environment with innovative data center management schemes. The data center has physical dimensions of $27.6" \times 28" \times 11.8"$. There are total of 288 servers with total rated power consumption of 150kW supplied with electricity from the grid. The chiller supplies cold air with a flow rate of $5m^3/s$ from a single CRAC and follows a regular hot-cold aisle structure. There are 16 Servertech Switched CW-24VD/VY-3Ph PDUs, each delivers power to 18 Dell PowerEdge 1855 servers, with six servers in each of power-phases. The maximum current delivered by PDUs is 60 A.

We did two sets of experiments in the data center. In the first set, we selected three servers numbered 3, 2, and 4 on each phase of the three phase delta connected PDU. We shut down all other servers in the data center. We measured the power consumption of the server for a period of 1.5 minutes (one reading every 30 seconds) after allowing a 10 minute stabilization time. We then took the average of the five readings. The power consumption of the servers are comparable as shown in Table I. We then considered three servers on phase 1 that had identical power consumptions ($< 1\%$ difference) with the servers 3, 2, and 4. We kept these servers on and all other servers were turned off. We then measured the power consumption of these servers using the same methodology. Table I shows that due to phase imbalance there is a difference of 4 W. Thus, the percentage change in power consumption is $100 \times (4/551) = 0.62\%$. Surprisingly we found a decrease in power consumption due to phase imbalance for the same load. The PDU power sensors in data center typically report a power measurement accuracy of $\pm 1\%$. Since this variation is less than 1% we assume that the decrease in power consumption is due to inaccuracies in PDU power sensors. Thus, it shows that line loss is such small that it may not be measurable by power sensors in the PDU.

The second set of experiment we chose one PDU which had eight servers in each phase. We turned on all the servers and had a nearly balanced assignment with phase 1, 2 and 3 drawing 5.17 A, 5.81 A and 5.95 A, respectively. We observed a minimum power factor of 0.99. We then shutdown all the servers of phase 3 and measured the current at the server inlets and at the PDU inlet. The PDU inlet does not measure the power factor and hence only gives the apparent power. However, the server inlets measure the rms power and also

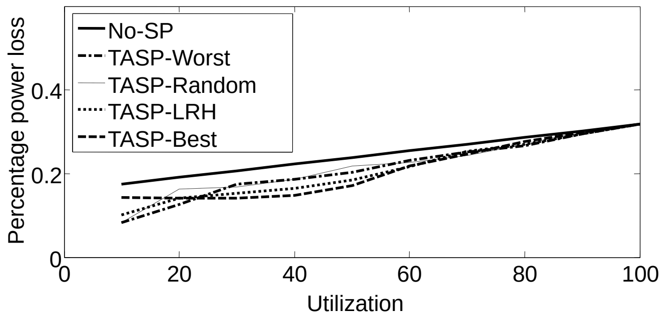

Fig. 4. **Percentage of power that is lost due to heating effect of line resistance for different phase aware workload scheduling algorithms.**

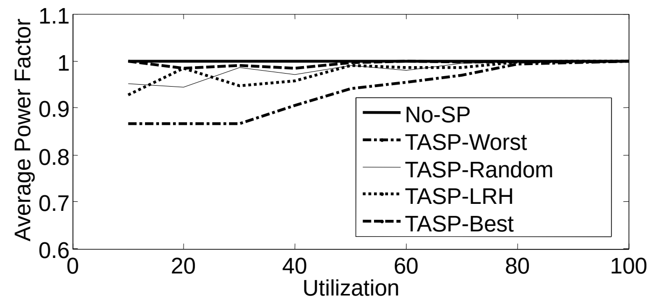

Fig. 5. **Power factor degradation at different utilizations due to phase imbalance caused by workload allocation.**

the power factor. With all servers in phase 3 turned off we obtained a power factor of 0.89. We then turned off the servers in phase 2 and 3 and measured the power factor for each phase. We obtained a minimum power factor of 0.87.

B. Verification

We simulate a data center similar to BlueCenter server and power infrastructure and study TASP along with the phase allocation algorithms. A data center with physical dimensions of $27.6' \times 28' \times 11.8'$ is considered. There are two rows of two 35U racks in each row, laid out in a typical hot-aisle/cold-aisle configuration. These racks consist of five 7U chassis each. There are 10 servers in each chassis which consume 1720 W in total when they all are idle and 3440 W when they all are operating at their peak. TASP performs server provisioning at the granularity of a chassis.

The cold air is supplied at a flow rate of $5m^3/s$ from a single cooling system (CRAC). The heat recirculation coefficients of servers are provided using BlueSim [15] which is used to design TASP and LRH [3], [4]. We also consider 4 PDUs for the entire data center where the maximum amperage delivered to each phase of a PDU is assumed to be $60A$. This translates into 6 chassis per PDU (i.e., two chassis per each phase of a PDU). We run simulations under a constant workload rate such that the utilization of the entire data center varies from 10% to 100% in 10% increment.

For each workload assignment technique, we compute the power loss due to line impedance using Equation 6 and the power factor using Equation 8. The total power loss due to line impedance is measured as a percentage of the total power consumed by the servers. Since there are four PDUs in the system we consider the overall power loss by summing up the power loss for all the PDUs. Thus, the percentage power loss for the data center, μ, is computed using Equation 9.

$$\mu = \frac{\sum_{i=0}^{N} P_{loss}^i}{\sum_{i=0}^{N} P_{del}^i} \tag{9}$$

where, N is the number of PDUs and P_{loss}^i and P_{del}^i are the power loss and delivered power per PDU.

The percentage power loss for all the phase aware workload allocation algorithms discussed in Section VIII for different utilizations are shown in Figure 4. In this simulation, we used a standard 208 V supply for the servers. Since the PDU rating is 60 A [17], the maximum current that each line of the PDU can deliver is $60/\sqrt{3} = 34.64A$. Thus, the maximum phase imbalance that each PDU can have is when one line is at maximum current and the other two is at near zero current. Further, we chose a line impedance of 0.04 ohms as suggested by several PDU data sheets [18]. Our simulation results show that line loss has a complicated relationship with phase imbalance but the magnitude of line loss is very minimal $< 1\%$, hence may not be a significant factor to consider for workload scheduling.

However, since line loss is considered a big concern in existing research on phase aware workload scheduling we discuss the characteristics of line loss in detail. We observe two very important characteristics from Figure 4.

1. Higher loss in balanced phase for low utilization (in this case very few number of active server are required): The TASP-Best algorithm which always tries to balance phases apparently has more line loss than all other TASP based algorithms which have higher phase imbalance.

2. Sweet utilization spot where the loss for balanced phase is minimum: From around 25% total data center utilization, where one fourth of servers need to be active, the TASP-Best algorithm has lowest line loss and the difference between the algorithms diminishes as utilization increases.

The main reason for this behavior is that the line loss is a non-linear function of both phase imbalance and current magnitude. If we consider the case where the data center is utilized at 10%, we have three chasis operating at 14.85 A each for a total load of 43.74 A to be distributed across four PDUs. The best case phase balance as obtained by the TASP-Best algorithm is to use one PDU and distribute the load equally to all three phases. This will result in a configuration where one PDU has 14.85 A load current in each phase and all other PDUs have zero load. The total line current as obtained from Equation 4 is 25.72 A on each line. This is a higher current magnitude than the case where we have imbalanced load. Thus the total line loss for all three lines is $3 \times 661.6 \times R_{line} = 1984.8R_{line}$ on one PDU and zeros on all others.

However, in the worst case phase imbalance we have a configuration where each of three PDUs have 14.85 A current on one phase and zero on the other two phases. The last PDU goes without a load. Thus, for each PDU total line current on any two lines is 14.85 and the third line has zero current. This is a lower current magnitude than the balanced case. Hence,

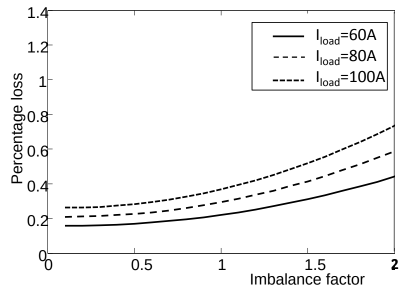

Fig. 6. **Line loss for different imbalance factor and load currents.**

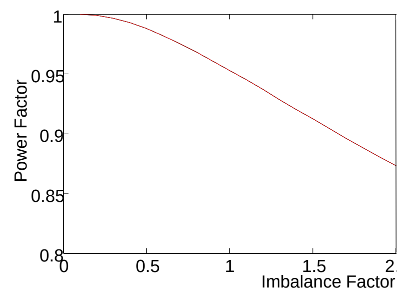

Fig. 7. **Power factor for different imbalance factors.**

the line loss on each of the first three PDUs is $2 \times 220.5225 \times R_{line}$ and the last PDU has zero loss. Thus, total power loss is $1323.135 R_{line}$, which is less than the balanced case!!

This happens due to the load consolidation on lesser number of PDUs, which increases the current on each line. But as the utilization of the data center increases the opportunities for load consolidation is reduced and the beneficial effects of phase balancing becomes more conspicuous. However, consideration of these trade-offs are subject to the significance of the loss. From our analysis we found out that a rigorous algorithm for phase balancing can at best reduce the losses by $< 1\%$, which is not a significant return of investment. We also computed the power factor degradation in at different utilizations for the different phase aware workload assignment algorithms. Figure 5 shows some obvious trends. The TASP-Best algorithm always achieves better power factors. However, the power factor gets closer to 1 as the load in the data center increases. The minimum power factor due to maximum imbalance as seen in our simulation results is 0.87 (i.e., the worst case power factor value for delta configuration).

To obtain a maximum value of the line loss we do simulation on a hypothetical case where we vary the load current from 60 A to 100 A in steps of 20 A. We characterize the phase imbalance using the imbalance factor as a parameter. The *imbalance factor* is defined as the deviation from average current and the current due to phase imbalance to the average current. We varied the phase imbalance factor from 0, balanced case, to 2, maximum imbalance case. The line loss and the power factor for these cases are shown in Figure 6 and 7. The results show that in the hypothetical worst case the line loss is 1.5% which is still insignificant and may not be measured by contemporary power meters which have an accuracy of 1% [18]. Also the minimum power factor obtained in the hypothetical case is 0.87 (i.e., the worst case power factor value for delta configuration).

X. Significance of Phase Imbalance

Phase imbalance causes several inefficiencies in power distribution of data center. In this paper, we considered the effect of phase imbalance on line loss, UPS power delivery efficiency and workload scheduling. Recent research have recognized energy loss due to phase imbalance as an important factor in data center operation. Researchers have even reported a savings as much as 14% by just scheduling workload so as to balance phases in a PDU. However, their assessment of the extent of the effects of phase imbalance is based on empirical observation and not on an objective measure through well established theory. In this paper, we consider the well established theory of three phase circuit and obtain an objective measure of the line loss, power factor degradation, and effect of workload scheduling on powe distribution efficiency of data centers. Below is a list of our findings:

Line loss due to phase imbalance does not contribute significantly to the overall power consumption - For a single PDU, although line loss increases due to phase imbalance in a non-linear fashion, it is typically less than 1% of the total power consumption. However, researchers report a maximum savings of 14% due to phase balancing [1]. One possible explanation of the 14% savings can be that the PDU input lines of the Servertech Switched CW-24VD/VY-3Ph unit used in their experiments (the same ones are used in BlueCenter) do not report power factor at the input. A fully imbalanced load causes input power factor to fall to 0.87 ($= \frac{\sqrt{3}}{2}$), and the apparent power to increase by 15%, thus leading to impression that power consumption has increased by 15%. In reality, power factor degradation does not cause increase in overall power drawn from the grid. However, the utility company may impose an additional cost if the overall power factor decreases below a specified threshold.

Power factor degradation has a significant effect on the efficiency of UPS power supply - Phase imbalance degrades the power factor of the PDU. This reduces the active power delivered by the UPS and hence the UPS supports lesser load. Researchers have considered this power factor degradation [2] effect and schedule workload so that the maximum phase imbalance factor does not exceed 20%. From our simulation results we find that in the worst case i.e., highest current rating a 20% imbalance factor gives a power factor of 0.98. From our BlueCenter experiments we observed that at maximum possible imbalance in a practical setting the power factor degrades to 0.89. Thus, assigning workload to balance phase does improve the power factor significantly and is important to reduce facility cost (reduces the need for higher capacity UPSes). In cases where the UPS capacity has been provisioned for the maximum load of the DC, the UPS will always suffice for the part of the load, even when it is phase imbalanced. For example, if a UPS is sized to power a DC having 300 servers, with 100 servers on each phase, it will still suffice if, say, only 200 servers are powered on, only on two phases, with one phase remaining completely unloaded.

However, If a setup has under-provisioned UPS, it can power more servers if the load is balanced, as the UPS capacity per phase is limited. In the above example, lets say the UPS is sized only for 270 servers. It can power maximum 90 servers on each phase. Obviously, we can power on a maximum of 270 servers, only if the load is fully balanced.

Workload scheduling in multiple PDUs is significantly different from scheduling in a single PDU - The power loss due to phase imbalance although very small has an interesting characteristics. Since the power loss depends on both magnitude and imbalance factor it is not always true that balancing three phases of a PDU will give minimum power loss specifically when we have multiple PDUs to which we can distribute workload (the practical case). If we consolidate the entire load into one PDU and distribute it evenly across the three phases, then the total current flowing through each line is one-third of the total load current. Instead if we distribute the load evenly across three maximally imbalanced PDUs then the current in each PDU is reduced leading to lesser line loss despite imbalance. This case is shown in Section IX through a numerical example.

For a single PDU, phase balancing always results in power savings as shown in previous research [1]. However, consideration of multiple PDUs makes the problem more difficult since the line loss due to phase imbalance is a non-linear function of the imbalance factor. Instead of minimization of phase imbalance researchers have suggested to include the phase imbalance as a constraint in the optimization framework to determine phase aware workload allocation. Such solutions should consider the constraint as a non-linear function leading to complexities in the solution.

In our experimental and simulation analysis we found that with simple solutions we can reduce the effect of phase imbalance to a large extent (i.e., improving power factor), and further savings with sophisticated optimization may lead to insignificant return of investment. This is especially true if the optimization aims to minimize line loss since it is only 1% of the total power consumption.

XI. Conclusions

Phase imbalance is considered as an important factor for data center workload and power allocation in recent research. Phase imbalances can cause line loss and also reduction in power factor, requiring an increase in UPS capacity. However, most of the research do not verify their hypothesis with objective measures of power factor and line loss. This paper considers a theoretical characterization of the line loss and power factor degradation due to phase imbalance and uses simulation and experiments on real data center deployment to evaluate the significance of these effects. Our results indicate that line loss is an insignificant fraction of the total power consumption and phase aware workload assignment can at best achieve 1.5% power savings. The worst case power factor observed in our experiments was 0.87, which matches with the theoretical calculations. The more dominant effect is the power factor degradation due to unbalanced phase which can increase losses in the UPS and decrease its efficiency. Particularly, in cases where the UPS capacity has been provisioned for the maximum load of the DC, the UPS will always suffice for the part of the load, even when it is phase imbalanced. However, if a setup has under-provisioned UPS, more servers can be powered if the load is balanced.

References

[1] P. Lama, Y. Li, A. M. Aji, P. Balaji, J. Dinan, S. Xiao, Y. Zhang, W. c. Feng, R. Thakur, and X. Zhou, "pVOCL: Power-aware dynamic placement and migration in virtualized gpu environments," in *The 33rd International Conference on Distributed Computing Systems*. IEEE, 2013.

[2] S. Pelley, D. Meisner, P. Zandevakili, T. F. Wenisch, and J. Underwood, "Power routing: dynamic power provisioning in the data center," in *ACM Sigplan Notices*, vol. 45, no. 3. ACM, 2010, pp. 231–242.

[3] Z. Abbasi, G. Varsamopoulos, and S. K. S. Gupta, "Thermal aware server provisioning and workload distribution for internet data centers," in *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing*, ser. HPDC '10. ACM, 2010, pp. 130–141.

[4] ——, "TACOMA: Server and workload management in internet data centers considering cooling-computing power trade-off and energy proportionality," *ACM Trans. Archit. Code Optim.*, vol. 9, no. 2, pp. 11:1–11:37, jun 2012.

[5] P. Ranganathan, P. Leech, D. Irwin, J. S. Chase, and H. Packard, "Ensemble-level power management for dense blade servers," in *In Proceedings of the International Symposium on Computer Architecture (ISCA*, 2006, pp. 66–77.

[6] Y. Chen, A. Das, W. Qin, A. Sivasubramaniam, Q. Wang, and N. Gautam, "Managing server energy and operational costs in hosting centers," *SIGMETRICS Performance Evaluation Review*, vol. 33, no. 1, pp. 303–314, 2005.

[7] D. Meisner, C. M. Sadler, L. A. Barroso, W.-D. Weber, and T. F. Wenisch, "Power management of online data-intensive services." in *ISCA*. ACM, 2011, pp. 319–330.

[8] D. Meisner, B. T. Gold, and T. F. Wenisch, "Powernap: eliminating server idle power," *SIGPLAN Not.*, vol. 44, pp. 205–216, March 2009.

[9] S. Govindan, D. Wang, A. Sivasubramaniam, and B. Urgaonkar, "Aggressive datacenter power provisioning with batteries," *ACM Transactions on Computer Systems (TOCS)*, vol. 31, no. 1, p. 2, 2013.

[10] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, "Energy storage in datacenters: what, where, and how much?" in *12th ACM SIGMETRICS*. ACM, 2012, pp. 187–198.

[11] V. Kontorinis, L. E. Zhang, B. Aksanli, J. Sampson, H. Homayoun, E. Pettis, D. M. Tullsen, and T. S. Rosing, "Managing distributed ups energy for effective power capping in data centers," in *39th International Symposium on Computer Architecture*. IEEE Press, 2012, pp. 488–499.

[12] D. S. Palasamudram, R. K. Sitaramanz, B. Urgaonkar, and R. Urgaonkar, "Using batteries to reduce the power costs of internet-scale distributed networks," in *Proceedings of 2012 ACM Symposium on Cloud Computing*. ACM, oct 2012.

[13] J. M. Gurrero, L. G. de Vicuna, and J. Uceda, "Uninterruptible power supply systems provide protection," *Industrial Electronics Magazine, IEEE*, vol. 1, no. 1, pp. 28–38, 2007.

[14] T. Mukherjee, A. Banerjee, G. Varsamopoulos, S. K. S. Gupta, and S. Rungta, "Spatio-temporal thermal-aware job scheduling to minimize energy consumption in virtualized heterogeneous data centers," *Computer Networks*, vol. 53, no. 17, pp. 2888 – 2904, dec 2009.

[15] S. K. S. Gupta, R. R. Gilbert, A. Banerjee, Z. Abbasi, T. Mukherjee, and G. Varsamopoulos, "GDCSim - an integrated tool chain for analyzing green data center physical design and resource management techniques," in *International Green Computing Conference (IGCC)*, 2011.

[16] S. K. S. Gupta, G. Varsamopoulos, A. Haywood, P. E. Phelan, and T. Mukherjee, "Bluetool: Using a computing systems research infrastructure tool to design and test green and sustainable data centers," in *Handbook of Energy-Aware and Green Computing*, 1st ed., I. Ahmad and S. Ranka, Eds. Chapman & Hall/CRC, 2012.

[17] "Switched pdu, http://www.servertech.com/products/switched-pdus/switched-pdu-cw-24vd-vy-3ph."

[18] B. Davies, "Analysis of inrush currents for dc powered it equipment," in *Telecommunications Energy Conference (INTELEC), 2011 IEEE 33rd International*, 2011, pp. 1–4.