

Dynamic Allocation of Power Delivery Paths in Consolidated Data Centers based on Adaptive UPS Switching

Fawaz AL-Hazemi^{1,2}, Yuyang Peng¹, Chan-Hyun Youn¹, Josip Lorincz³, Chao Li⁴, Guo Song⁵, Raouf Boutaba⁶

¹*School of Electrical Engineering, Korea Advanced Institute of Science and Technology, Daejeon, South Korea. E-mail: fawaz,yuyangpeng,chyoun@kaist.ac.kr*

²*College of Computer Science and Information Technology, University of Prince Mugrin, Madinah, Saudi Arabia. E-mail: f.alhazemi@upm.edu.sa*

³*Faculty of Electrical Engineering, Mechanical Engineering and Naval Architecture (FESB), University of Split, Croatia, E-mail: josip.lorincz@fesb.hr*

⁴*Department of Computer Science and Engineering at Shanghai Jiao Tong University (SJTU), China. E-mail: lichao@cs.sjtu.edu.cn*

⁵*Department of Computing, The Hong Kong Polytechnic University, Hong Kong. E-mail: song.guo@polyu.edu.hk*

⁶*Faculty of Computer Science at the University of Waterloo, Canada. E-Mail: rboutaba@uwaterloo.ca*

Abstract

Although technique known as server consolidation approach in a data center can reduce the overall power consumption, the Power Usage Effectiveness (PUE) of the data center will still be negatively affected with presence of distributed Uninterruptible Power Supplies (UPSs). The impact on the PUE arises from the fact that all UPS modules are kept running to maintain power availability for only a few active servers during off-peak periods. To address this problem, in this paper technique for reducing power consumption in a data center by consolidating the UPSs used during off peak periods is proposed. The proposed technique achieves power savings by leveraging a micro Automatic Transfer Switch (micro-ATS) at the server end. The novelty of this work lies in developed adaptive algorithm that continuously looks for opportunities to reduce the number of UPSs by offloading under-loaded UPSs to a neighboring UPS whenever that neighboring UPS can handle the extra load. In various simulated scenarios involving corporate data centers, our approach demonstrates the ability to save more power and achieve lower PUE degradation compared with state-of-the-art approaches such as server consolidation. Specifically, the proposed approach achieves a savings of approximately 20% to 40% in a data center's power consumption, depending on the data center's off-peak periods, which can be accomplished using only 80% of the UPS modules in the data center.

Keywords: Data center, Energy efficiency, Power management, UPS

1. Introduction

Approximately 10% of world's power consumption is due to Information and Communication Technology (ICT), and according to Koomey's report [1], approximately 14% of that power is

Preprint submitted to Computer Networks *August 15, 2018*

consumed by data centers. Small, medium, and corporate data centers account for 95% of the total data center power consumption because of their inefficient usage of electric power [2]. The urgency of prioritizing this inefficient electricity usage among the various ICT challenges to be addressed is illustrated by a recent survey conducted by *Uptime Institute*, which shows that 65% of corporate workloads are processed by on-premises data centers [3].

We highlight the following causes of the inefficient usage of electric power: First, the majority of servers in these data centers are under-utilized or idle during off-peak hours. Second, the legacy design of rack-based energy storage backups used in data centers [4] is increasing the complexity of addressing power inefficiency. The problem of under-utilized or idle servers can be mitigated by applying the server consolidation technique [5, 6]; however, the legacy design of data centers (as shown in Fig. 1) apparently exacerbates the power wastage at the level of the Uninterruptible Power Supplies (UPSs). Moreover, the size of the data center along with the legacy design is significantly contributing to the implication of electricity power waste. In fact, the larger the data center in size the higher possibility that server could be turned off during off peaks; which increases the power wastage at UPSs. Hence, as the size of the data center is increased, we are expecting an increase of power wastage. Therefore, toward preserving the benefits of saving energy through consolidating the number of active servers in data centers, we have to coordinate and couple that with consolidating data center's UPSs.

Leveraging recent contributions in energy and power savings such as server consolidation for corporate data centers, while keeping the UPSs operating at the maximum possible load is a challenge. To address this challenge, it is necessary to extend the server consolidation technique [5, 6] to include UPS consolidation. In this way, it is possible to dynamically adjust the required number of operational UPSs depending on the number of active servers in the data center. To do so, modification of the electricity power supply delivery path at the server end and application of suitable control flow scheme to monitor each UPS must be done. Hence, in this paper, technique for reducing power consumption in a data center by consolidating the UPSs used during off peak periods is proposed.

The remainder of this paper is organized as follows. In sections 2 and 3, we briefly describe the background on UPS limitations and present relevant contributions from related works, respectively. Next, in section 4, we mathematically analyze the impact of the UPSs on a data center's Power Usage Effectiveness (PUE). Then, in section 5, we introduce proposed UPS consolidation approach. Subsequently, in section 6, we present a preliminary evaluation, followed by a discussion of the outcomes in section 7. We discuss possible trade-off issues related to UPS consolidation design in section 8. Finally, we conclude the paper with closing remarks in section 9.

2. Background

2.1. Challenges in Data Centers

The power consumption of a data center is influenced by users' activities, which pose various challenges in data center power management, as illustrated in Fig. 1 (left side). The first potential challenge is based on the fact that users generate (predictable or unpredictable) workloads for the applications hosted in the data center; these workloads place servers and interconnected (network) appliances in the busy state. Depending on the workloads assigned to the hosted applications, these servers and network appliances will experience certain levels of utilization. For instance, if the user workload is high, these servers and network appliances are fully utilized;

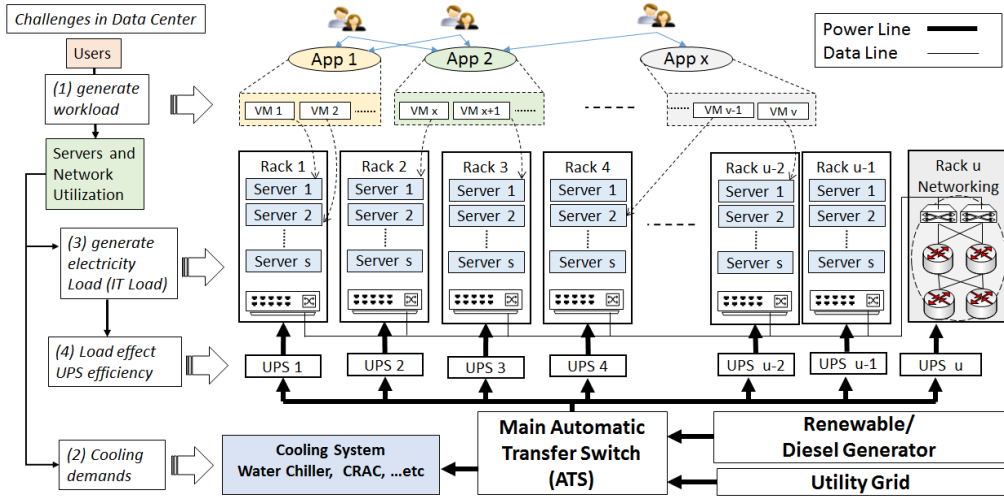


Figure 1: Data center infrastructure design and power delivery paths.

otherwise, they are under-utilized. The second potential challenge is that the data center's servers and network appliances (or IT resources) generate an electrical power load depending on their utilization. This electrical power load is regarded as direct power consumption by the data center. The third potential challenge is that depending on their utilization, these servers and network appliances also gradually generate heat, which will ultimately require substantial cooling efforts. The cooling demand is another source of electrical power consumption; although it represents indirect power consumption and some means of free cooling might be available, it still plays a role in the overall power consumption of the data center.

The challenges of data center power management mentioned above have been considered in the literature for the past decade, and they have been approached from two major aspects. One aspect is data center workload management, which focuses on the first challenge (user workloads), and the other aspect is the PUE, which focuses on the other two challenges (IT load and cooling demands). On the one hand, workload management strategies have been widely proposed for managing the power consumed by data centers by means of (1) workload balancing and re-distribution among servers and network appliances within a single data center, (2) server consolidation (reducing the number of servers) within a single data center, and (3) workload admission control and geo-distributed data centers. On the other hand, a number of proposals have addressed the PUEs of data centers by means of (1) free cooling and water chillers to address cooling demands, (2) thermal and hot spot management with assistance from server consolidation, and (3) renewable energy sources and energy storage devices.

Although the above contributions have had a significant impact on the power consumption challenges posed by data centers, the existing literature seems to ignore another power saving option, namely, the efficient management of electrical power appliances. In a data center, the power appliances are essential components on the power delivery path because of their high availability characteristics. These appliances have certain power delivery efficiency characteristics related to their electrical loads. Therefore, in this paper, we consider the following potential power consumption challenge in a data center: the efficiency of the power appliances that supply the electrical power for the generated IT loads. In the following subsection, we briefly present the background on these power appliances and provide more details on UPSs.




Figure 2: Electrical conversion along the power delivery path used in a data center.




Figure 3: UPS efficiency curve [8]

2.2. Uninterruptible Power Supplies (UPSs)

The equipment operating in a data center typically includes Power Distribution Units (PDUs), Automatic Transfer Switches (ATSs), and UPSs. All of them have certain power losses depending on their efficiencies, but the most critical of these is the UPS efficiency, as will be explained in detail in this subsection.

UPS systems, particularly those of the double-conversion type, are very important for handling power surges in data centers because of their high availability. However, this type of UPS unit has a power loss similar to that of any other electrical power appliance in which electrical power is converted from Alternate Current (AC) to Direct Current (DC) to charge batteries and from DC to AC to power racks and servers in the data center, as shown in Fig. 2. In addition, this type of UPS has a bypass circuit that will activate when the UPS system is overloaded to protect its batteries from damage. It has been shown that efficiency curve of double conversion UPSs depends on applied electrical load (Fig. 3). Hence the double-conversion UPS design limits the efficiency of the delivery of electric power based on the applied load, i.e., the rack server load, as reported by Zhang and Shi in [7] and [8] and as shown in Fig. 3. Consequently, AC/DC conversion results in a significant amount of wasted power, especially when the IT (server or network appliance) load is low.

Besides significant amount of power consumption in a data center due to double conversion topology (refer to Fig. 3), according to a study released in 2016 by *Emerson* and the *Ponemon Institute*, UPSs also ranks among the top causes of data center outages [9]. This study [9] reported that UPS systems are considered to be the leading cause of unplanned outages in data centers and are estimated to be responsible for one-fourth (25%) of all data center outages. On the other hand, the presence of UPS modules on the electrical power delivery path is vital for power failover. However, this need has become a major bottleneck hindering the expansion of corporate data centers. The power consumption complexity introduced by UPS systems might be considered straightforward; however, the existing literature seems to ignore both the power waste caused by

UPSs and the potential for power surges. As an example of the potential implications of such power surges, industrial sectors such as airlines are routinely suffering from power surges caused by UPS systems [10]. To date, industries such as airlines, which rely on back-end data centers for their core business (airport operations and aviation), remain very susceptible to service outages due to UPS problems. Two recent examples of airlines that have suffered data center power outages caused by UPSs are Delta Airlines (August 2016) and British Airways (May 2017) [10]. The British Airways (May 2017) incident left more than 75,000 passengers in London without the ability to travel, whereas the Delta Airlines (August 2016) incident prevented hundreds of flights from operating. Both incidents were rooted in UPS systems [10].

3. Related Works

In the literature, there are diverse contributions that share common goals related to power consumption, energy consumption, energy awareness, energy storage, power shaving (or peak shaving), and green data centers (refer to Fig. 1). Here, we focus on works closely related to the scope of this paper, such as energy and power efficiency using UPSs and energy storage techniques.

The literature on UPS efficiency techniques has noticeably expanded in recent years, and this increased attention to UPSs in data centers can be related to the following three aspects of the problem. The first is when power is over-subscribed or over-provisioned [11]. The second is the realization of “hot & fast” energy storage to help a data center to avoid a high peak electricity price or to address the difficulties faced by a data center that is powered by renewable energy resources [12, 13]. The final aspect is the power efficiency (or power loss) of the UPS systems in a data center [7, 8].

The PUE is not ideal energy efficacy metric for data centers, because it does not take into account other elements of data centers such as power distribution and cooling losses inside IT equipment. In the [15], authors propose two new metrics for addressing energy efficacy benchmarking of data centers as communication network entity. However, this work addresses dynamic allocation of power delivery paths for servers during consolidation process in data centers. Proposed dynamic allocation of power delivery paths based on adaptive switching of UPSs is as a concept related to only server consolidation. This consolidation process does not take into account other data center equipment. Since PUE as the metric is broadly accepted for expressing data center energy-efficiency from perspective of servers as main energy consumers, the PUE is selected as metric in this paper because it is the most appropriate for presenting effects of adaptive UPS switching.

3.1. Power Over-Provisioning and Capping

Power over-provisioning refers to the case in which a data center is run on an existing power infrastructure and a power-capping technique is implemented to avoid power budget violations [14, 16, 17, 18, 19, 20, 21]. Petoumenos et al. [22] reviewed the most popular power-capping techniques that are implementable at the server level. Islam et al. [14] proposed a coordinated power manager (called COOP) that uses supply function bidding to handle power violations in over-provisioned multi-tenant data centers. The supply function used in COOP [14] rewards tenants by reducing their power consumption during power emergencies. Liu et al. [16] and Zheng et al. [17] suggested the use of a hybrid UPS and supercapacitor (SC) system to achieve over-provisioning and peak shaving for data center power management. Other power-capping

techniques used in data centers include hierarchical power-capping designs such as Facebook’s Dynamo [23], which implements higher levels of power coordination and distribution at the data center level and uses leaf agents to cap the power at the server level.

3.2. Energy Storage as an Enabler

“Hot & fast” energy storage can be used to enable more effective power management. Kontorinis et al. [11] suggested a distributed per-server UPS system with the ability to store electrical power during times of low demand (i.e., low workload on the servers) and to supply that stored electricity during times of peak power demand (i.e., peak workload on the servers). This approach is known as Google’s Customized server [24], whose configuration differs from standard server configurations, as shown in Fig. 4. Recently, a distributed per-server UPS system was proposed as part of the Power Attack Defense (PAD) design by Li et al. [25], which protects data center servers from any potential power threat. Kontorinis et al. [11] adopted a distributed UPS system for peak power shaving; by contrast, Aksanli et al. [26, 27] suggested a modified version of the architecture used by Kontorinis et al. [11] in which a central UPS system is used for peak power shaving.

Aksanli et al. [26, 27] introduced a grid-tie inverter into the power path design for a data center to allow the batteries of the central UPS system to be charged during off-peak times. Similarly, grid-tie inverters have also been suggested in other works [11, 12, 13, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37] to enable the utilization of renewable energy sources in data centers. Deng et al. [28] suggested using a grid-tie approach to power a data center with both utility grids and renewable energy, and they defined a Renewable-Powered Instance (RPI) in place of a server. Goiri et al. [29, 30, 31, 32] presented a similar grid-tie approach to address renewable intermittency, and they thoroughly investigated the utilization of solar panels to power small data centers (in a system design called *Parasol* [31]). The grid-tie approach has also been used in other proposals, such as Blink [38], SolarCore [39], and iSwitch [40]. However, Li et al. [12, 33, 34, 35] considered the use of grid-tie inverters with the UPS system topologies used in data centers and suggested a novel expansion strategy (called Oasis [12]) as an alternative means of utilizing renewable energy sources for electrical power while simultaneously minimizing the cost of scaling up a data center.

Meanwhile, other proposals have suggested the utilization of under-provisioned UPS systems instead of diesel generators [42]. Recent studies have demonstrated the effectiveness of utilizing fuel cells to power a data center [43, 44]; however, these fuel-cell-powered data centers require power capping. One limitation of fuel cells is their slowness in adapting to changes in power demands; therefore, Li et al. [18] considered the use of an SC to address such power surges and suggested a framework called SizeCap to address these surges through power-capping policies. These suggestions, which incorporate the use of UPS systems to power the data center for short periods of time (during renewable energy shortages or to assist in power peak shaving) and the use of a grid-tie inverter to recharge their batteries, are lacking in terms of the major limitations on UPS power losses and PUE degradation.

3.3. UPS Efficiency-Aware Control

The power efficiency (or power loss) of UPS systems is also crucial for power management in data centers. Zhang and Shi [7, 8] studied UPS efficiency and used it as an indicator (or quality factor) for efficient workload distribution in a data center. To the best of our knowledge, Zhang and Shi [8] practically benchmark a rack-level Double-Conversion UPS module used in a university data center. The outcome of their benchmark demonstrated a significant impact of UPS

efficiency in data center power consumption, precisely the power loss caused by UPSs. Their benchmark shows a relationship between the applied IT load on UPS and the UPS efficiency, Fig. 3. They managed the workload distribution in an effort to improve performance when running a mixture of applications; however, they did not consider server consolidation or servers with modified settings, such as CPU dynamic voltage and frequency scaling (DVFS) or server virtualization.

Another recent work, although one that is not focused on UPS efficiency, is the framework suggested by Islam and Ren [45], who proposed a policy for sharing the power consumed by non-IT systems (including UPS system power losses) among tenants in a multi-tenant data center. The suggested policy is based on the *Shapley value* from game theory, and they were able to divide the cost of power consumed by non-IT systems over multiple tenants; however, they could not reduce the power losses of inefficient UPSs.

In this paper, we focus on reducing UPS power losses by consolidating the number of active UPS modules used during off-peak periods and taking advantage of recent works on server consolidation for power savings.

4. Trade-offs in the UPS Consolidation Strategy

To the best of our knowledge, the application of UPS consolidation in an existing data center in combination with a power-saving technique such as server consolidation has not previously been proposed. The existing literature seems to ignore the possibility of reducing energy consumption by switching UPSs on/off in accordance with their workload-dependent efficiency; consequently, there is no reference model available with which to compare our approach. Therefore, in this section, we highlight possible trade-offs in the UPS Consolidation strategy.

In general, having a UPS Consolidation strategy would ensure that a UPS is always operating in its rated load state (Fig. 3), so there is no concern about shortening a given UPS's life. However, in the case of UPS Consolidation, it may occur that a single UPS unit is attached to two neighboring racks and a high workload is experienced during an off-peak time (in the case of an unexpected workload). When the load is higher than the rated capacity of the UPS, another UPS feature activates automatically via a built-in electric circuit bypass switch (as shown in Fig. 2). This built-in bypass switch protects the UPS and its battery from damage while continuing to feed electricity to the attached load without interruption.

The proposed UPS Consolidation strategy tends to connect (at most) two neighboring racks to the same UPS to save energy; however, it does not result in the deactivation of at most half (50%) of the data center's UPSs. In practice, our UPS Consolidation strategy deactivates UPSs in accordance with the number of active servers in the data center, which is controlled via the server consolidation approach, i.e. a higher consolidation layer. For example, if there are 4 racks, each has its own set of servers and one UPS unit; that is, there are 4 UPS units. If the number of active servers is reduced by half (from 4 racks to 2 racks) via server consolidation due to a sufficiently low workload, it is possible that our UPS Consolidation strategy might result in both of these active racks being operated using a single UPS (if the rated capacity of the UPS is sufficient). Therefore, the UPS Consolidation strategy could result in 3 UPS units being shut off which is 75% rather than 50% deactivation rate. For that reason, placement of servers and UPSs in relation to energy efficiency is important. Approach proposed in this work requests appropriate pairing of UPSs with corresponding servers in the racks.

The transitions between activation and deactivation for UPS modules can be fast. Unlike servers, a UPS does not require a booting procedure (powering on followed by the loading of

an operating system and firmware) to enter operation; therefore, activation (or deactivation) is simply a matter of attaching (or detaching) the UPS to (or from) the power delivery path. This attachment/detachment of a UPS is simply achieved based on a mechanical switch (electric circuit breaker/switch), which can be maintained remotely using a programmable interface such as *Yocto-Watt* from *Yoctopuce* [46]. Definitely, adding small electrical device as breaker/switch, i.e. *Yocto-Watt*, would have its own overhead cost on the proposed approach, nonetheless, for neutrality, we suggested the use of *Yocto-Watt* to be a remote breaker/switch in the proposed approach. Modern UPS modules could have a built-in breaker/switch that is accessible by Application Programmable Interface (API); which save the data center owners from such extra overhead costs.

N.B. UPS must supply all types of IT equipment including servers and interconnecting devices such as firewalls, load balancers, routers, switches, intrusion detections/prevention systems, etc. Apparently, during the data center operation, demand for constant activity of interconnecting devices is higher than for regular servers and applying a consolidation mechanism on these devices is less appropriate. In fact, servers are the most appropriate for consolidation in the data center; which means the consolidated (or reduced) portion of power consumption is coming from shutting down servers. Therefore, in this paper, we considered the power consumption of servers as $P_{IT,u}$.

5. Analyses of the UPS impacts

Distributed UPS modules have two main implications in data centers. First, a UPS has a non-negligible impact on the PUE of a data center. Second, UPS efficiency has an inversely proportional relationship with the number of supplied active servers in a data center. These two impacts are discussed and expressed mathematically in the following. All abbreviations used here and hereafter are listed in Table 1.

5.1. UPS Efficiency and Power Loss

A double-conversion UPS system has a disadvantage in terms of the efficiency of power conversion (from AC to DC and then back to AC) for providing electricity to the supplied load (i.e., servers). To date, double-conversion UPS modules have achieved a maximum power conversion efficiency of approximately 85 to 94%, and the recent Eaton UPS rack-based module has achieved a 96% conversion efficiency [4]. Let us assume that data center has set U of UPSs in order to be:

$$U = \{1, \dots, u, \dots, U\} \quad u \in \mathbb{N} \quad (1)$$

where U represents total number of UPSs in a data center. According to Zhang and Shi [8], UPS efficiency ($UPS_{eff,u}$) can be modeled as shown in Eq. (2):

$$UPS_{eff,u} = \alpha_u \times \ln\left(\frac{P_{IT,u}}{UPS_{Rate,u}}\right) + \beta_u, \quad (2)$$

where $UPS_{eff,u}$ is the efficiency of UPS u , $P_{IT,u}$ is the IT power load generated by the rack servers connected to UPS u , $UPS_{Rate,u}$ is the rated power of UPS u (the maximum power that can be supplied), and α_u and β_u are estimated coefficients of the model.

N.B. UPS must supply all types of IT equipment include servers and interconnecting devices such as switches, routers, firewalls, intrusion detections/preventions, load balancers... etc. Apparently, during the data center operation, interconnected devices are on higher demand than regular servers. Applying a consolidation mechanism on these devices can jeopardize its availability. In fact, only servers that could freely be consolidated in the data center; which means the consolidated (or reduced) portion of power consumption is coming from shutting down servers. Therefore, in this paper, we considered the power consumption of servers as $P_{IT,u}$

The UPS efficiency model expressed in Eq. (2) includes a ratio that takes a value between 0 and 1; this ratio is used to measure the amount of power that is lost at the UPS system due to conversion, as given below by Eq. (4):

$$P_{UPS,u} = \frac{P_{IT,u}}{UPS_{eff,u}} - P_{IT,u} \quad (3)$$

$$= \left(\frac{1 - UPS_{eff,u}}{UPS_{eff,u}} \right) P_{IT,u} [W], \quad (4)$$

Where the power loss caused by the UPS ($P_{UPS,u}$) is the calculated power quantity. Let us assume that each UPS supply a set of servers \mathbf{S}_u in order to be:

$$\mathbf{S}_u = \{s_{1u}, \dots, s_{iu}, \dots, s_{Mu}\} \quad i = 1, \dots, M \in \mathbb{N}, u \in \mathbf{U} \quad (5)$$

where M is the maximum number of servers in the set \mathbf{S}_u supplied by UPS u . The IT power load ($P_{IT,u}$) is the total power of the rack servers supplied by the UPS module, as follows

$$P_{IT,u} = \sum_{i=1}^M P_{S_{i,u}} \quad [W] \quad \forall u \in \mathbf{U} \quad (6)$$

Different servers use versatile amount of resources, and set of each server resources can be expressed as:

$$\mathbf{R}_{i,u} = \{r_{i,j,u}, \dots, r_{j,i,u}, \dots, R_{N,i,u}\} \quad j = 1, \dots, N \in \mathbb{N}, \quad i = 1, \dots, M \in \mathbb{N}, \quad u \in \mathbf{U} \quad (7)$$

where N is the maximal number of resources used by the server i supplied by UPS u .

Accordingly, server i located in rack u has instantaneous power consumption impacted by amount of used resources which can be expressed as:

$$P_{S_{i,u}} = P_{S_{i,u},Idle} + \sum_{j=1}^N (\alpha_{r_{j,i,u}} \times Ur_{j,i,u}) \quad [W] \quad \forall u \in \mathbf{U}, \forall i = 1, \dots, M \in \mathbb{N} \quad (8)$$

where $P_{S_{i,u},Idle}$ is the server's power consumption in the idle state, and $Ur_{j,i,u}$ is the utilization of resource j of server i supplied by UPS u . In addition, $\alpha_{r_{j,i,u}}$ is the estimated linear regression parameter for that resource, which is determined using system identification methods from control theory [47].

5.2. Impact of UPSs on the PUE

The PUE is a well-known efficiency metric developed by The Green Grid consortium [48] (For more details, the reader is referred to [55]). This metric is defined as the ratio between the

Table 1: Parameter and variable notations

Parameter	Description
$P_{UPS,u}$	Power loss caused by UPS u
$UPS_{Rate,u}$	Rated power for UPS u (maximum power capacity)
$UPS_{eff,u}$	Power conversion efficiency of UPS u
$UPS_{eff,u,Max}$	Maximum power conversion efficiency of UPS u
α_u, β_u	Estimated (logarithmic trending) parameters for the efficiency modeling of UPS u
R	Number of racks in the data center
$P_{IT,u}$	Power load due to active servers in the rack that is supplied by UPS u
P_{Total}	Total IT power load in the data center
$Ps_{i,u}$	Power consumed by server i supplied with UPS u
$Ps_{i,u,Idle}$	Power consumed by server i supplied with UPS u in the idle state
Ur_{iju}	Utilization of resource r of server i supplied by UPS u
αr_{iju}	Estimated (linear regression) parameter for resource r of a server i
RA_{upss}	Ratio between the total (all) UPS power loss and the total (all) IT power load
$RA_{ups,u}$	Ratio between the power loss of a single UPS u and the IT power load of the set of servers supplied by that UPS
$RA_{ups,u,MIN}$	The minimum ratio between the power loss of a single UPS u and the IT power load of the set of servers supplied by that UPS
Pb_u	Power budget for the rack supplied by UPS u
$P_{IT,u,NB}$	Power load imposed by active servers on neighboring racks of the rack supplied by UPS u
$excess$	Surplus power after the subtraction of the $load$ from a rack power budget Pb_u : $excess = Pb_u - load$
$Rack_{OL}$	Binary value indicating whether a rack's UPS is overloaded by a neighboring rack (0: not loaded, 1: loaded)

amount of electrical power (in watts) used by IT equipment (e.g., servers) and the total amount of electrical power consumed by the overall facility, i.e.,

$$PUE = \frac{P_{Total}}{\sum_{u=1}^U P_{IT,u}} \quad (9)$$

where the summation of all $P_{IT,u}$ represents the power consumed by all servers supplied over U different UPSs located in data center.; and P_{Total} is the total power consumed by the data center, including both IT and non-IT (facilities) equipment, and is defined in Eq. (10):

$$P_{Total} = \sum_{u \in U} P_{IT,u} + \sum_{u \in U} P_{UPS,u} + P_{Cooling} + P_{Others}, \quad (10)$$

where the summation of all $P_{UPS,u}$ represents the power loss of the multiple UPS systems, $P_{Cooling}$ is the power loss of the thermal control system, and P_{Others} is the power consumed by other elements in the data center. P_{Others} is referred to the amount of power consumed by non-IT equipment that are not mentioned in the equation (10), such as the power consumed by Main Switching Board (MSPs) for electricity power, lights, fire alarms and fire-detection sensors. . . etc. These equipments are ignorable compared to the ones considered in the equations, i.e. Cooling system and UPS, therefore we included them under single notation P_{Others} .

By substituting Eq. (10) into Eq. (9), we obtain

$$PUE = \frac{\sum_{u \in \mathbf{U}} P_{IT,u} + \sum_{u \in \mathbf{U}} P_{UPS,u} + P_{Cooling} + P_{Others}}{\sum_{u \in \mathbf{U}} P_{IT,u}} \quad (11)$$

or

$$PUE = \sum_{u \in \mathbf{U}} \frac{\overbrace{P_{IT,u}}^{RA_{load}=1}}{P_{IT,u}} + \frac{\overbrace{\sum_{u \in \mathbf{U}} P_{UPS,u}}^{RA_{upss}}}{\sum_{u \in \mathbf{U}} P_{IT,u}} + \frac{\overbrace{P_{Cooling}}^{RA_{cooling}}}{\sum_{u \in \mathbf{U}} P_{IT,u}} + \frac{\overbrace{P_{PowerSupply}}^{RA_{others}}}{\sum_{u \in \mathbf{U}} P_{IT,u}}, \quad (12)$$

where the superscripts (RA) denote the individual ratios between the power losses from each power consumer in the data center and the total IT power load (of the rack servers) P_{IT} :

$$P_{IT} = \sum_{u=1}^U P_{IT,u} [W]. \quad (13)$$

Although the efficiency of the cooling system also has a major impact on PUE improvement, in this paper, we restrict our consideration to the implications of UPS system efficiency in data centers. Therefore, we assume that the ratios between the power losses from other non-IT power consumers (except for the UPS systems) and the total IT power load, i.e., $RA_{Cooling}$ and RA_{Others} , are fixed but non-zero ratios. By contrast, the ratio between the UPS power loss (RA_{upss}) and the total IT power load is not fixed because of the inversely proportional relationship between them; this is our subject of study in this paper.

It is not possible to generalize the precise proportion of the power loss that is caused by UPS systems in terms of the PUE. For example, let us assume a data center with a PUE of 1.7 (which is the reported average PUE among data centers surveyed in [49]); this tells us that 41% of the total power consumed in the data center is consumed by non-IT equipment. This 41% cannot simply be assumed to be equally distributed among all non-IT equipment because each piece of non-IT equipment has its own operation mechanism and power consumption trends. Hence, we consider the efficiency of the UPS systems and their contribution to the UPS-to-load ratio (RA_{upss}) as follows:

$$RA_{upss} = \frac{\sum_{u \in \mathbf{U}} P_{UPS,u}}{\sum_{u \in \mathbf{U}} P_{IT,u}}. \quad (14)$$

We separate the above ratio RA_{upss} expressed in Eq. (14) into a per-UPS power loss by substituting the expression given in Eq. (4) for $P_{UPS,u}$ as follows:

$$RA_{upss,u} = \frac{\left(\frac{1-UPS_{eff,u}}{UPS_{eff,u}}\right)P_{IT,u}}{P_{IT,u}} \quad (15)$$

$$= \frac{1-UPS_{eff,u}}{UPS_{eff,u}} \quad \forall u \in \mathbf{U}, \quad (16)$$

where each $RA_{upss,u}$ is the ratio of the power loss of UPS u to the power load fed to its rack's servers. The ratio $RA_{upss,u}$ takes values in the range of $(0, \infty)$; a ratio that is closer to 0 indicates a better UPS efficiency, and vice versa.

5.3. Problem Definition

We divide the problem presented in this paper into two parts as follows. First, we define the UPS power loss problem and its relationship with UPS efficiency. Second, we identify the correlation between server consolidation (shutting down servers) and UPS efficiency.

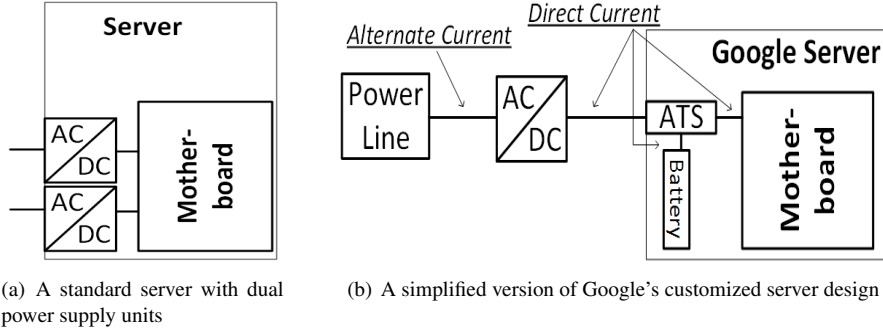


Figure 4: A standard server configuration and Google's customized server configuration.

5.3.1. UPS Power Loss

The following statement is generally understood to be true for UPS systems. *Any active UPS module has a small amount of power loss and (consequently) less impact on the data center's PUE.* However, in this paper, we strongly argue that UPS modules cause an unavoidable amount of power loss in a data center, which will exert a significant impact on the data center's PUE. Our proposition is as follows. First, the ratio $RA_{ups,u}$ expressed in Eq. (16) has the following property:

Property 1: *The ratio $RA_{ups,u}$ has a lower bound or minimum value (denoted by $RA_{ups,u,MIN}$) that is greater than zero:*

$$RA_{ups,u,MIN} > 0. \quad (17)$$

PROOF. See Appendix A.1.

Definition 1 (UPS Power Loss - P_{UPS}): *Any active double-conversion UPS module that is either fully or partially loaded with a proper electrical power load (i.e., an IT load or physical server load) must have a non-negligible power loss that will affect the PUE of the data center.*

Lemma 1: *Consider a double-conversion UPS module with maximum capacity (rated power) of $UPS_{Rate,u}$ and a maximum efficiency of $UPS_{eff,u,MAX}$ (in %). Any IT power load ($P_{IT,u}$) applied to the UPS system will result in a UPS power loss ($P_{UPS,u}$) that is equal to or greater than the IT power load ($P_{IT,u}$) times the minimum UPS-to-load ratio ($RA_{ups,u,MIN}$) for that UPS module, as shown in inequality (18):*

$$P_{UPS,u} \geq RA_{ups,u,MIN} \times P_{IT,u}. \quad (18)$$

PROOF. See Appendix A.2.

To apply property 1 and definition 1 to a real-world example, let us consider the latest Eaton UPS rack-based module [4], which has an efficiency of 96%, that is, $UPS_{eff,u,MAX} = 0.96$. By property 1, the minimum value $RA_{ups,u,MIN}$ of the ratio defined in Eq. (16) is

$$RA_{ups,u,MIN} = \left(\frac{1 - 0.96}{0.96} \right) = 0.04167 = 4.167\%. \quad (19)$$

By definition 1, the latest Eaton UPS module will consume an amount of electrical power that is equal to at least 4.167% of the electrical power consumed by the rack servers when these servers

are operating at peak. In this paper, this minimum amount of electrical power wasted by the UPS when the IT load is at its peak is called the *UPS lower bound*. Hence, in the case of the Eaton UPS, the UPS lower bound is 4.167%. This UPS lower bound imposes a constraint on the PUE of the data center, which means that the PUE will never reach its optimal value of 1 even when the data center is operating at its peak. Consequently, the power wastage will be worse than this UPS lower bound constraint when the data center is not running at its peak.

5.3.2. Server Consolidation and UPS Efficiency

When data center servers are working at their peak, there is no need for server consolidation. However, in an under-utilized data center, i.e., when the servers are operating off their peak, server consolidation is needed to reduce electricity costs. On the one hand, shutting down servers reduces the electrical power load; on the other hand, it also reduces the UPS efficiency.

Property 2: *The server consolidation approach limits the efficiency of distributed double-conversion UPS modules by reducing the current IT power load (P_{IT}) via the shutting down of under-utilized servers.*

PROOF. See Appendix A.3.

6. UPS Consolidation for a Data Center

In practice, there are different power supply configurations of the servers in data centers. Fig. 4(a) presents a standard server power supply configuration with dual power supply units, while in Fig.4(b), a simplified version of Google’s customized power supply server design has been presented. In the case of Google’s customized server configuration [24], the advantages of such configuration include a reduced number of AC/DC conversion phases and a local battery to avoid the need for double-conversion UPSs in the data center. However, it would be enormously costly to replace all existing (traditional) servers with Google’s customized server design depicted in Fig. 4(b), which uses an embedded battery as a server-based UPS [24].

For that reason, most of today data centers use standard server power supply configuration presented in Fig. 4(a). It is nearly impossible to reconstruct the full power delivery paths in existing data centers with distributed UPS systems for the management of power loads.

However, by leveraging emerging server-level micro-ATSS to connect each server to multiple (UPSs) power feeds, we can gain the ability to power a data center with fewer UPS systems. Therefore, in this paper, we suggest an efficient operational design for distributed UPS systems for use in data centers that includes the following steps. In the **first** step, we leverage emerging server-level micro-ATSS to modify the power delivery paths to connect each server to two power supply lines from two adjacent distributed UPS systems, as shown in Fig. 5(a). We activate one of these power lines while keeping the other inactive (on standby) using a server-level ATS. As workloads are redistributed across active servers with the goal of shutting down under-utilized servers, we can change the power delivery path for an active server from one distributed UPS system to another neighboring active UPS system and remove the offloaded UPS system from the power delivery path. Using this approach, we can eliminate inefficient UPS systems while increasing the efficiency of the other systems and simultaneously preserving the number of consolidated servers. By doing this, we may place the active UPS at risk, because the active servers’ peak power may be greater than the UPS’s rated capacity. Therefore, as the **second** step, we need to cap the power consumption of active servers when they are running on a single UPS module

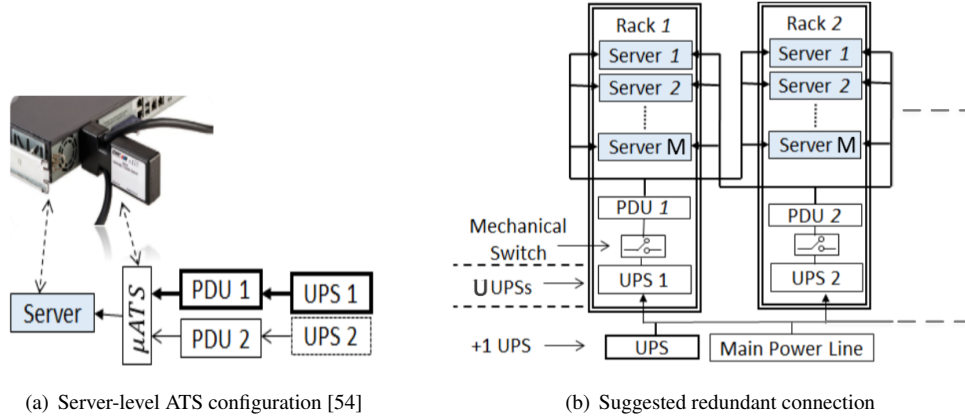


Figure 5: Suggested power delivery path in a data center.

with a lower peak power capacity. For this power capping, we suggest using a UPS-level power-capping manager. For instance, if we switch the power delivery path for an active server from an under-utilized UPS module to another neighboring UPS module, then the number of servers now supplied by the neighboring module or their new peak power consumption may exceed the UPS's power capacity ($UPS_{Rate,u}$). The task of power capping is not within the scope of this paper; for this purpose, we suggest using one of the many power-capping approaches available in the literature[5, 6, 11, 14, 16, 17, 18, 19, 22, 23, 42, 51, 52]. The **third** step relies on the two previous steps. It is managed by an Operational Power Manager at the data-center scale (called UPS Consolidation), which coordinates between the server-level ATSs and the power-capping manager in addition to interfacing and cooperating with any server consolidation systems and power-budgeting managers that are operating at the data-center scale (Fig. 1).

6.1. Server-Level Automatic Transfer Switch (ATS)

A server-level ATS is a small electrical device that has multiple power input lines and a single power output line attached to its corresponding server (Fig. 5(a)). This ATS is used to protect the server from any power failures that may occur in the main power line connected to the server by enabling fall-back to a secondary power line. Commercial versions of such electrical devices are intended to connect a server to two sources of power, namely, the main power line from the utility grid and a backup UPS module; typically, power is supplied only by the power line coming from the utility grid. An example of such a server-level ATS is the μ ATS from Zonit [54]. We suggest the use of such server-level ATSs to connect servers to two UPS modules, such that only one of the UPS modules supplies power during operation while the other remains on standby; see Fig. 5(b). Unlike the 2N UPS configuration (Fig. 4(a)), in which both UPS modules feed electricity to a server simultaneously, this server-level ATS ensures that only one UPS module at a time will feed electricity to a server in an U+1 configuration (Fig. 5(b)) [53].

6.2. Activation Control Procedure for UPS

In order to reduce the consumed power, server consolidation is applied during off peak periods, what results in data center's racks having some servers that are in turn off state. The results of this consolidation process are racks with few active servers. Without disturbing the resulted active

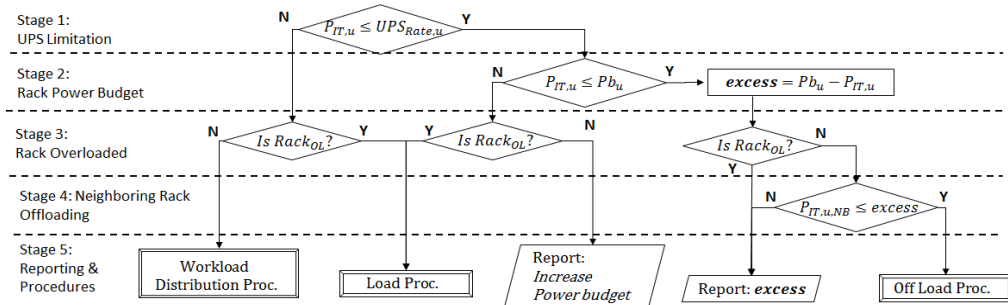


Figure 6: A five-stage control flow for UPS consolidation. The control flow is applied for each UPS in the data center.

servers activity, it is necessary to maintain an appropriate UPS loads, i.e. IT or servers load on UPSs supplying racks. Therefore, if there are two neighboring racks where summation of their loads can be supplied by single UPS, then off-loading of UPS power-supply of one rack to the UPS of the other rack (whichever can handle the combined load) and deactivating of the off-loaded UPS can be done. Intuitively, in case of the opposite process where the combined load is greater than the capacity of the UPS, then the reloading each UPS with the local rack's IT load (servers and network appliances) must take place.

Such off loading process requires a procedure to follow, therefore we proposed a simple yet effective control actions on UPS activation and deactivation (called UPS Consolidation). When putting a UPS into an active state (adding it to the power delivery path) or an inactivate state (removing it from the power delivery path), we need checkpoints and markers, as shown in Fig. 6. These checkpoints and markers form proposed control flow for UPS Consolidation for each UPS system in the data center, and this control flow consists of five stages as follows. In the first stage (called the **UPS Limitation stage** - Fig. 6 Stage 1), the UPS Consolidation control flow checks the power consumed by the active servers (called the IT load and denoted by $P_{IT,u}$) against the rated power capacity of the UPS ($UPS_{Rate,u}$). If there is no violation of the manufacturer specifications and no UPS failure [9], then the control flow moves to the next stage. In the second stage (called the **Rack Power Budget stage** - Fig. 6 Stage 2), the IT power consumed is checked against the power budget (P_bu) allocated as part of the overall data center power budget (such as in [18, 19, 23, 42, 51]). This stage checks whether the allocated power is sufficient to operate the rack and the UPS system. If so, the power surplus (*excess*) is calculated. Then, the third stage will be performed. The third stage (called the **Rack Overloading stage** - Fig. 6 Stage 3) is the most important stage of proposed control flow and checks whether the rack (in particular, the UPS system) has already been overloaded with two sets of active servers (two neighboring racks). Based on the results from the preceding stage and whether the rack is overloaded ($Rack_{OL}$), this stage results in various outcomes. When the preceding stage was the UPS Limitation stage (1st stage), if the rack (UPS) is overloaded, then in the final stage (the **Reporting and Procedure stage** - Fig. 6 Stage 5), the execution of a *Load Procedure* will be invoked; otherwise, a *Workload Distribution Procedure* will be invoked. When the preceding stage was the Rack Power Budget stage (2nd stage), if the rack (UPS) is overloaded, then in the final stage, the execution of the *Load Procedure* will be invoked; otherwise, a report of the power demands will be passed to the (overall) data center power budget manager. When the preceding stage was the calculation of *excess* relative to the power budget, if the rack (UPS) is overloaded, then a report of the *excess* power will be sent to the (overall) data center power budget manager; otherwise, a fourth stage (called the **Neighboring Rack Offloading stage** - Fig. 6 Stage 4) must

Algorithm 1 UPS Consolidation

Input: U

```
1: for ( $u = 1; u \leq U; u++$ ) do
2:   if ( $u$  is odd) then
3:     RETRIEVE the current status of  $P_{IT,u}, P_{IT,u+1}, UPS_{Rate,u}, Rack_{OL,u}$ 
4:      $ACTION = UPS\_Controller(UPS_{Rate,u}, Pb_u, Rack_{OL,u}, P_{IT,u+1})$ 
5:   else
6:     RETRIEVE the current status of  $P_{IT,u}, P_{IT,u-1}, UPS_{Rate,u}, Rack_{OL,u}$ 
7:      $ACTION = UPS\_Controller(UPS_{Rate,u}, Pb_u, Rack_{OL,u}, P_{IT,u-1})$ 
8:   end if
9:   switch ( $ACTION$ )
10:  case 'Workload Distribution':
11:    INVOKE Zhang and Shi Algorithm
12:  case 'Report Excess':
13:    UPDATE  $Pb_u$ 
14:    NOTIFY Data Center Power Budget about the excess
15:  case 'Report Increase Power Budget':
16:    NOTIFY Data Center Power Budget about the shortage
17:  case 'Load':
18:    ACTIVATE the mechanical switch at the local rack  $u$ 
19:  case 'Off Load':
20:    DEACTIVATE the mechanical switch at the local rack  $u$ 
21:  end switch
22: end for
```

be performed. In this 4th stage (Fig. 6 Stage 4), the UPS Consolidation control flow will check the neighboring rack's IT load ($P_{IT,u,NB}$) and compare it against the available *excess* power (Fig. 6 Stage 1). If there is sufficient power available to supply the neighboring rack's active servers, then an *Offloading Procedure* is performed in the final stage, and the UPS of the neighboring rack will be removed from the power delivery path; otherwise, no offloading procedure will be performed, and the *excess* power budget will be reported to the overall manager.

The *Loading Procedure* and the *Offloading Procedure* presented in stage 5 (Fig. 6 Stage 5 bottom) are opposite processes relative to one another. Both procedures basically trigger a mechanical switch attached to the electrical circuit between the UPS system and the PDU. On the one hand, if the *Loading Procedure* is invoked, then the mechanical switch will be switched on and connect the UPS to the PDU. The micro-ATS attached to each server will automatically source power from the local UPS. On the other hand, if the *Offloading Procedure* is invoked, then the mechanical switch will be switched off and disconnect the original UPS supplier from the PDU. As a result, the micro-ATSs will automatically source power from the neighboring UPS system.

Notably, there are three more procedures that are not in the scope of this paper as they are considered an incremental contribution in future works. These procedures are *Workload Distribution Procedure*, *Report: Increase Power Budget* and *Report: excess* (Fig. 6 Stage 5 bottom).

The *Workload Distribution Procedure* is basically inherited from a previous state-of-the-art work by Zhang and Shi [8] where the workload has to be redistributed toward balancing load over racks' server evenly. According to Zhang and Shi [8], when the IT power load applied

Algorithm 2 UPS Controller

Input: $UPS_{Rate,u}$, Pb_u and $Rack_{OL,u}$, $P_{IT,u}$, $P_{IT,u,NB}$ **Output:** *ACTION*

```
1: if ( $P_{IT,u} \leq UPS_{Rate,u}$ ) then
2:   if ( $P_{IT,u} \leq Pb_u$ ) then
3:      $excess = Pb_u - P_{IT,u}$ 
4:     if ( $Rack_{OL,u} = TRUE$ ) then
5:       ACTION : Activate ReportExcess Procedure
6:     else
7:       if ( $P_{IT,u,NB} \leq excess$ ) then
8:         ACTION : Activate OffLoad Procedure
9:       else
10:        ACTION : Activate ReportExcess Procedure
11:      end if
12:    end if
13:  else
14:    if ( $Rack_{OL,u} = TRUE$ ) then
15:      ACTION : Activate Load Procedure
16:    else
17:      ACTION : Activate ReportIncreasePowerBudget Procedure
18:    end if
19:  end if
20: else
21:   if ( $Rack_{OL,u} = TRUE$ ) then
22:     ACTION : Activate Load Procedure
23:   else
24:     ACTION : Activate WorkloadDistribution Procedure (Zhang and Shi)
25:   end if
26: end if
27: return ACTION
```

to the UPS is extremely high (and that UPS is not handling multiple racks as we suggested in this work), a workload redistribution is required so we could improve the all UPS modules efficiencies in the data center.

The *Report: Increase Power Budget* and the *Report: excess* are intuitive procedures that either request power budget increase or reports excess on power budget, respectively (Fig. 6). The first report *Increase Power Budget* can be invoked when the UPS experience an IT power load that is more than the allocated power budget for the rack. Conversely, the second report *Excess* is invoked when the rack's UPS module experience an IT power load that is lesser than the allocated power budget. Those two procedure reports could be potentially used in the trade off on how to "redistribute power budget" among racks (UPSs) to control their power loss. Following this potential trade off, we could gain more power usage effectiveness (PUE) via merging both state-of-the-art contributions namely the workload redistribution by Zhang and Shi [8] and the server consolidation while we stress the existing UPS modules with the proper power loads.

The above conceptual procedure is implemented in data center through two algorithms, namely *UPS Consolidation Algorithm* and *UPS Controller Algorithm*. On one hand, the *UPS Consolida-*

tion Algorithm is periodically checking the UPS load and interacting with both data center power manager (power budget controller) and the mounted mechanical switch. On the other hand, the *UPS Controller Algorithm* is responsible to implement the five-stages control flow checks presented in Fig. 6.

The *UPS Consolidation Algorithm* (Algorithm 1) description is as follows. The algorithm basically requires only the set of U UPSs to re-consolidating the active UPSs in data center. The algorithm is continuously invoked every time t . From step 1 to 22, the algorithm will check each UPS in the data center and apply the proper actions (as explained in the control flowchart - Fig. 6). First, in step 1, it selects the UPS index. And then, in steps 2 to 8, the algorithm is retrieving the needed informations about the local IT power budget (Pb_u), the current local IT power load of the rack ($P_{IT,u}$), the neighboring rack's IT power load ($P_{IT,NB}$), the current UPS power rate or capacity ($UPS_{Rate,u}$) and the status of the UPS overloading ($Rack_{OL,u}$). It is important to note that the neighboring rack's IT power load is related to the UPS index, so if the UPS index is odd, then the neighboring rack's IT power load is the next in index ($u + 1$) otherwise it is the previous in index ($u - 1$). After that, the algorithm invokes the *UPS Controller Algorithm* and wait for the needed actions suggested by that algorithm. From step 9 to 21, and based on the suggested action, the algorithm will apply one of the control flowchart actions mentioned above (Fig. 6).

The *UPS Controller Algorithm* (Algorithm 2) is basically the pseudo-code of the control flowchart presented in Fig. 6. The algorithm requires activation of different procedures in evaluating an UPS on/off activity, which are: *Report Excess Procedure*, *Off Load Procedure*, *Report Excess Procedure*, *Load Procedure*, *Report Increase Power Budget Procedure*, *Load Procedure*, *Workload Distribution Procedure*. From step 1 to step 26, IF-ELSE selection statements are used, which allows to simply recompile the control flowchart (Fig. 6). In step 27, the algorithm would return the suggested action on the current UPS (u).

7. Performance Evaluation

We simulate a daily operation workload, as presented in Fig. 13(a), for a corporate data center with a peak power consumption of 3 MW and a PUE of 1.7. We choose the following subjects for testing: First, considering that the global average PUE for data centers is 1.7, as reported in [49], we choose a representative data center with a PUE of 1.7 as the **Baseline** for our comparison. Further, we recommend readers to revise the report to U.S. Congress on server and data center energy efficiency for more understanding on the status of relevant energy efficiency in data centers [50]. In addition, we evaluate three alternative approaches to power management. The first one is Workload distribution, called the **Workload** approach, as proposed by Zhang and Shi [8]. Workload approach is based on continuous activity of all servers and UPSs independently of the load. The second one is **Server consolidation** or dynamic right-sizing, called the Server consolidation approach, as proposed by Xiao et al. [56] and Lin et al. [5, 6]. Server consolidation approach is based on shutting-down of some servers during off-peak periods and transferring workload to other servers which remain active. Those active servers will continue activity and will be supplied with electricity power by UPSs which supply servers independently of workload. The third approach is in this paper proposed the **Server and UPS Consolidation** approach, which is based on the server consolidation approach. More specifically, we utilize server consolidation as an overall data center management strategy to control the average server utilization in the data center, and we apply proposed UPS Consolidation strategy to optimize the active distributed UPS modules (to reduce UPS power losses) in the data center. In the following subsections, we

Table 2: Testbed configuration scenarios

Parameters	1	2	3	4	5	6
Data Center Type	Homo-geneous	Homo-geneous	Homo-geneous	Homo-geneous	Homo-geneous	Hetero-geneous
UPS Type	Unified	Unified	Unified	Unified	Unified	Mixed
Server Type	Unified	Unified	Unified	Unified	Unified	Mixed
VMs per Server	4	4	5	8	10	4 & 8
UPS Capacity Rate (kW)	1	2	1	1.2	1.2	1 & 1.2
No. of servers (cores) supported by UPS	10 (4)	20 (4)	10 (4)	10 (8)	10 (8)	10 (4) & 10 (8)
No. of Racks	2,000	1,000	2,000	1,000	1,000	1,000
No. of UPSs	2,000	1,000	2,000	1,000	1,000	500 & 500
No. of Servers	20,000	20,000	20,000	10,000	10,000	5,000 & 5,000
No. of VMs	80,000	80,000	100,000	80,000	100,000	20,000 & 40,000

Table 3: UPS and server resources parameters used in analyses

Parameter	Value
α	0.1341
β	0.9408
$Ps_{4iu, idle}(W)$	79.921
$Ps_{8iu, Idle}(W)$	93.58954
$\alpha_{r_{4ju}}$	0.0233
$\alpha_{r_{8ju}}$	0.020126
N	4,8

introduce the experimental testbed (including all considered scenarios), the workloads, and the evaluation metrics.

7.1. Testbed Environment and Scenarios

The testbed environment used for evaluation analysis is based on a simulated data center with tens of thousands of virtual machines (VMs) hosted on thousands of servers located in racks; these racks are equipped with distributed UPS modules. The environment parameters (for both servers and UPSs) used in the simulated data center were generated from real benchmark; where the parameters are listed in Table 3.

N.B. We used the simulation testbed to observe rare cases (incidences) that occur when server consolidation [5, 6, 56] is applied in large-scale data centers. These incidences occur when the **“Server Consolidation”** approach designates under-utilized servers to be shut down and these marked servers happen to be located in the same rack. In this case, the UPS module for this rack will have no load applied to it; eventually, if the UPS is fully charged, it will have no

current passing through it and will have no power loss. Such incidences are very difficult to observe in small-scale testbeds (with 2 to 5 racks); therefore, we evaluated the approaches using a simulation procedure. To observe these incidences, we simulated a data center with 1,000 and 2,000 racks instrumented with UPS modules under different configuration settings. In particular, we simulated a data center with 1,000 racks in the EXP 2, EXP 4, EXP 5 and EXP 6 settings and a data center with 2,000 racks in the EXP 1 and EXP 3 settings, as shown in Table 2.

The simulator used for simulation of tested data center is a Java Program designed to simulate the workload traces per individual server. We used Multi-threads to independently configure both rack servers and UPS modules with the parameters generated from the benchmarks. Despite that CloudSim [57, 58, 59] (a Java-based simulator) could be used, it doesn't focus on the goal of our work in this paper, i.e. the UPS consolidation. Therefore, we preferred to have our own simple yet focused simulator.

We embedded both **Workload** and **Server Consolidation** approaches as modeled codes in the simulator, and we use the server consolidation approach model (code) to be extended to include our **Server and UPS Consolidation** approach. These modeled (coded) approaches were used when we simulated thousands of servers with various simulated workloads (more details in subsection 7.2)

There are several configurations that may be considered for a data center. For example, a data center may be either homogeneous or heterogeneous with respect to the hardware types used, i.e., servers and UPS modules (Table 3). Another configuration setting is related to the characteristics of the workloads handled in the data center, for example, normal or high workload levels and variable or working-hour workload traces. Moreover, we consider for analyses various UPS capacities ($UPS_{Rate,u}$), i.e., 1, 1.2 and 2 kW.

The mathematical models used for the UPS modules and servers are as follows. Parameters and corresponding values defining analyzed UPS system are indicated in Table 3. For the UPS model, we have modified the original UPS efficiency model according to relation (2) to achieve a maximum UPS efficiency of 94%; the modified model is as follows:

$$UPS_{eff,u} = 0.1341 \times \ln\left(\frac{P_{IT,u}}{UPS_{Rate,u}}\right) + 0.9408. \quad (20)$$

Regarding the servers, we consider two types of servers: servers with 4-core CPUs and servers with 8-core CPUs. The corresponding power consumption models according to relation (8) are as follows:

$$P_{S_{4,i,u}} = 79.921 + 0.0233 \times \sum_{j=1}^4 U_{r_{j,i,u}} [W] \quad \forall u \in \mathbf{U}, \forall i = 1, \dots, M \in \mathbb{N}, \quad (21)$$

$$P_{S_{8,i,u}} = 93.58954 + 0.020126 \times \sum_{j=1}^8 U_{r_{j,i,u}} [W] \quad \forall u \in \mathbf{U}, \forall i = 1, \dots, M \in \mathbb{N}, \quad (22)$$

where the constant values in the above two power models are based on real benchmarking experiments. These two power models were generated by benchmarking two real rack-mounted IBM System x3650 servers, one with a 4-core CPU and the other with an 8-core CPU [60, 61]. The hardware specifications of these two servers are as follows: an Intel Xeon E5-2600 processor (4 or 8 cores); 16 GB of memory; 2 TB of hard disk space; 4 Ethernet interfaces, each operating at 1 Gbps; running the Linux CentOS operating system and the Xen hypervisor. In the expressions

above, $U_{r_{ju}}$ denotes the (average) utilization of core j in CPU of server S_{iu} at time t . The benchmarking power models are simplified to consider only the average CPU utilization because the CPU represents the majority of the variability in a server’s power consumption.

7.2. Data Center Workloads and Evaluation Metrics

Our evaluation consists of two phases based on the workloads used: the **quantitative analysis** phase and the **daily operation analysis** phase. We consider the VM demands on a server’s CPU as the workload in our simulations¹, and we define two different workload traces to be used as the VM demands in our evaluations. The first workload trace is based on randomness; which is used for each and every VMs in the **quantitative analysis** phase. Thus, in this evaluation phase, every VM has a random utilization trace at any instance of time; and in this evaluation phase we focused on the number of VMs and the data center capacity; where the size of handled workload is varied over the scenarios (Table 2). The second workload trace, which is a working-hour workload trace, is used for all VMs in the **daily operation analysis** phase. The working-hour workload trace is well-known model suggested by Clazarossa et al. [65]. In addition, there are useful workload traces that could be found on [66]. In this evaluation phase, we focused on data center daily operation analysis rather than data center capacity; which is more likely happening in small, medium and corporate data centers.

For the **quantitative analysis** phase, we evaluate the three approaches on six scenarios, each with different configuration settings, as shown in Table 2. The workloads used in this phase are basically multiple (per VM) randomly generated workloads. For the **daily operation analysis** phase, we use only the first scenario considered in the previous phase (experiment 1 in Table 2) and apply working-hour workloads.

The **evaluation metrics** are defined to quantify the effectiveness of an approach in terms of two main aspects: data center **power reduction** and **PUE** improvement. These two metrics are intuitively important, as all data center owners would like to reduce the operational cost (OPEX) of their data centers and maximize the utilization of that cost to achieve a better Return on Investment (ROI). More precisely, the power reduction is linked to the data center OPEX, and the PUE improvement is related to maximizing the effective utilization of the data center.

8. Results and Discussion

8.1. Quantitative Analysis Phase

The outcomes of simulated scenarios in terms of metrics are summarized in two main figures as follows. The **power reductions** achieved by the three approaches are shown in Fig. 7, and the data center’s **PUEs** under the three approaches are compared against the baseline PUE of 1.7 in Fig. 8.

All approaches achieve power reductions (savings) in all six scenarios, as shown in Fig. 7. The higher value of power savings in Fig. 7 means better data center energy-efficiency and vice versa. However, significant power reductions are observed with the **Server Consolidation** approach and proposed **Server and UPS Consolidation** approach. These two approaches achieve power savings of greater than 40% for a normal workload (EXP 1, 2, and 4 in Table 2)

¹To date, VM profiling and application assignment in data centers have remained a debated issue. A recent study by Vasudevan et al. [62] shows that profiling is feasible and can improve a data center’s energy efficiency, similar to other studies [63, 64]. However, VM profiling is beyond the scope of this paper.




Figure 7: Power reductions where savings are expressed as percentages.




Figure 8: PUEs of the three approaches.

and approximately 30% for a high workload (EXP 3 and 5). These significant savings do not change when the data center includes heterogeneous UPS modules and servers (as in EXP 6); the normal-workload savings remain greater than 40%.

The PUE comparison among the three approaches yields different results, as shown in Fig. 8. The lower value of the PUE in Fig. 8 means better data center energy-efficiency and vice versa. Only the workload distribution method proposed by Zhang and Shi [8] (**Workload**) improves the PUE of the data center, reducing it from the baseline of 1.7 to 1.67. This is a vital improvement compared with the other approaches, as the PUEs for **Server Consolidation** and our proposed **UPS Consolidation** method are degraded compared with the original (baseline) PUE of the data center. **Server Consolidation** achieved the worst PUEs among all approaches, with values of approximately 1.81 PUE for a normal workload (EXP 1, 2, 4 and 6, of which the last represents the “heterogeneous” case) and 1.77 for a high workload (EXP 3 and 5). By comparison, our proposed **Server and UPS Consolidation** method results in less degradation in the PUE, causing the PUE to increase to only 1.72 to 1.75 from the original baseline PUE of 1.7.

We can draw three main observations from our results as follows: 1) the **Workload** approach improves the PUE, 2) the **Server Consolidation** approach and our proposed **Server and UPS Consolidation** approach achieve major power savings, and 3) our proposed **Server and UPS Consolidation** approach exhibits an advantage over **Server Consolidation** alone in terms of both higher power savings and reduced PUE degradation.

The first observation is the improvement in the data center’s PUE achieved using the **workload** distribution approach. This improvement is mainly due to the redistribution of the workloads among the servers to improve the efficiency of the distributed UPS modules. The objective of




Figure 9: Number of active servers i data centers optimized via Server consolidation [56].




Figure 10: Number of UPSs on the power delivery path.

this approach is to efficiently reduce the power losses of the distributed UPS modules, which ultimately results in a significant improvement in the data center’s PUE and marginally contributes to power savings for the data center.

The second observation, namely, the major power savings achieved under **Server Consolidation** and our proposed **Server and UPS Consolidation** approach, can be attributed to the significant number of servers that are shut down under the **Server Consolidation** strategy. This strategy aggressively offloads workloads from under-utilized servers and designates them for possible shutdown; therefore, we highlight the results of its actions, specifically the shutting off of servers, in Fig. 9. Under a **normal** workload, this approach shuts down approximately half of the servers in both a homogeneous data center (EXP 1, 2, and 4 in Table 2) and a heterogeneous data center (EXP 6). Because our proposed **Server and UPS Consolidation** approach also incorporates the server consolidation approach for overall data center management, the proposed approach achieves a relative power reduction similar to that of the **Server Consolidation** approach presented in Fig. 7.

The third observation is the additional power savings and reduced PUE degradation achieved by our **Server and UPS Consolidation** approach compared with the **Server Consolidation** approach. This dual improvement is due to the reduced number of UPS modules on the power delivery path. Figure 10 shows that our **Server and UPS Consolidation** approach significantly reduces the number of active distributed UPS modules. The results indicate that approximately 80% of the active UPS modules are needed to operate a data center with a normal workload regardless of whether the data center has a homogeneous (as in EXP 1, 2, and 4 in Table 2) or heterogeneous (as in EXP 6) configuration. Meanwhile, if the data center has a high workload, our **Server and UPS Consolidation** approach reduces the number of active UPS modules to approximately 88~95%, as seen for EXP 3 and 5 in Fig. 10.




Figure 11: UPS modules eliminated from the delivery path with our proposed UPS Consolidation method.





Figure 12: Active UPS modules and overloaded UPS modules under our UPS Consolidation method.

When the number of UPS modules deactivated by our approach is considered, as shown in Fig. 11, our **Server and UPS Consolidation** approach is able to eliminate approximately 760 of the 2,000 distributed UPS modules (approximately 62% of the active UPS modules) in the case of EXP 1. This significant result is achieved for a data center with a homogeneous configuration, unified small-scale servers (servers with 4-core CPUs), and a normal workload. By contrast, for a data center with a high workload, a homogeneous configuration and either large-scale servers of a unified type or a mix of server types, our **Server and UPS Consolidation** approach reduces the number of UPS modules only slightly to avoid the risk of power surges under high workload demands (EXP 5 and 6 in Fig. 11).


Based on our UPS Consolidation control flow process, there is one possible case in which the IT power load of one rack (UPS) will be offloaded to a neighboring rack (UPS), namely, when the sum of the IT power loads of both racks is within the rated capacity of the UPS (UPS_{Rate}). In this case, our method will offload (switch the power delivery path for) each individual active server on the rack (using the server-level ATSS) and mark the loaded rack (specifically, the loaded UPS) as “OverLoaded” or “OL”. In Fig. 12, we show the numbers of active and overloaded UPS modules that result from the application of our **Server and UPS Consolidation** approach, denoted by **Active UPS** and **OverLoaded**, respectively, as well as the original (baseline) number of UPS modules for comparison. As shown in Fig. 12, there are a considerable number of overloaded UPS modules that are capable of supplying power to two racks simultaneously.

8.2. Daily Operation Analysis Phase


The outcomes of this phase are presented in Fig. 13 as follows. First, the power reduction results for the traffic pattern presented in Fig. 13(a) are shown in Fig. 13(b). With conventional




(a) Workload (data center utilization)



(b) Data center power consumption



(c) UPS power loss



(d) PUE

Figure 13: A one-day simulation of a data center using the EXP 1 scenario configuration and a two-peak daily workload trace.

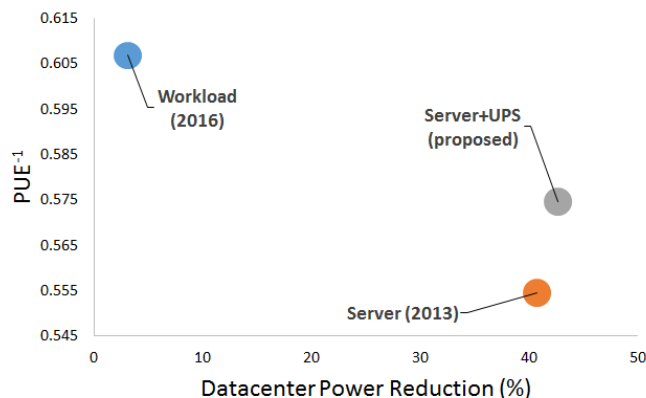


Figure 14: Overall comparison, in which the top-right corner corresponds to the optimal performance.

data center design (Baseline), the power consumption of the data center is enormous; by contrast, both **Server Consolidation** and our **Server and UPS Consolidation** approach result in a significant reduction in power consumption during off-peak operation.

Second, the PUE results are presented in Fig. 13(d). When the **Server Consolidation** approach is applied, the PUE of the data center is dramatically degraded compared with conventional data center design (increased from the original PUE of 1.7 to 1.83). This is because **Server Consolidation** reduces the energy consumption of storage/computing equipment (e.g., servers) without impacting other facility equipment (e.g., UPSs, lighting, cooling). This is a drawback in terms of the PUE metric; approximately 13% of the PUE increase is a consequence of the optimization of server power consumption. By comparison, our **Server and UPS Consolidation** approach reduces the PUE degradation caused by **Server Consolidation**, restricting the data center’s PUE (during off-peak times) to below 1.73 (i.e., an additional PUE waste of less than 3%).

Finally, the UPS power loss is addressed in Fig. 13(c). According to this figure, when the **Server Consolidation** approach is applied, the UPS power loss in some periods of a day increases to almost 20%. However, when our **Server and UPS Consolidation** approach is applied, the UPS power loss decreases by between 20 and 40% of the original UPS power loss. This is because the proposed **UPS Consolidation** model optimizes the number of active UPSs in accordance with the server load variations in the data center.

Conclusions regarding the three investigated approaches are summarized in Fig. 14. We plot the performance of each approach in terms of the **power reduction** and **PUE** evaluation metrics, where the inverse PUE^{-1} is used to highlight the fact that a lower PUE is better. Consequently, the top-right corner corresponds to the optimal performance, and as seen in the figure, there is no approach that achieves that result. Nevertheless, each approach achieves a good position relative to its scope; for example, the **workload** distribution approach (proposed by Zhang and Shi [8]) achieves the best PUE among all approaches. However, our **Server and UPS Consolidation** approach achieves the closest-to-optimal performance in terms of both metrics. This achievement can be realized in typical small-, medium-, and corporate-scale data centers with insufficient energy awareness.

A final remark on our approach is that the UPS Consolidation strategy proposed here is more a heuristic strategy than an optimization strategy for the power budget of a data center, and it simultaneously helps to preserve the data center’s PUE. However, if a data center owner is

interested in both preserving the PUE and controlling the power budget, then a good choice is to enable coordination between (1) the workload distribution, server consolidation, and UPS consolidation approaches on the one hand and (2) the power-capping and workload-balancing approaches for geo-distributed data centers on the other hand.

9. Conclusion

Data centers are key back-end communication entities in ICT, and reducing the amount of power they waste would greatly contribute to achieving efficient ICT operations. Adopting the server consolidation approach in small-, medium-, and corporate-scale data centers can greatly reduce their power consumption. However, the implementation of server consolidation in a data center with distributed UPS modules will severely degrade the data center's PUE. As illustrated in this paper, server consolidation reduces the number of active servers (the major IT power consumers); however, it causes the IT power load to become unbalanced over the distributed UPSs. This unbalancing of the IT power load limits the efficiency of the UPSs (non-IT power consumers) and eventually increases their power loss.

Therefore, in this paper, we have proposed the UPS Consolidation strategy to cope with the degradation in a data center's PUE caused by server consolidation. Our UPS Consolidation strategy is an adaptive process that continuously looks for opportunities to reduce the number of UPSs by offloading the workload of an under-loaded UPS to a neighboring UPS whenever that neighboring UPS can handle the additional load. This offloading between UPSs can be achieved by means of micro-ATSs at the server end to enable failover to a neighboring UPS.

The key features of the proposed UPS Consolidation strategy are summarized as follows. First, in combination with server consolidation, our UPS Consolidation strategy achieves higher power savings compared with the workload distribution approach and the server consolidation approach alone. Second, under our approach, the data center's racks are operated with only 80% of the total UPS modules, whereas the other approaches do not focus on reducing the number of active UPS modules. Finally, the UPS Consolidation strategy is promising in terms of achieving significant cost savings during data center operation. Our preliminary evaluation shows that applying only a state-of-the-art technique such as server consolidation in a conventional data center will increase the power waste due to UPS power loss by up to 13%. Our approach is able to reduce the data center's UPS power loss by 20 to 40% under off-peak workloads while preserving the advantages of previous approaches, i.e., the advantages of server consolidation.

In future works, we will deploy our UPS Consolidation strategy in small and mid-sized data centers to evaluate the process of transitioning between UPS modules and the deployment of power-capping techniques. Further extensions of our UPS Consolidation strategy could include coordination with thermal control systems such as those based on hot spots and free cooling in data centers.

- [1] Koomey, "Growth in data center electricity use 2005 to 2010," *A report by Analytical Press, completed at the request of The New York Times*, p. 9, 2011.
- [2] J. Whitney and P. Delforge, "Data center efficiency assessment—scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers," *NRDC and Anthesis, Rep. IP*, 2014.
- [3] M. Stansberry, "The seventh annual uptime institute data center industry survey," May 2017.
- [4] M. Smolaks, "Eaton launches compact rackmount version of 5p ups," *Datacenter Dynamics*, January 2017.
- [5] M. Lin, A. Wierman, L. L. Andrew, and E. Thereska, "Dynamic right-sizing for power-proportional data centers," in *Proceedings IEEE INFOCOM*, 2011.
- [6] M. Lin, A. Wierman, L. L. Andrew, and E. Thereska, "Dynamic right-sizing for power-proportional data centers," *IEEE/ACM Transactions on Networking (TON)*, vol. 21, no. 5, pp. 1378–1391, 2013.

- [7] Q. Zhang and W. Shi, "Ups-aware workload placement in enterprise data centers," *IEEE Computer Magazine*, 2015.
- [8] Q. Zhang and W. Shi, "Energy-efficient workload placement in enterprise datacenters," *Computer*, vol. 49, no. 2, pp. 46–52, 2016.
- [9] E. N. Power, "Study says unplanned data center outages cost companies nearly \$9,000 per minute," *BUILDINGS.com*, January 2016, columbus, Ohio.
- [10] K. Heslin, "Here's why major airline outages keep happening, and what you can learn from them," *ITProPortal*, June 2017. [Online]. Available: <https://www.itproportal.com/features/heres-why-major-airline-outages-keep-happening-and-what-you-can-learn-from-them/>
- [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 *Computer Architecture (ISCA), 2012 39th Annual International Symposium on*. IEEE, 2012, pp. 488–499.
- [12] C. Li, Y. Hu, J. Gu, J. Yuan, and T. Li, "Oasis: Scaling out datacenter sustainably and economically," *IEEE TPDS*, 2016.
- [13] L. Liu, H. Sun, C. Li, Y. Hu, J. Xin, N. Zheng, and T. Li, "Re-ups: an adaptive distributed energy storage system for dynamically managing solar energy in green datacenters," *the Journal of Supercomputing*, vol. 72, no. 1, pp. 295–316, 2016.
- [14] M. A. Islam, X. Ren, S. Ren, A. Wierman, and X. Wang, "A market approach for handling power emergencies in multi-tenant data center," in *2016 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2016, pp. 432–443.
- [15] Patterson M.K. et al., "TUE, a New Energy-Efficiency Metric Applied at ORNL's Jaguar," in *International Supercomputing Conference, ISC 2013*, pp:372-382, 2013
- [16] L. Liu, C. Li, H. Sun, Y. Hu, J. Gu, T. Li, J. Xin, and N. Zheng, "Heb: deploying and managing hybrid energy buffers for improving datacenter efficiency and economy," in *ACM SIGARCH Computer Architecture News*, vol. 43, no. 3. ACM, 2015, pp. 463–475.
- [17] W. Zheng, K. Ma, and X. Wang, "Hybrid energy storage with supercapacitor for cost-efficient data center power shaving and capping," *IEEE Transactions on Parallel and Distributed Systems*, vol. 28, no. 4, pp. 1105–1118, 2017.
- [18] L. et al., "Sizecap: Efficiently handling power surges in fuel cell powered data centers," in *HPCA'16*. IEEE, 2016, pp. 444–456.
- [19] T. et al., "Nipd: Non-intrusive power disaggregation in legacy datacenters," *IEEE Transactions on Computers*, vol. 66, no. 2, pp. 312–325, Feb 2017.
- [20] L Gu, D Zeng, A Barnawi, S Guo, I Stojmenovic, "Optimal task placement with QoS constraints in geo-distributed data centers using DVFS," *IEEE Transactions on Computers*, vol. 64, no. 7, pp. 2049–2059, (2015)
- [21] L Gu, D Zeng, S Guo, Y Xiang, J Hu, "A general communication cost optimization framework for big data stream processing in geo-distributed data centers," *IEEE Transactions on Computers*, vol. 65, no. 1, pp. 19–29, (2016).
- [22] P. Petoumenos, L. Mukhanov, Z. Wang, H. Leather, and D. S. Nikolopoulos, "Power capping: what works, what does not," in *ICPADS'15*. IEEE, 2015, pp. 525–534.
- [23] W. et al., "Dynamo: facebook's data center-wide power management system," in *ISCA'16*. IEEE, 2016, pp. 469–480.
- [24] R. Miller, "Google files patent on ups architecture," February 2008.[Online]. Available: <http://www.datacenterknowledge.com/archives/2008/02/15/google-files-patent-on-ups-architecture>
- [25] C. Li, Z. Wang, X. Hou, H. Chen, X. Liang, and M. Guo, "Power attack defense: Securing battery-backed data centers," in *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*, June 2016, pp. 493–505.
- [26] B. Aksanli, E. Pettis, and T. Rosing, "Architecting efficient peak power shaving using batteries in data centers," in *2013 IEEE 21st International Symposium on Modelling, Analysis and Simulation of Computer and Telecommunication Systems*. IEEE, 2013, pp. 242–253.
- [27] B. Aksanli, T. Rosing, and E. Pettis, "Distributed battery control for peak power shaving in datacenters," in *Green Computing Conference (IGCC), 2013 International*. IEEE, 2013, pp. 1–8.
- [28] N. Deng, C. Stewart, and J. Li, "Concentrating renewable energy in grid-tied datacenters," in *Proceedings of the 2011 IEEE International Symposium on Sustainable Systems and Technology*. IEEE, 2011, pp. 1–6.
- [29] Í. Goiri, K. Le, M. E. Haque, R. Beauchea, T. D. Nguyen, J. Guitart, J. Torres, and R. Bianchini, "Greenslot: scheduling energy consumption in green datacenters," in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*. ACM, 2011, p. 20.
- [30] Í. Goiri, K. Le, T. D. Nguyen, J. Guitart, J. Torres, and R. Bianchini, "Greenhadoop: leveraging green energy in data-processing frameworks," in *Proceedings of the 7th ACM european conference on Computer Systems*. ACM, 2012, pp. 57–70.
- [31] Í. Goiri, W. Katsak, K. Le, T. D. Nguyen, and R. Bianchini, "Parasol and greenswitch: managing datacenters powered by renewable energy," in *ACM SIGARCH Computer Architecture News*, vol. 41, no. 1. ACM, 2013, pp.

- 51–64.
- [32] Í. Goiri, W. Katsak, K. Le, T. D. Nguyen, and R. Bianchini, “Designing and managing datacenters powered by renewable energy,” *IEEE Micro*, p. 1, 2014.
 - [33] L. Liu, C. Li, H. Sun, Y. Hu, J. Xin, N. Zheng, and T. Li, “Leveraging heterogeneous power for improving datacenter efficiency and resiliency,” *IEEE Computer Architecture Letters*, vol. 14, no. 1, pp. 41–45, Jan 2015.
 - [34] C. Li, R. Wang, T. Li, D. Qian, and J. Yuan, “Managing green datacenters powered by hybrid renewable energy systems,” in *11th International Conference on Autonomic Computing (ICAC 14)*. Philadelphia, PA: USENIX Association, Jun. 2014, pp. 261–272.
 - [35] C. Li, R. Wang, D. Qian, and T. Li, “Managing server clusters on renewable energy mix,” *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, vol. 11, no. 1, p. 1, 2016.
 - [36] L. Liu, H. Sun, C. Li, Y. Hu, N. Zheng, and T. Li, “Towards an adaptive multi-power-source datacenter,” in *Proceedings of the 2016 International Conference on Supercomputing*. ACM, 2016, p. 11.
 - [37] X. Li, R. Wang, Z. Luan, Y. Liu, and D. Qian, “Coordinating workload balancing and power switching in renewable energy powered data center,” *Frontiers of Computer Science*, pp. 1–14, 2016.
 - [38] N. Sharma, S. Barker, D. Irwin, and P. Shenoy, “Blink: managing server clusters on intermittent power,” in *ACM SIGPLAN Notices*, vol. 46, no. 3. ACM, 2011, pp. 185–198.
 - [39] C. Li, W. Zhang, C.-B. Cho, and T. Li, “Solarcore: Solar energy driven multi-core architecture power management,” in *2011 IEEE 17th International Symposium on High Performance Computer Architecture*. IEEE, 2011, pp. 205–216.
 - [40] C. Li, A. Qouneh, and T. Li, “iswitch: coordinating and optimizing renewable energy powered server clusters,” in *Computer Architecture (ISCA), 2012 39th Annual International Symposium on*. IEEE, 2012, pp. 512–523.
 - [41] W. et al., “Underprovisioning backup power infrastructure for datacenters,” in *ACM SIGARCH’14*, vol. 42, no. 1. ACM, 2014, pp. 177–192.
 - [42] W. et al., “Underprovisioning backup power infrastructure for datacenters,” in *ACM SIGARCH’14*, vol. 42, no. 1. ACM, 2014, pp. 177–192.
 - [43] L. Zhao, J. Brouwer, J. Liu, S. James, J. Siegler, A. Kansal, and E. Peterson, “Fuel cells for data centers: Power generation inches from the server,” Microsoft Research, Tech. Rep. MSRTR-2014-37, Tech. Rep., 2014.
 - [44] S. Sevenscan, G. Lindbergh, C. Lagergren, and P. Alvfors, “Economic feasibility study of a fuel cell-based combined cooling, heating and power system for a data centre,” *Energy and Buildings*, vol. 111, pp. 218–223, 2016.
 - [45] M. A. Islam and S. Ren, “A new perspective on energy accounting in multi-tenant data centers,” in *USENIX Workshop on Cool Topics on Sustainable Data Centers (CoolDC 16)*, 2016.
 - [46] “Yoctopuce,” <http://www.yoctopuce.com/>, accessed June 25, 2015.
 - [47] Ljung, Lennart, “System Identification Toolbox for Use with MATLAB,” *The MathWorks, Inc.*, 2007.
 - [48] Mingay, Simo, “Green IT: the new industry shock wave,” *Gartner RAS Research Note G*, vol. 153703, 2007.
 - [49] “2014 Uptime Institute Report,” <https://journal.uptimeinstitute.com/2014-data-center-industry-survey/>.
 - [50] “Report to Congress on Server and Data Center Energy Efficiency: Public Law 109-431,” available in https://eta.lbl.gov/sites/default/files/publications/pdf_3.pdf
 - [51] W. et al., “Ship: A scalable hierarchical power control architecture for large-scale data centers,” *IEEE TPDS*, vol. 23, no. 1, pp. 168–176, 2012.
 - [52] X. Shi, C. A. Briere, S. M. Djouadi, Y. Wang, and Y. Feng, “Power-aware performance management of virtualized enterprise servers via robust adaptive control,” *Cluster Computing*, vol. 18, no. 1, pp. 419–433, 2015.
 - [53] “Data center site infrastructure tier standard: Topology,” Uptime Institute, LLC, 2010, accessed on February 23, 2017. [Online]. Available: www.gpxglobal.net/wp-content/uploads/2012/08/tierstandardtopology.pdf
 - [54] “mats - micro automatic transfer switch,” Zonit Structured Solutions, Jan. 2017, accessed on January 3, 2017. [Online]. Available: <http://www.zonit.com/products/micro-ats>
 - [55] W. Torell, “Pue relies on continuous monitoring,” April 2017. [Online]. Available: <http://www.datacenterdynamics.com/content-tracks/power-cooling/pue-relies-on-continuous-monitoring/98214.fullarticle>
 - [56] Z. Xiao, W. Song, and Q. Chen, “Dynamic resource allocation using virtual machines for cloud computing environment,” *IEEE TPDS*, vol. 24, no. 6, pp. 1107–1117, 2013.
 - [57] Buyya, Rajkumar, Rajiv Ranjan, and Rodrigo N. Calheiros, “Modeling and simulation of scalable Cloud computing environments and the CloudSim toolkit: Challenges and opportunities,” *IEEE International Conference on High Performance Computing & Simulation, (HPCS’09)*, 2009.
 - [58] Calheiros, R. N., Ranjan, R., De Rose, C. A., and Buyya, R. “Cloudsim: A novel framework for modeling and simulation of cloud computing infrastructures and services,” *arXiv preprint*, arXiv:0903.2525, (2009).
 - [59] Calheiros, R. N., Ranjan, R., Beloglazov, A., De Rose, C. A., and Buyya, R., “CloudSim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms,” *Software:*

- Practice and experience*, 41(1), 23-50, 2011.
- [60] F. Al-Hazemi, Y. Peng, and C.-H. Youn, "A miso model for power consumption in virtualized servers," *Cluster Computing*, vol. 18, no. 2, pp. 847–863, 2015.
 - [61] F. AL-Hazemi, D.-K. Kang, S.-H. Kim, Y. Peng, S. H. S. Newaz, and C.-H. Youn, "LpcFreqSchd: A local power controller using the frequency scheduling approach for virtualized servers," *Cluster Computing*, pp. 1–16, 2016.
 - [62] M. Vasudevan, Y.-C. Tian, M. Tang, and E. Kozan, "Profile-based application assignment for greener and more energy-efficient data centers," *Future Generation Computer Systems*, vol. 67, pp. 94 – 108, 2017.
 - [63] A. Kansal, J. Liu, D. C. Burger, and A. A. Bhattacharya, "Virtual machine power consumption measurement and management," Oct. 14 2014, uS Patent 8,862,914.
 - [64] B. Krishnan, H. Amur, A. Gavrilovska, and K. Schwan, "Vm power metering: feasibility and challenges," *ACM SIGMETRICS Performance Evaluation Review*, vol. 38, no. 3, pp. 56–60, 2011.
 - [65] Calzarossa, Maria, and Giuseppe Serazzi. "A characterization of the variation in time of workload arrival patterns." *IEEE Transactions on Computers* 2 (1985): 156-162.
 - [66] Beloglazov Anton, and Rajkumar Buyya. "Optimal online deterministic algorithms and adaptive heuristics for energy and performance efficient dynamic consolidation of virtual machines in cloud data centers." *Concurrency and Computation: Practice and Experience*, 24.13 (2012): pp. 1397-1420.

Appendix A.

In this appendix, we prove property 1, lemma 1 and property 2 as follows.

Appendix A.1. Proof of Property 1

PROOF. This proof is based on the concept of the *mathematical limit* $\{\lim_{x \rightarrow \infty} f(x)\}$. According to the UPS efficiency model given in Eq. (2), $UPS_{eff,u}$ is maximized when the IT power load ($P_{IT,u}$) is equal to the rated power capacity of the UPS ($UPS_{Rate,u}$). To apply the *mathematical limit* concept to infer the maximum possible efficiency of the UPS when the IT power load ($P_{IT,u}$) reaches the rated capacity ($UPS_{Rate,u}$), we consider the following function $f(x)$:

$$f(x) = UPS_{eff,u}(P_{IT,u}) = \alpha_u \times \ln\left(\frac{P_{IT,u}}{UPS_{Rate,u}}\right) + \beta_u. \quad (A.1)$$

$$\lim_{x \rightarrow 1} f(x) = \lim_{P_{IT,u} \rightarrow UPS_{Rate,u}} UPS_{eff,u}(P_{IT,u}) \quad (A.2)$$

$$= \lim_{P_{IT,u} \rightarrow UPS_{Rate,u}} \left[\alpha_u \ln\left(\frac{P_{IT,u}}{UPS_{Rate,u}}\right) + \beta_u \right] \quad (A.3)$$

$$= \left[\alpha_u \times \ln\left(\frac{UPS_{Rate,u}}{UPS_{Rate,u}}\right) + \beta_u \right] = \beta_u, \quad (A.4)$$

where β_u is the estimated parameter that represents the maximum efficiency level ($UPS_{eff,u,MAX}$) that a UPS could achieve, that is,

$$UPS_{eff,u,MAX} = \beta_u. \quad (A.5)$$

By substituting the maximum achievable UPS efficiency into the ratio expressed in Eq. (16), we can obtain the smallest (minimum) ratio between the UPS power loss and the corresponding IT power load as follows:

$$RA_{ups,u,MIN} = \left(\frac{1 - UPS_{eff,u,MAX}}{UPS_{eff,u,MAX}} \right), \quad (A.6)$$

where

$$0 < UPS_{eff,u,MAX} < 1. \quad (A.7)$$

Appendix A.2. Proof of Lemma 1

PROOF. First, we know that the power loss of a double-conversion UPS is given by Eq. (A.8):

$$P_{UPS,u} = RA_{ups,u} \times P_{IT,u}. \quad (A.8)$$

We are interested in determining the minimum amount of electrical power that a double-conversion UPS module could consume ($P_{UPS,u,MIN}$) due to the conversion of electrical power from AC to DC and from DC to AC.

According to property 1, the minimum bound can be achieved when the ratio $RA_{ups,u}$ is minimized. In addition, the ratio $RA_{ups,u,MIN}$ is achievable if and only if the IT power load ($P_{IT,u}$) is equal to the maximum UPS rated capacity ($UPS_{Rate,u}$). Thus,

$$P_{UPS,u,MIN} \longleftrightarrow RA_{ups,u} = RA_{ups,u,MIN}, \quad (A.9)$$

$$RA_{ups,u,MIN} \longleftrightarrow P_{IT,u} = UPS_{Rate,u}. \quad (A.10)$$

Therefore, by substituting Eq. (A.6) into Eq. (A.8) and substituting the UPS rated power capacity $UPS_{Rate,u}$ for the IT power load $P_{IT,u}$, we can obtain the minimum UPS power loss as follows:

$$P_{UPS,u,MIN} = \left(\frac{1 - UPS_{eff,u,MAX}}{UPS_{eff,u,MAX}} \right) \times UPS_{Rate,u}, \quad (A.11)$$

where

$$\left(\frac{1 - UPS_{eff,u,MAX}}{UPS_{eff,u,MAX}} \right) > 0 \quad (A.12)$$

and

$$UPS_{Rate,u} > 0. \quad (A.13)$$

As a result,

$$P_{UPS,u,MIN} > 0. \quad (A.14)$$

Appendix A.3. Proof of Property 2

PROOF. Suppose that we have a distributed UPS module that has a maximum rated power capacity ($UPS_{Rate,u}$) that is sufficient to provide power up to that required for all of a rack's servers during peak operation. The number of servers in the rack is M , and the rack's peak power ($\overline{P_{IT,u}}$) is the sum of the peak powers for all servers ($\overline{P_{S_{iu}}}$). Therefore, the ideal (maximum) power consumption of the rack is as follows:

$$UPS_{Rate,u} = \overline{P_{IT,u}}, \quad (A.15)$$

where

$$\overline{P_{IT,u}} = \overline{P_{S_u}} = \sum_{i=1}^M \overline{P_{S_{iu}}} = UPS_{Rate,u} \quad [W] \quad \forall u \in \mathbf{U}. \quad (A.16)$$

Starting from Eq. (8), we revise the server power consumption expression as follows:

$$\overline{P_{S_{iu}}} = P_{S_{iu},idle} + \sum_{j=1}^M (\alpha_{r_{ju}} \times \overline{U_{r_{ju}}}) \quad [W] \quad \forall u \in \mathbf{U}, \forall i = 1, \dots, M \in \mathbb{N}, \quad (A.17)$$

where $\overline{U_{r_{ju}}}$ represents 100% utilization of resource j of server s_{iu} , $\alpha_{r_{ju}}$ is the estimated regression parameter for resource j , and $P_{S_{iu},idle}$ is the power consumed in the idle state.

Because servers are being shut down and the number of active servers is consequently fewer than the actual number of servers (M) in the rack, we define $m < M$ to denote the number of shutdown servers. Thus, the current power consumption of the rack (after server consolidation) will be as follows:

$$\overline{P_{IT,u}} = \sum_{i=1}^M \overline{P_{S_{iu}}} - \sum_{i=1}^{m < M} P_{S_{iu}} \quad [W] \quad \forall u \in \mathbf{U}, \forall i = 1, \dots, M \in \mathbb{N}, \quad (A.18)$$

where the server power $P_{S_{iu}}$ is as defined in the original Eq. (8). By applying the IT power load in Eq. (A.18) in the UPS efficiency model given in Eq. (2), we obtain the updated UPS efficiency ($UPS_{eff,u}^*$) as follows:

$$UPS_{eff,u}^* = \alpha_u \times \ln \left(\frac{\sum_{i=1}^M \overline{P_{S_{iu}}} - \sum_{i=1}^{m < M} P_{S_{iu}}}{UPS_{Rate,u}} \right) + \beta_u \quad \forall u \in \mathbf{U}, \forall i = 1, \dots, M \in \mathbb{N}. \quad (A.19)$$

Because α_u and β_u are fixed parameters, we can evaluate the argument of the natural logarithm as a comparator. Thus, we compare the cases of the rack with the ideal (peak) IT power load and the rack with consolidated servers (off-peak) as follows:

$$\frac{\sum_{i=1}^M \overline{P_{Siu}} - \sum_{i=1}^{m < M} P_{Siu}}{UPS_{Rate,u}} = \frac{\sum_{i=1}^M \overline{P_{Siu}}}{UPS_{Rate,u}}, \quad (\text{A.20})$$

where the left-hand side of the equality corresponds to the IT power load of the consolidated servers and the right-hand side corresponds to the ideal (peak) IT power load. From Eq. (A.15), we have

$$\frac{\sum_{i=1}^M \overline{P_{Siu}}}{UPS_{Rate,u}} = \frac{UPS_{Rate,u}}{UPS_{Rate,u}} = 1. \quad (\text{A.21})$$

Therefore,

$$\frac{\sum_{i=1}^M \overline{P_{Siu}} - \sum_{i=1}^{m < M} P_{Siu}}{UPS_{Rate,u}} < 1, \quad (\text{A.22})$$

$$\frac{\sum_{i=1}^M \overline{P_{Siu}}}{UPS_{Rate,u}} - \frac{\sum_{i=1}^{m < M} P_{Siu}}{UPS_{Rate,u}} < 1, \quad (\text{A.23})$$

$$1 - \frac{\sum_{i=1}^{m < M} P_{Siu}}{UPS_{Rate,u}} < 1. \quad (\text{A.24})$$

The above comparison shows us that the quantity inside the natural logarithm in the UPS efficiency model in Eq. (A.19) is less than 1, which causes the value of the natural logarithm to be strictly negative. According to Eq. (A.5), when the rack is operating under the ideal (peak) IT power load, the UPS efficiency is $UPS_{eff,u} = \beta_u$. By contrast, under the consolidated server (off-peak) IT power load,

$$UPS_{eff,u}^* = \alpha_u \times \ln \left(1 - \frac{\sum_{i=1}^{m < M} P_{Siu}}{UPS_{Rate,u}} \right) + \beta_u \quad \forall u \in \mathbf{U}, \forall i = 1, \dots, M \in \mathbb{N}. \quad (\text{A.25})$$

Therefore,

$$UPS_{eff,u}^* \leq UPS_{eff,u}. \quad (\text{A.26})$$