

# A Survey on Power Management Techniques for Oversubscription of Multi-Tenant Data Centers

SULAV MALLA, University of South Florida

KEN CHRISTENSEN, University of South Florida

---

Power management for data centers has been extensively studied in the past ten years. Most research has focused on owner-operated data centers with less focus on Multi-Tenant Data Centers (MTDC) or colocation data centers. In an MTDC, an operator owns the building and leases out space, power, and cooling to tenants to install their own IT equipment. MTDC's present new challenges for data center power management due to an inherent lack of coordination between the operator and tenants. In this paper, we conduct a comprehensive survey of existing MTDC power management techniques for demand response programs, sustainability, and/or power hierarchy oversubscription. Power oversubscription is of particular interest as it can maximize resource utilization, increase operator profit, and reduce tenant costs. We create a taxonomy to classify and compare key works. Our taxonomy and review differ from existing works in that our emphasis is on safe power oversubscription, which has been neglected in previous surveys. We propose future research for prediction and control of power overload events in an oversubscribed MTDC.

CCS Concepts: • **General and reference** → **Surveys and overviews**; • **Hardware** → *Power and energy*; • **Computer systems organization** → *Availability*;

Additional Key Words and Phrases: Multi-tenant data center, colocation data center, power management, oversubscription

---

## 1 INTRODUCTION

The proliferation of smart consumer devices and Internet of Things (IoT) has placed an increasing demand on data processing, storage, and management in the cloud. Increasingly, companies are migrating towards cloud based services [Columbus 2017] as such services are easy to manage and highly accessible to clients. Moreover, popular Internet services, such as search, email, movie/music streaming, etc. inherently require a cluster of interconnected computing and storage nodes rather than a single machine. This rise in server-side computation has driven the need for more data centers, the physical home of the cloud, both in terms of number as well as size.

Data centers can range from small server closet and rooms to large buildings consuming tens of megawatts of power (equivalent to that of a small city). In 2014, electricity consumption by data centers accounted for approximately 1.8% (70 TWh) of the entire electricity used in the U.S. [Shehabi et al. 2016]. Of the total energy consumed by data centers, on an average, only 60% of it is used to power the IT equipment [Shehabi et al. 2016; Uptime Institute 2014] including servers, storage, and networking equipment while the rest, two fifths of the energy, is wasted in mechanical (cooling) and electrical (lighting) systems of the data center. Given such a large energy footprint and massive energy waste, adoption of energy efficiency techniques in data centers is necessary. Consequently, significant research has been done in the past decade to tackle this problem, which are described in many past surveys [Beloglazov et al. 2011; Kong and Liu 2014; Liu and Zhu 2010; Rahman et al. 2014; Shuja et al. 2016].

However, the focus of past research has mostly been towards owner-operated data centers where a single entity owns and manages the data center infrastructure as well as the IT equipment in them. Equally important, but far less explored, is the multi-tenant data center (MTDC), also referred to as colocation data center or simply as “colos.” In an MTDC, the data center infrastructure is

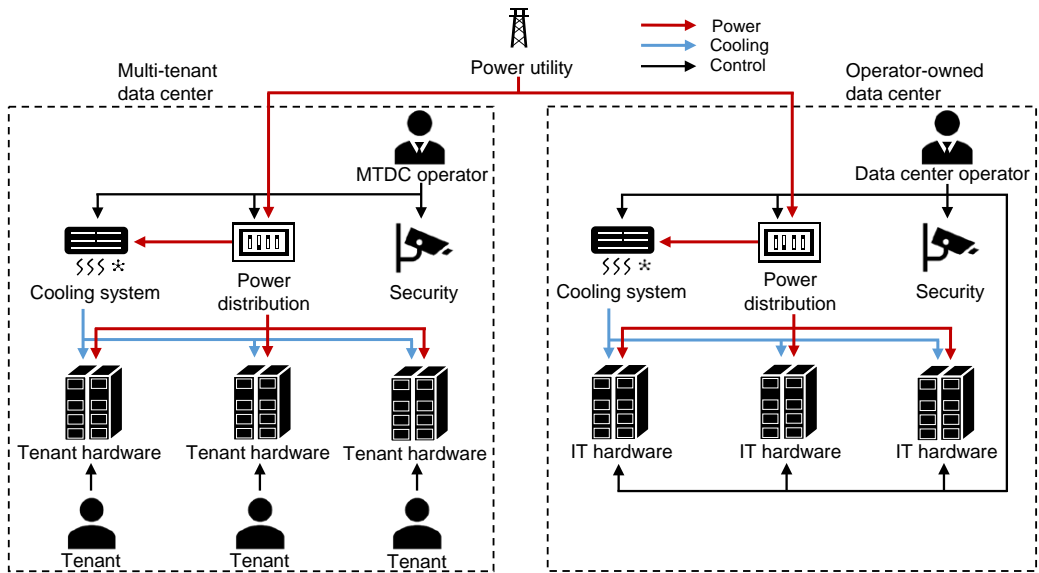



Fig. 1. Difference between multi-tenant data center (where the operator owns the data center infrastructure while tenants own the IT equipment) and operator-owned data center (where the operator has central control over all components).


owned by an operator who leases out space and power capacity in the form of racks, cages, or rooms to tenants. The operator provides power, cooling, and security while tenants install and maintain their own IT equipment to provide services to the tenant's customers, as illustrated in Fig. 1. Unlike in operator-owned data center where the operator has control over all aspect of the data center, in an MTDC the operator does not control the IT hardware that belong to tenants. Many businesses opt for MTDC as building their own data center is expensive and can take years to complete, they might want to place their servers closer to the client locations, move servers from inefficient office rooms to MTDC with higher power efficiency to minimize carbon footprint, and/or expand globally to multiple locations around the world. MTDC offers a quick and scalable solution to businesses at relatively small capital expenditure.

Major websites, e-commerce, banks, and even IT giants like Google, Microsoft, IBM, and Apple use MTDC to complement their own data centers [Islam et al. 2016; Ren 2017]. For example, Apple owns and operates five data centers but also leases data center space at different MTDCs worldwide, all of which consumed a total of 778 GWh of electricity in 2016 [Apple 2017]. As shown in Fig. 2a, about one fifth of this electricity use was from MTDCs. The trend of Apple's increasing reliance on MTDCs can be seen in Fig. 2b. The energy use footprint by MTDCs has grown by more than four fold in four years.

There are 4,000+ colocation data centers worldwide with 40% of them in the U.S. [Data Center Map 2017]. The MTDC market size is estimated at \$31.5 Billion and increasing at a 14.6% annual rate [MarketsandMarkets 2017]. Some of the biggest data centers in the world are multi-tenant data centers. For example, the Digital Reality data center in Chicago has a total area of 1.1 million square feet and draws more than 100 MW power from the utility [Digital Realty 2017]. More importantly, MTDCs consume about one fifth of the entire data center electricity use in the U.S. This is about five times more electricity than Google type hyperscale data centers [NRDC 2014]. Therefore, the



(a) Electricity consumption percentage



(b) MTDC energy use

Fig. 2. (a) Proportion of electricity consumption by different Apple data centers in Fiscal Year 2016. (b) Apple’s multi-tenant data center energy consumption trend [Apple 2017].

MTDC is an important type of data center and making them energy efficient is equally, if not more, important. However, there is a gap between the operator and the tenants as identified by Ren et al. in their seminal paper [Ren and Islam 2014]. The operator might want to manage power or adopt energy efficient techniques in their data center but they lack control over the physical servers or the workload, which are owned and managed by the tenants independently. Same power management techniques that have been proposed for data centers with a single owner and operator are not directly applicable to an MTDC.

Only very recently has there been some research to address the problem of power management in MTDC [Islam et al. 2014, 2016; Malla and Christensen 2017; Ren and Islam 2014; Tran et al. 2015b; Zhang et al. 2016]. Since, an MTDC operator sells data center space/power to tenants, an MTDC can be viewed as a market with the operator as the seller and the tenants as the buyers. Hence, some kind of market mechanism is required to incentivize tenants in order to have coordination between the MTDC operator and the tenants. In this paper, we do a comprehensive survey of power management techniques for MTDC with a special focus on power oversubscription.

Oversubscription is essential for efficient use of resources. For example, telephone lines and Internet are oversubscribed, airline tickets are overbooked. Power oversubscription, that is, installing servers such that aggregated peak power is greater than the power capacity limit, is desirable in data centers as it increases the overall utilization of the power infrastructure and reduces the amortized cost associated with power provisioning. MTDC operators may want to oversubscribe power for greater profit while tenants themselves may want to oversubscribe power they have purchased to reduce cost. We formally define what power oversubscription is as well as outline in more detail the profit, risk, and challenges associated with oversubscribing power capacity of MTDC in section 2.

The outline of the remainder of this paper is as follows. In section 2, we discuss the different aspects of MTDC ranging from its power hierarchy to different leasing schemes. In section 3, we describe our new taxonomy and use it to classify the existing literature. In the following three sections, 4, 5, and 6, we introduce research work in each of the three main categories. In section 7, we discuss related surveys in power management of data centers. In the last section we discuss open problems and future directions.

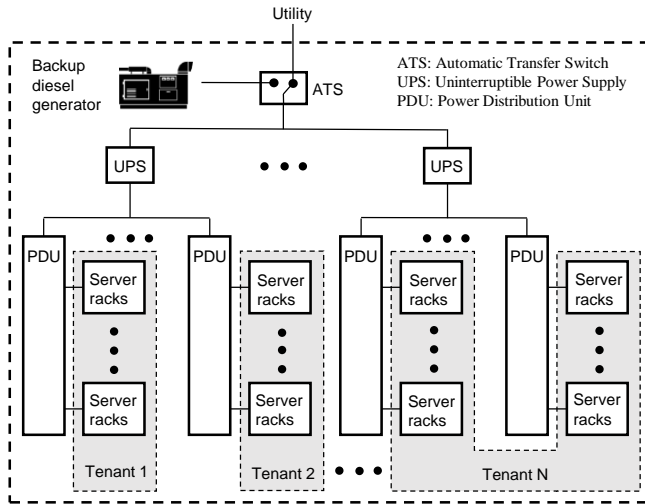


Fig. 3. Power delivery infrastructure of an MTDC with  $N$  tenants.

## 2 BACKGROUND

In this section we describe the power hierarchy of MTDC, concept of oversubscribing the power hierarchy, demand response programs from the utility, and different types of costs for tenants and the operator.

### 2.1 MTDC Power Fundamentals

*2.1.1 Power Hierarchy.* Fig. 3 shows a simple MTDC power infrastructure hierarchy without any redundancy. After a series of voltage step-downs from high to low voltage (typically, 480 Volts), power from the utility enters the data center. On-site backup diesel generators are generally present to provide power in case of a utility power outage. An Automatic Transfer Switch (ATS) chooses between the two inputs automatically. Uninterruptible Power Supply (UPS) with batteries regulate power from the utility removing any sags or spikes as well as performing power factor correction of the load side. UPS's also enhance reliability by supplying stored energy in case of utility power outages where backup diesel generators take a few seconds to start and bear the full load. These are generally double conversion UPS with a rectifier (AC-DC) step which also charges the batteries and an inverter (DC-AC) step.

The conditioned output of the UPS's goes to multiple Power Distribution Units (PDU) spread across the server room. Each PDU transforms voltage to that suitable for the installed IT equipment (typically 110 volts in the U.S.) and splits to multiple outputs with circuit breakers for protection (similar to a residential circuit breaker panel) which ultimately go to racks and servers [Barroso et al. 2013]. Alternative configurations are also possible, for example having redundant components (UPS and PDU) and a power bus for greater reliability and availability at increased construction cost. Multiple distributed UPS's per rack or server instead of a centralized UPS or a high-voltage DC power distribution system can be used to avoid double conversion losses for increased efficiency. Increasingly, data centers also have local renewable energy sources, like wind and solar, to minimize their dependency on the grid and reduce their carbon footprint [Deng et al. 2014].

As mentioned earlier, the MTDC operator owns and manages the power infrastructure. Multiple tenants can rent out a single server rack, a cage (with several racks), or even the entire room/building depending upon their needs.

**2.1.2 Power Consumption.** Two major sources of power consumption in an MTDC are the IT equipment of tenants and the cooling system used to maintain the temperature and humidity inside the data center.

**IT equipment power:** Tenants typically have IT equipment including, servers for computing, storage disks for storing data, and networking equipment for communication. Servers are the predominant power consumers and their power consumption varies according to their utilization. Hence, power consumption by tenants is generally modeled as depending upon the number of servers they have powered on and the average utilization of these servers. For example, consider an MTDC with  $N$  tenants with each tenant,  $i \in \{1, 2, \dots, N\}$ , having  $M_i$  homogeneous servers. A server consumes  $P_i^{idle}$  power when completely idle,  $P_i^{max}$  power when fully utilized, and tenants can choose to power down some of their servers when not needed, to reduce the power use. The time under consideration (usually, the monthly billing cycle) can be divided into  $T$  time slots, with a time interval  $t \in \{1, 2, \dots, T\}$ . If, for each time interval  $t$  and tenant  $i$ , the average workload arrival rate is  $\lambda_i(t)$ , the average service rate of a server is  $\mu_i$  and the tenant has  $m_i(t)$  servers powered on, then the mean utilization of a server is  $\lambda_i(t)/m_i(t)\mu_i$ . A linear power model is a popular way to estimate a server's power consumption [Fan et al. 2007; Tran et al. 2016d]. The total power consumption by the tenant,  $x_i(t)$ , is given as

$$x_i(t) = m_i(t) \left[ P_i^{idle} + \frac{\lambda_i(t)}{m_i(t)\mu_i} (P_i^{max} - P_i^{idle}) \right] \quad (1)$$

and the total IT power consumption of the entire MTDC,  $p_{IT}(t)$ , is the sum of IT power consumption by all tenants

$$p_{IT}(t) = \sum_{i=1}^N x_i(t). \quad (2)$$

**Cooling system power:** The heat generated from IT equipment must to be removed for reliable operation of an MTDC. In a typical MTDC, computer room air conditioning (CRAC) units are located in server rooms with raised floor to take in the hot air produced by servers. CRAC units have heat exchangers, where the hot air is cooled, typically using chilled water from a water-cooled chiller [Barroso et al. 2013]. The cold air is pumped back to the server room through the underfloor plenum. Rooms with raised floor have perforated tiles in front of the server racks such that cool air blown from beneath enters the servers, and the process repeats. Cooling systems consume a significant amount of power (up to 40% of entire data center power [Barroso et al. 2013]). One way of modeling the cooling power of an MTDC,  $p_{cool}(t)$ , is as [Zhang et al. 2013]

$$p_{cool}(t) = \frac{m * C_p * (T_{in} - T_{out})}{COP(T_{out})} \quad (3)$$

where  $m$  is the mass flow rate,  $C_p$  is the specific heat,  $T_{in}$  and  $T_{out}$  are the temperatures at the inlet and outlet of the CRAC unit respectively.  $COP(T_{out})$  is the coefficient of performance of the CRAC unit which depends on the outlet temperature. For example, the COP of a typical chilled-water CRAC unit at HP Lab's Utility data center is given as [Moore et al. 2005]

$$COP(T_{out}) = 0.0068T_{out}^2 + 0.0008T_{out} + 0.458 \quad (4)$$

which shows that COP increases with increase in outlet temperature.

**Power usage effectiveness:** The total power used by an MTDC,  $p_{MTDC}(t)$ , can be approximated from equations 2 and equation 3 as

$$p_{MTDC}(t) = p_{IT}(t) + p_{cool}(t). \quad (5)$$


A popular metric to depict the total power consumption of a data center is through the power usage effectiveness (PUE). PUE of a data center is defined as “the ratio of total building power to IT power” [Barroso et al. 2013]. The ideal value is 1 where all the power consumed by the MTDC goes to the IT equipment. But there is always some power overhead due to mechanical (e.g., cooling) and electrical systems (e.g., lighting), with power for cooling being the most significant portion. Hence, real-world data centers have PUE greater than 1 and can range up to 2 [Barroso et al. 2013; Ren and Islam 2014]. PUE can also be used to compare the energy efficiency between data centers, with lower values indicating more efficient infrastructures. For example, PUE for small and inefficient data centers is about 2.0 while large data centers with advanced airflow and cooling management have PUE of 1.1 or lesser and the average for all data centers is about 1.8 [Shehabi et al. 2016]. Instantaneous PUE of a data center will exhibit diurnal or seasonal variations [Islam et al. 2015], for simplicity we consider the annualized average PUE. The total power used by an MTDC can be expressed as

$$p_{MTDC}(t) = PUE * p_{IT}(t). \quad (6)$$


## 2.2 Power Oversubscription

The critical power capacity of a data center is defined as the actual power available for IT equipment ignoring other power consumption overhead (e.g., cooling). Infrastructure cost of a data center scales linearly with its critical power capacity and can cost between \$11 to \$25 per provisioned critical watt [Barroso et al. 2013; Karidis et al. 2009]. The amortized capital expenditure for building the power infrastructure of a data center can be 1.5 times higher than its operational cost over its entire lifetime [Wang et al. 2012b]. Given the high cost of building the power infrastructure of a data center, 100% utilization would be ideal. But, this isn’t the case in actual data centers. This is due to the power use of a server being proportional to the workload and servers are not 100% utilized all the time. The average server utilization in well-tuned hyperscale data centers is about 45% [Shehabi et al. 2016]. Furthermore, it has been found that the conservative approach of installing servers such that the sum of maximum power use by individual servers is less than the critical power capacity, leads to further inefficiencies and wasted resources. Even though a server may frequently consume peak power, it is rare for a large group of servers to peak simultaneously.

Fan et al. [2007] characterized the power usage of a cluster of 5000 servers at Google for different real world workloads over a period of six months. The main finding is summarized in Fig. 4a. Fan et al. found that, as we move up the power hierarchy from rack to PDU to cluster, dynamic range of power usage is narrower and the actual maximum power use compared to potential maximum power use decreases. For example, at the rack level the actual peak is 93% of maximum possible, but at the cluster level, the actual power use did not exceed 72% of the aggregated peak power. This is to say that more than a quarter of the critical power capacity of the data center never got utilized. The authors identified and explored the opportunity for maximizing the utilization of data center power infrastructure through *power oversubscription*, that is, installing more servers than allowed by the critical power capacity. Similar observations were made by Islam et al. [2016] when they studied the power usage characteristics of 10 tenants in an MTDC. Fig. 4b shows the CDF of power use by 1, 5 and 10 tenants normalized to the maximum possible power use. We see that as the number of tenants increases, the maximum normalized aggregate power use of all tenants decreases. This is due to the fact that, different tenants in an MTDC have different workloads and the event of correlated spikes is very rare. Hence, it is possible to oversubscribe the power hierarchy



(a) Operator-owned data center (Google)



(b) Multi-tenant data center

Fig. 4. (a) CDF of normalized power usage of Google data center at rack, PDU and cluster level [Fan et al. 2007]. (b) CDF of normalized power usage of an MTDC for different group of tenants [Islam et al. 2016].

Table 1. Profit and risk of power oversubscription in MTDC [Islam et al. 2016]

<b>Oversubscription</b>	Extra revenue (\$/kW/year)	Probability of overloading
<b>10%</b>	180	1%
<b>15%</b>	270	1%
<b>20%</b>	360	2%
<b>25%</b>	450	3%

in an MTDC for increased utilization and efficiency. This observation is key to the future work proposed in this paper.

**2.2.1 Profit and risk of power oversubscription.** MTDC power oversubscription is profitable to the operator as they are able to host more tenants than allowed by the power capacity. Tenants are generally charged at a fixed rate based on their power subscription [Islam et al. 2015]. So, power oversubscription leads to increase in the operator’s profit (revenue increases without additional cost). However, anytime the power infrastructure is oversubscribed, there exists the risk of having a *power overload*, that is, the total power consumption is above the power capacity. This can have disastrous results. For example, at the rack level or PDU, the circuit breakers may trip, leading to server downtime and service outage. Similarly, at the cluster or data center level, where the power limit is contractual rather than physical, there may be fines from the utility [Malla and Christensen 2017]. Considering a typical leasing cost of \$150/kW/month and practical power use characteristic of tenants, the potential profit to operator and the associated risk (power overloading probability) at different power oversubscription level is shown in Table 1 [Islam et al. 2016].

In order to avoid power overload situation and ensure safe power oversubscription, some kind of *power capping* mechanism is required. Power capping is a method of restricting the power use of a server or group of servers such that certain power limit/budget isn’t exceeded. Power management technologies in modern servers have this power capping capability. An example is Intel’s Running Average Power Limit (RAPL) interface which limits a server’s average power use over a time window [David et al. 2010; Rountree et al. 2012]. Capping the power of a server may affect its performance. Power capping at the server level forms the building block of power capping mechanisms at higher levels (for example, PDU and data center) in operator owned data centers

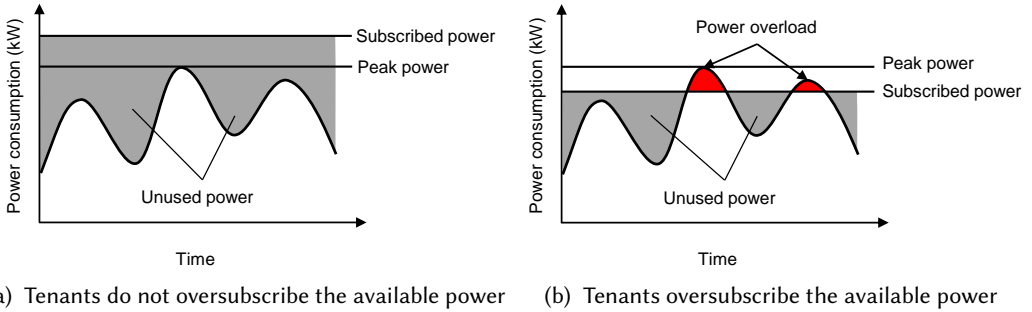


Fig. 5. Power overload may occur when tenants oversubscribe the power capacity to reduce their power subscription cost.

[Azimi et al. 2017; Wang et al. 2012a; Wu et al. 2016], and hence enable safe power oversubscription. But in an MTDC, the operator lacks control over the servers for any kind of power capping.

**2.2.2 Tenant-side power oversubscription.** It is possible for tenants to oversubscribe (putting in more servers than allowed by power capacity) their own power limit as well, without the operator's knowledge, in order to reduce the cost for power subscription. Tenants are generally charged by the MTDC for power they subscribe to, rather than the actual energy they use or for their peak power consumption [Palasamudram et al. 2012]. Fig. 5 shows the time varying nature of tenant's power consumption. Normally, tenants subscribe to power to ensure that peak power use is below the power capacity (fig. 5a). This can lead to wastage of available power as peak power consumption by tenants may be rare. Tenants can oversubscribe power to better utilize the available power and reduce their operating cost as well (lower leasing cost). However, this can lead to power overload as shown in fig. 5b if proper control mechanisms are not in place. In practice, MTDC operators do not allow tenants to consume more than 75% [Verizon 2014] or 80% [TekLinks 2017] of the rated power circuit limit to ensure safety. Here, the oversubscription is in lower (rack or cage) levels in the power hierarchy.

If tenant  $i$  has  $M_i$  servers, each with maximum power rating of  $P_i^{max}$ , then the maximum possible power use by the tenant is,  $x_i^{max} = M_i * P_i^{max}$ . Assuming that the tenant subscribes to a power level of  $S_i$  from the operator, we say the tenant has oversubscribed the available power when the maximum possible power consumption by the tenant is greater than its subscribed capacity, that is,

$$x_i^{max} > S_i. \quad (7)$$

Tenants may use power capping techniques to avoid power overload when oversubscribing in such situations. Alternatively, tenants may deploy batteries of reasonable capacity in order to handle local peak power usage [Palasamudram et al. 2012].

**2.2.3 Operator-side power oversubscription.** The MTDC operator can oversubscribe the power hierarchy similar to power oversubscription in operator owned data center. This can be completely transparent to the tenants. Let  $C_0$  be the critical power capacity of an MTDC, and each tenant  $i$  has a power subscription of  $S_i$ . The MTDC is said to be oversubscribed by the operator when the total allocated peak power to the tenants is greater than the critical power capacity, that is,

$$\sum_{i=1}^N S_i > C_0. \quad (8)$$



Furthermore, the operator can oversubscribe at different levels in the power hierarchy. Let there be  $R$  PDU in the MTDC each with a power limit of  $C_j$  for  $j \in \{1, 2, \dots, R\}$  and denote  $\mathcal{R}_j$  as the set of tenants drawing power from PDU  $j$ . Assuming a tenant only belongs to one PDU, i.e.,  $\mathcal{R}_j \cap \mathcal{R}_k = \emptyset$ , if  $j \neq k$  and  $\mathcal{R}_0$  is the set of all tenants, we can have oversubscription at the PDU level in addition to the entire data center level, that is,

$$\sum_{i \in \mathcal{R}_j} S_i > C_j, \text{ for } j = 1, 2, \dots, R. \quad (9)$$

Safe oversubscription of MTDC is not straightforward compared to operator owned data center because the operator lacks control over the tenant servers. Tenants will be unaware of the power oversubscription, which is an operator created artifact. There is no mechanism in place to avoid power consumption from tenants peaking simultaneously and creating a power overload situation. We look into existing research work that focuses on enabling safe power oversubscription in MTDC by an operator [Islam et al. 2016; Malla and Christensen 2017]. This limited existing work is covered in section 6 of this paper.

### 2.3 Electric Utility

The main source of power for an MTDC is typically the electric utility grid. Power from the grid is considered stable and reliable. In case of grid power failure, UPS with batteries help in transition to on-site diesel generators.

**2.3.1 Utility Pricing.** Being a large power consumer, MTDCs are charged industrial rates by the utility. The electricity bill generally has two components, volume (or energy) charge and demand (or peak power) charge [Wang et al. 2014].

**Volume charge:** This component of the bill is for the energy consumption by the MTDC, charged as dollar per kWh. The charge may vary throughout the billing period (for example, one month) with higher price when the demand is high. For example, in time of use (TOU) pricing we could have different time blocks in a day marked as on-peak with higher rates and off-peak with lower rates. A more complex version is real time pricing (RTP) where the electricity price varies every hour according to demand and the customers are informed an hour-ahead or a day-ahead [Albadi and El-Saadany 2008] of time.

**Demand charge:** This part of the bill is for the peak power consumption by the MTDC during the billing period, charged as dollar per kW. The power use by the MTDC is monitored by the utility and the highest average power consumption in a sampling interval (for example, 15 minutes) is considered as the peak demand (measured in kW) for that billing period. Thus, demand charge is only determined at the end of the billing cycle. Utilities have a demand charge in order to compensate for the expensive overhead infrastructures (for example, generators, transformers, and substations) they have to build to meet the power demand at all times [Islam et al. 2015]. Demand charge can constitute up to 40% of the entire electricity bill of the MTDC [Govindan et al. 2011].

If we denote  $\alpha(t)$  as the volume charge at time  $t$ ,  $\beta$  as the demand charge and time interval length to be  $\tau$ , then the total electric bill of an MTDC in a billing cycle is

$$E_{bill} = \sum_{t=1}^T \alpha(t) \cdot p_{MTDC}(t) \cdot \tau + \beta \cdot \max_t p_{MTDC}(t). \quad (10)$$

An MTDC has a contract with the utility on the maximum power it can draw. There is usually a fine involved for power use by MTDC that exceeds the contracted demand [MTEMC 2017].

**2.3.2 Demand Response.** In order to make power grids more reliable, efficient, sustainable, and to incorporate highly varying renewable energy sources into the grid, utilities offer different demand

response (DR) programs to their customers. DR refers to reshaping the consumer’s electricity consumption profile using time-varying price or offering incentives for power reduction when requested, for example, when there is peak demand the wholesale market price for electricity is high, or the reliability of the grid is in danger [Siano 2014]. DR programs help match demand to supply by coordinating power use with power generation, rather than the traditional way of matching supply to demand [Liu et al. 2014].

As noted in a survey [Wierman et al. 2014], “data centers are large loads, but are also flexible” which makes them a good candidate for demand response programs. Data centers are generally viewed as power hungry customers, but can alternatively be viewed as energy storage devices due to their flexible demand. Utilities can extract this flexibility using demand response programs. A data center can participate in demand response programs in various ways as described below.

- **Temporal load balancing:** Data centers have a mix of delay-sensitive interactive workloads (such as online shopping and search requests) and delay-tolerant batch jobs (such as web indexing, MapReduce jobs, and scientific calculations). Exploiting the temporal flexibility of delay-tolerant workload to scheduling it to a different time is a well-studied technique [Liu et al. 2012; Luo et al. 2014] for data center load shaping.
- **Geographical load balancing:** In addition to shifting workload in time, it is possible to shift it in space for large internet-scale systems. In order to reduce load on a particular data center, workload can be rerouted to a different data center. Geographical load balancing has been used for reducing cost [Qureshi et al. 2009] and reducing carbon footprint [Le et al. 2010] of a data center.
- **IT equipment power management:** Power consumption of servers depends upon their utilization. Putting idle or lightly loaded servers to sleep states [Lin et al. 2013], varying CPU power states through Dynamic Voltage and Frequency Scaling (DVFS) [Wu et al. 2014], and capping server power use [Lefurgy et al. 2008] are some of the ways in which a data center can reduce its power consumption. This, however, can result in performance degradation and must be done in a controlled fashion.
- **Using stored energy:** Data centers have on-site backup diesel generators and energy storage devices (batteries) to handle utility power outages and provide a reliable operation. They can be used to reduce the electricity drawn from the power grid for demand response [Wang et al. 2012b].

In an MTDC, the operator can participate in DR programs by either using stored energy (that may only last for a few minutes), or using backup diesel generators (that are environmentally unfriendly) only. The IT control mechanisms described above are under the tenant’s control, who do not have any incentive for DR. Hence, participation in DR programs by MTDC operator is limited and poses a unique challenge. Recent research [Guo et al. 2017; Ren and Islam 2014; Tran et al. 2015a; Zhang et al. 2015b] has focused on this problem to design incentive programs for tenants so that MTDC can participate in DR programs. We describe this existing work in section 4.

## 2.4 MTDC Economics

*2.4.1 Operator Leasing Schemes.* There are various ways in which an MTDC operator can charge tenants a monthly fee for power and space. Three major pricing models are described below [Islam et al. 2015]. The operator can also charge extra for other additional services, such setup fees and network fees.

**Space-based pricing:** In this pricing model, the operator charges a monthly fee according to the space occupied. It may be a per square foot charge or a per rack space charge, measured in units of “U” (1.75 inches) [Optimal Networks 2015]. Modern data centers are, however, moving

from space-based pricing to power-based pricing as critical power is increasingly becoming the primary source of cost [Dines 2011; Islam et al. 2015].

**Power-based pricing:** MTDC generally charge tenants based up on their power subscription. All costs for space and energy use are bundled into the flat monthly fee charged for each kW of power supplied to the tenant. At the time of contract, the tenant may specify the amount of power they want to subscribe to (their power limit). Power subscription charges are in the range of \$150-200 per kW per month [Islam et al. 2015]. This is a widely used pricing model in MTDCs.

**Energy-based pricing:** In both space-based and power-based pricing model, the tenant isn't being charged for metered powered use. Rather, the monthly bill is a flat fee. This kind of pricing scheme does not encourage tenants to be energy efficient in their operation. On the contrary, in energy-based pricing model, tenants are charged for actual energy use on top of a flat power subscription fee (which is less than the charge in power-based pricing [Islam et al. 2015]). Such a pricing model encourages tenants to be energy efficient and are usually found in wholesale MTDCs where the operator leases power in excess of 100 kW to a single tenant [Dines 2011].

**2.4.2 Tenant Costs.** Tenants may adopt energy efficient techniques, like powering down some servers during periods of light workload, if they have proper incentive to do so. Powering down servers or putting servers in some kind of power saving sleep mode can cause inconvenience to tenants [Oo et al. 2017], like performance degradation, time overhead of powering back up, and wear-and-tear of equipment. This is generally modeled using monotonically increasing cost [Islam et al. 2014] or monotonically decreasing utility [Malla and Christensen 2017] as a function of number of servers turned off. Such cost functions are private to individual tenants and are determined at their own discretion. Two types of costs generally associated with tenants are delay performance cost and switching cost [Lin et al. 2013; Oo et al. 2017; Ren and Islam 2014].

**Delay performance cost:** Tenants may have computational tasks or may be providing some kind of interactive service to their own customer. Powering down servers can cause delay in computation and service time performance to degrade. Delay performance cost models the cost associated with an increase in delay that may also cause violation of Service Level Agreements (SLA) [Lin et al. 2013; Ren and Islam 2014]. It is a common practice to model each active server as M/M/1 queue, for which the average delay for tenant  $i$  with  $m_i$  active servers

$$d_i = \frac{1}{\mu_i - \frac{\lambda_i}{m_i}} \quad (11)$$

increases as the number of available servers decreases. A common approach is to model cost as a linearly increasing function of the delay after it exceeds certain threshold,  $d_i^{th}$ , defined in the SLA of the tenant with their respective customers [Islam et al. 2014; Lin et al. 2013; Oo et al. 2017; Ren and Islam 2014]

$$\text{delay cost} = \kappa_i^d \cdot \lambda_i \cdot [d_i - d_i^{th}]^+ \quad (12)$$

where  $\kappa_i^d$  is a constant factor that converts experienced delay into some monetary value (dollars) and  $[\cdot]^+ = \max\{\cdot, 0\}$ .

**Switching cost:** When switching servers between power saving modes, some additional cost such as server wear-and-tear cost, state migration cost, and startup energy overhead cost are incurred [Guo et al. 2017; Islam et al. 2014; Lin et al. 2013; Ren and Islam 2014]. We can model these cost components as switching cost which grows linearly with the number of servers powering down. For a tenant  $i$  with total  $M_i$  servers but only  $m_i$  servers active, we can express switching cost as

$$\text{switching cost} = \kappa_i^s \cdot (M_i - m_i) \quad (13)$$

where  $\kappa_i^s$  is a constant scaling factor that denotes the cost of switching one server.

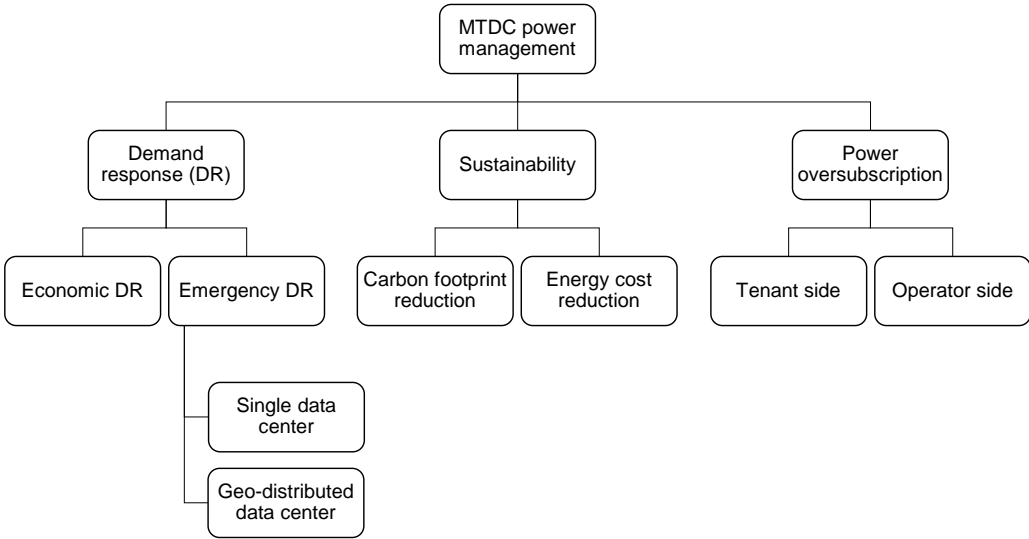


Fig. 6. Taxonomy of power management in multi-tenant data centers.

### 3 TAXONOMY

In this survey paper, we focus on works in power management of MTDC and categorize them according to the research problem they try to solve. Fig. 6 shows our taxonomy and in Table 2 we categorize all surveyed works according to our taxonomy. These works can be categorized as follows.

#### 3.1 Demand response participation

Demand response (DR) programs aim to make power grids more stable and reliable. During periods of peak power demand or when the wholesale electricity price is high, utilities may offer monetary rewards to consumers who decrease their electricity usage. We consider such programs as *economic demand response* programs, where the consumer can voluntarily participate and reduce their power demand at their own will and to their economic benefit [Tran et al. 2015a,b, 2016a]. Another type of DR program is *emergency demand response* which refers to demand response in case of an emergency (for example, extreme weather that stresses a grid) when the grid is about to fail. Participants usually sign a contract with the utility for a fixed energy reduction during emergency (signaled by the utility) in return for a monetary reward [Chen et al. 2015]. Emergency DR is the most widely adopted DR program representing 87% of the total DR capability across all reliability regions in the U.S. [Managan 2014], and it acts as last line of defense against a cascading power failure. Furthermore, data centers have been recognized as key participants in emergency DR programs by the EPA [EnerNOC 2013]. The main difference between economic and emergency DR is that the former is voluntary with flexible energy reduction while the latter is mandatory and has a fixed amount of energy to reduce. While there have been a number of studies to enable DR participation for operator owned data centers [Liu et al. 2014, 2013; Wierman et al. 2014], they cannot be applied to MTDC directly due to the lack of coordination between the operator and tenants. An indirect control method is required. Participation by MTDC in these DR programs will have some reward mechanism to incentivize tenants to reduce power use. The operator actively

Table 2. Classification of existing work using our taxonomy

Category	Subcategory		Works
Demand response (DR)	Economic DR		[Ren and Islam 2014; Tran et al. 2015a,b, 2016a; Xu et al. 2017]
	Emergency DR	Single data center	[Ahmed et al. 2015; Chen et al. 2015; Guo et al. 2017; Niu et al. 2016; Sun et al. 2016b, 2015; Tran et al. 2016b, 2015b, 2016c,d; Wang et al. 2015, 2017; Zhang et al. 2015b; Zhao et al. 2016]
		Geo-distributed data center	[Sun et al. 2016a; Zhang et al. 2016]
Sustainability	Carbon footprint reduction		[Islam et al. 2014; Mahmud and Iyengar 2016]
	Energy cost reduction		[Guo and Pan 2015; Guo et al. 2018, 2017; Islam et al. 2015; Zhang et al. 2015a]
Power oversubscription	Tenant side		[Palasamudram et al. 2012]
	Operator side		[Islam et al. 2016; Malla and Christensen 2017]

passes on the financial gain from the electric utility (reward for participating in DR programs) to tenants for coordinated power demand reduction.

### 3.2 Making MTDC more sustainable

An MTDC might want to become environmentally friendly and green to have a good public image, attract pro-sustainable tenants or simply reduce energy cost by consuming less energy. Being green, an MTDC can also earn LEED (Leadership in Energy and Environmental Design) certification which may have tax benefits [Islam et al. 2014; USGBC 2017]. In this category, we review works that focus on reducing the carbon footprint of an MTDC [Islam et al. 2014; Mahmud and Iyengar 2016] or cutting their electricity bill [Guo and Pan 2015; Guo et al. 2018, 2017; Islam et al. 2015; Zhang et al. 2015a]. In both cases, coordination among tenants is necessary to make the MTDC energy efficient and hence, more sustainable.

### 3.3 Enabling power oversubscription

Power oversubscription can lead to power overload which may result into power outages. On average a data center power outage costs is estimated at \$740,357 [Ponemon Institute 2016]. To ensure safe power oversubscription, tenants can employ most of the control mechanisms that have been proposed for operator-owned data centers [Azimi et al. 2017; Bhattacharya et al. 2013; Wang et al. 2012a; Wu et al. 2016] to reduce their power use. However, if the operator wants to oversubscribe the MTDC, the operator cannot directly reduce IT power consumption during power overload as they lack control over tenant servers and workload. One way is to tap into the energy stored in batteries of UPS [Govindan et al. 2011] or use the on-site diesel generator. This prevents overdrawing power from the utility. However, even if the power infrastructure can handle the excess power, the cooling capacity may be exceeded [Islam et al. 2016] leading to overheating or thermal shutdown of servers. The key idea is to coordinate power consumption of multiple tenants,

in an otherwise uncoordinated MTDC. We review works that develop incentive mechanisms that MTDC operators can adopt to encourage tenants to reduce power consumption, while minimizing the impact on their performance, to avoid power overload in an oversubscribed MTDC.

## 4 DEMAND RESPONSE PARTICIPATION

In this section we review works that focus on sustainable MTDC participation in different types of demand response programs.

### 4.1 Economic Demand Response

Various works have looked into MTDC participation in economic DR programs.

*4.1.1 Problem description.* MTDC operators may have financial incentive to participate in economic DR programs. Generally, data centers can lower the total power use of the data center by powering down servers or throttling them to lower power consumption states and/or shifting workload in time or space (to a different data center). As tenants have no incentive to take part in DR program, the problem is finding incentive mechanisms such that the operator can pass on the compensation from utility to its tenants to participate in economic DR programs.

*4.1.2 Existing solutions.* Ren et al. [2014] were among the first to study demand response participation for MTDC and to identify the split-incentive in MTDC – the operator may want to participate in demand response programs but don't have control over the servers to reduce the power consumption, while tenants who can reduce IT power use have no incentive to participate in demand response programs. Ren et al. propose a reverse auction mechanism called iCODE to break this split incentive by rewarding tenants for power reduction. The authors formulate the problem as energy reduction maximization problem subject to the constraint that total reward to tenants is less than what operator receives from utility for demand response participation. During a demand response period, the utility notifies the MTDC of the compensation rate. The MTDC operator in turn notify tenants to reduce their power consumption. Tenants, then, voluntarily prepare a set of bids (energy reduction and payment requested for it) and submit it to the operator. The operator in turn selects the winning bids for each tenant by trying to maximize the energy reduction while keeping the payments lower than what it receives as compensation from the utility. The problem of selecting winning bids is NP-hard, hence the authors use branch and bound technique to approximate a close to optimum solution.

It should be noted that iCODE cannot guarantee truthfulness of strategic tenants. Game-theoretic approaches for economic demand response have been proposed to tackle strategic tenants. Tran et al. [2015a; 2015b] design a reward program in which the operator first sets the reward rate to maximize their total profit (revenue minus cost), then tenants reduce power consumption to maximize their own profit. The interaction between tenants and operator is modeled and analyzed as a two-stage Stackelberg game. The reward is set by the operator (leader) and tenants (followers) respond to it, where both the tenants and the operator act to maximize their own utility (profit). Using backward induction, the authors examine the game and find the Stackelberg equilibrium. However, in this incentive mechanism, tenants need to communicate all their cost and workload information to the operator beforehand, such that operator knows the tenant's best response to the reward rate. The operator then solves the profit maximization problem in a central fashion to determine the optimum reward rate. This may not be practical in an MTDC as tenants do not want to reveal their cost and workload information to the operator. This work is extended in [Tran et al. 2016a] to account for the role of utility in demand response program. Rather than assuming utility as an outside component, utility and its compensation rate determination is incorporated as part

Table 3. Comparison of existing works in economic demand response

Paper	Objective	Solution	Technique	Result
Ren and Islam [2014]	Max. energy reduction	iCODE	Bidding	50% hourly energy reduction
Tran et al. [2015a]	Max. utility	A linear time algorithm	Operator sets reward rate	Optimal reward rate and operator profit
Tran et al. [2015b]	Max. utility	A linear time algorithm	Operator sets reward rate	Optimal reward rate and operator profit
Tran et al. [2016a]	Max. utility	Reward-to-Reduce	Operator sets reward rate	Stackelberg equilibrium exists
Xu et al. [2017]	Max. utility	Stackelberg equilibrium	Operator sets reward rate iteratively	Stackelberg equilibrium convergence

of the game. The authors propose a Reward-to-Reduce (R2R) mechanism in which the utility first determines the compensation rate for the operator, who then decides the reward rate for tenants.

Similarly, Xu et al. [2017] also analyze the interaction between tenants and operator as a non cooperative Stackelberg game where operator is the leader and tenants are the followers. But in their mechanism, tenants do not have to reveal all their cost parameters to the operator. Tenants bid energy reduction capability and the operator sets a reward rate in an iterative fashion. The authors prove that the iteration converges to a unique Stackelberg equilibrium. Through practical simulation under different scenarios the authors show that tenant's bids, operator reward rate and profit for all of them converges after 20 to 30 iterations. All these works are summarized and compared in Table 3.

## 4.2 Emergency Demand Response

Most of the works on power management of MTDC have focused on DR participation, more specifically on emergency DR.

*4.2.1 Problem description.* This DR program differs from the economic DR program in that while economic DR program is flexible and participation is voluntary, emergency DR program is mandatory where the utility define a fixed energy to be reduced and non compliance can result in a financial penalty. Since servers and their corresponding workload are controlled by the tenants in an MTDC, they need incentives to take part in emergency DR programs. The problem is finding incentive programs that are efficient, robust, not susceptible to cheating, minimize the impact to tenants, and achieve total energy reduction greater or equal to that demanded by the utility.

*4.2.2 Existing solutions.* Several works that have focused on MTDC power management have tried to solve the problem of sustainable MTDC emergency DR participation. The majority of these works have focused on a single MTDC while some have focused on geo-distributed MTDCs.

**Single data center:** Zhang et al. [2015b] proposed Truth-DR, a reverse auction based incentive mechanism to reward tenants for power reduction. During an emergency DR event, the operator solicits bids as an amount of energy reduction and cost associated with such energy reduction, from tenants. Then, the operator selects winners and calculates payment amounts. The problem of deciding the winning bid is stated as a social cost minimization problem, which is NP-hard. Social cost of an MTDC is defined as the sum of aggregate cost of all tenants (cost for energy reduction) and the cost of the operator (cost of using backup generators in the case where energy reduction by tenants is not sufficient). The authors come up with a primal-dual based 2-approximation algorithm.

Using realistic simulations they show how Truth-DR can incentivize tenants for energy reduction in a close-to-optimal way. Wang et al. [2015; 2017] propose a similar incentive mechanism of tenants bidding for energy reduction but they consider the case where all tenants coordinate among each other in addition to the coordination of tenants with the operator.

Truth-DR is not fair in the sense that two tenants reducing the same amount of energy may be rewarded differently. Sun et al. [2015] proposed FairDR, a reverse auction similar to Truth-DR, but solve the fairness issue. In contrast to Truth-DR which considers the case of a single emergency DR event in isolation, FairDR considers the case of multiple emergency signals (for example, multiple consecutive emergency DR events in a single day). Only one bid is requested from tenants for the entire time window which may consist of multiple time slots with energy reduction signals. Tenant energy reduction decisions for each time slot are made in an online fashion without the knowledge of future energy reduction requirement. Through theoretical analysis they show that their mechanism is fair, truthful, and has a bounded competitive ratio in social cost saving. The authors extend their work in [Sun et al. 2016b] for tenants with delay-tolerant batch jobs. They design an online reverse auction mechanism where both tenant jobs and emergency DR signals can extend to multiple time slots. Tenants bid their valuation function (cost for energy reduction) and looking at tenant valuation functions, the operator assigns energy reduction amounts to tenants and diesel generator such that operator's cost is minimized. The operator then calculates the reward for each participating tenant. The authors also develop an online scheduling algorithm for tenants to maximize their individual utility. Zhao et al. [2016] developed a reverse auction called TECH that considers the cooling infrastructure of the MTDC as well. The authors use a server heat recirculation model to find the optimum temperature of the supply air such that the server temperature does not exceed a threshold, while consuming minimum energy. Tenants submit bids containing servers they want to power down (energy reduction) and the desired compensation for it (cost) similar to other bidding mechanisms. The operator then selects winners and calculates their payment such that energy reduction target, cost constraints (operator cost is less than using diesel generator for energy reduction), and temperature constraints (server temperature below a thresholds) are met.

Chen et al. [2015] propose ColoEDR, a pricing mechanism based on parameterized supply function bidding. An advantage of using the supply function is that tenants do not have to reveal their private cost information to the operator as they did in Truth-DR and FairDR. The basic idea is that tenants bid for energy reduction during emergency situations using a parameterized supply function. Through a parameter in this supply function, they can express their energy reduction flexibility/willingness. The operator, then, chooses a market clearing price and communicates it to tenants. Tenants can now use the price from the operator in their supply function to calculate individual required energy reductions and payments for it. The authors also present mathematical analysis of the efficiency of ColoEDR when adapted to economic demand response, where no fixed energy reduction is required.

Niu et al. [2016] set-up the problem as an optimization problem with the goal of maximizing the social welfare (sum of profits of tenants as well as profit of operator). The incentive mechanism resembles other works in that the operator rewards tenants for power reduction but differ in the sense that they model an MTDC as a cooperative environment. Tenants and operator cooperate to achieve the goal of emergency DR. During an emergency DR event, the operator starts bargaining with tenants the amount of power they need to reduce and the desired compensation for it. They propose a solution based on Nash bargain solution under concurrent bargaining of operator with all tenants. They prove that their solution is max-min fair, Pareto efficient, and maximizes social welfare. The authors extend their work in [Guo et al. 2017] to consider the case of sequential bargaining where the operator bargains with tenants one-by-one in a sequential fashion. They compare



sequential and concurrent bargaining to show that concurrent bargaining always maximizes social welfare while in sequential bargaining it may not be the case and the output depends on the order in which bargaining takes place. They also argue that reverse auction type mechanisms [Chen et al. 2015; Sun et al. 2015; Zhang et al. 2015b] may not be suitable for MTDC since the operator, who is the auctioneer, is self-interested and the output is highly dependent on the auctioneer selection.

Contrary to tenants bidding for energy reduction, Ahmed et al. [2015] proposed Contract-DR, in which the operator determines a set of contracts (energy reduction and reward amount) directly that are pushed to tenants. A tenant may participate by accepting one or none of the contracts. Similarly, Tran et al. [2015b; 2016d] propose incentive mechanisms where an operator starts out by broadcasting a reward rate. Tenants then let the operator know about their energy reduction level (that maximizes the tenant's profit) and the operator updates the announced reward rate based on this knowledge. This continues in an iterative and distributed fashion until convergence. This mechanism works for non-strategic (price-taking) tenants whereas they develop a bidding game to deal with strategic (price-anticipating) tenants. The authors use a similar mechanism in [Tran et al. 2016c] where they consider the case of a multi-tenant data center situated in shared building, that is the entire building isn't dedicated to the data center but has other office spaces as well. This non-data center portion occupies a significant area and shares the power infrastructure of the building. They refer to such buildings as mixed-use buildings (MUB). The main differentiating factor of this work is that the authors additionally consider energy reduction by non-data center space by increasing the set point of HVAC (heating, ventilation and air conditioning). The operator initially provides reward rates along with energy reduction targets and a target deviation penalty to offices, data center tenants, and diesel generator. Based on this information, all participants decide their own energy reduction to maximize their own profit. The operator now updates reward rates and energy reduction targets and the process repeats until convergence. The authors extend their work in [Tran et al. 2016b] where they also consider strategic tenants.

**Geo-distributed data center:** An MTDC operator might own more than one data center and there may be tenants who have their servers in one or more of these data centers. In this case, tenants can migrate their workload on top of powering down idle servers to reduce energy use in a particular data center. Sun et al. [2016a] have proposed BatchEDR to enable emergency DR in geo-distributed MTDC. BatchEDR is an online incentive mechanism. As in the previous work by the authors [Sun et al. 2016b], they consider the case of multiple emergency DR events and model tenants as having delay-tolerant workloads that can be shifted in time and processed later. During an emergency DR time period, the tenants bid with the following information at each data center: workload that can be deferred, a factor to convert workload deferment to energy reduction, and claimed cost of energy reduction. The tenant also bids the total workload that can be deferred across all data centers. Additionally, during the first bid the tenant mentions the total workload deferment throughout the entire time periods DR events occur. The authors develop a Vickrey-Clarke-Groves (VCG) mechanism to determine the energy reduction and rewards for each tenants. Zhang et al. [2016] also extend their incentive mechanism, Truth-DR [Zhang et al. 2015b] for the case of geo-distributed MTDC with similar formulation, solution, and results.

All works in emergency demand response of single or geo-distributed MTDC are summarized and compared in Table 4.

## 5 SUSTAINABILITY

MTDCs have lagged behind operator-owned data centers in their adoption of renewable energy sources and making their data centers green [Greenpeace 2017]. With growing concerns of climate change and push from tenants (like Apple [Apple 2017], Facebook [Facebook 2017] and Google [Google 2017]) committed to become 100% renewable powered, it is increasingly important for

Table 4. Comparison of existing works in emergency demand response

<b>Paper</b>	<b>Objective</b>	<b>Solution</b>	<b>Technique</b>	<b>Result</b>
Ahmed et al. [2015]	Min. operator cost	Contract-DR	Operator sets contracts directly	Cost minimized
Chen et al. [2015]	Min. social cost	ColoEDR	Operator chooses a market clearing price	Close to the optimal social cost
Guo et al. [2017]	Max. social welfare	Nash bargaining solution	Cooperatively bargain between operator and tenants	Bargaining improves social welfare
Niu et al. [2016]	Max. social welfare	Nash bargaining solution	Cooperatively bargain between operator and tenants	Concurrent bargaining maximizes social welfare
Sun et al. [2015]	Max. social cost saving	FairDR	Bidding	Fairness is guaranteed
Sun et al. [2016a]	Min. social cost	BatchEDR	Bidding	32% higher social cost saving
Sun et al. [2016b]	Min. social cost	Online mechanism design	Bidding	Close to offline optimal social welfare
Tran et al. [2015b]	Min. social cost	EPM and SWO	Operator sets reward rate iteratively	Optimal social cost
Tran et al. [2016b]	Min. mixed-use building cost	MECH-SA and MECH-NA	Operator sets reward rate iteratively	Optimal mixed-use building cost
Tran et al. [2016c]	Min. mixed-use building cost	DAMESH	Operator sets reward rate iteratively	Cost is smaller compared to baseline
Tran et al. [2016d]	Min. social cost	EPM and SWO	Operator sets reward rate iteratively	Optimal social cost
Wang et al. [2015]	Min. weighted tenant cost	P-RAA and D-RAA	Operator selects a set of tenants for energy-saving	78% average energy-saving
Wang et al. [2017]	Min. tenant cost	Co-Colo	Bidding	Improved resource utilization/energy efficiency
Zhang et al. [2015b]	Min. social cost	Truth-DR	Bidding	Close to the optimal social cost
Zhang et al. [2016]	Min. social cost	Truth-DR	Bidding	20% to 60% cost saving
Zhao et al. [2016]	Max. tenant energy reduction	TECH	Bidding	20% higher energy reduction

MTDCs to decrease their carbon footprint and become sustainable. MTDC can become sustainable by reducing “brown” energy use and increasing “green” energy use or by reducing their overall energy use by being more energy efficient. However, only using green energy, but being very inefficient in their use isn’t sustainable because we would be wasting energy that could be used by others.

## 5.1 Carbon Footprint Reduction

*5.1.1 Problem description.* Research works in this category try to minimize the carbon footprint of MTDC by taking into account the carbon emission per kilowatt-hour (measured in g/kWh) of the fuel mix being used by the grid to produce electricity. There is, however, a cost associated with being green, for example, an operator would have to give out financial reward to tenants for energy reduction required for carbon emission reduction, or invest into on-site renewable energy sources to decrease reliance on grid electricity. These works that focus on minimizing carbon emission must also take operating budget constraints into account.

*5.1.2 Existing solutions.* Islam et al. [2014] looked into the problem of reducing carbon footprint of an MTDC by reducing energy use whenever possible. They proposed a reverse auction incentive mechanism called GreenColo where the operator rewards tenants for energy use reduction. Tenants submit energy reduction bids (amount of energy reduction and proposed payment for it) and the operator decides winning bids to minimize carbon footprint of the MTDC based on the operator’s long term budget. Carbon emission is calculated from the type of fuel and fuel mix (which changes by time of day) used by the utility that provides electricity to the MTDC. An online algorithm is developed to keep track of the reward as well as minimize the carbon footprint at each time interval such that cumulative reward is less than the operator’s budget. The authors evaluate their algorithm through trace-based simulation and demonstrate that GreenColo decreases carbon emission by 18% while tenants can lower their cost by 25% without the operator increasing its operational expense. This is due to the fact that tenants reduce energy consumption lowering the electricity bill, and the savings are eventually rewarded back to tenants but the operator’s total cost remains the same. Mahmud et al. [2016] investigated geographical load balancing (GLB) to minimize overall carbon footprint in hybrid (operator-owned + multi tenant) data center setting. The focus is on large IT organizations who have servers in their own data centers as well as in leased MTDCs. The authors develop a distributed resource management algorithm called CAGE (Carbon and Cost Aware GEographical Job Scheduling) which can be used by an organization for making GLB decisions based on reduction of carbon emission and cost of electricity while satisfying performance guarantees. Price of electricity generally does not correlate with carbon emission, so the authors introduce a carbon-cost parameter in their problem formulation to capture the trade-off between carbon emission minimization and electricity cost minimization. This parameter can be set by the organization independently, to match their carbon reduction target and budget.

## 5.2 Energy Cost Reduction

*5.2.1 Problem description.* Works in this category focus on reducing the electricity use of an MTDC. These works have primary objective of minimizing the operating cost of an MTDC by cutting the electricity bill. Overall, this translates into a greater goal of reducing the energy consumption of MTDC through coordination, making them more sustainable.

*5.2.2 Existing solutions.* Islam et al. [2015] proposed RECO (REward for COst reduction), an incentive mechanism where the operator offers a reward rate to encourage tenants to reduce their energy use. The objective is to minimize the operating expense of MTDC: electricity cost and the rewards to tenant. The authors identify three challenges in designing such a reward program, 1)

dynamic operating condition of MTDC (changing outside temperature, cooling efficiency, and on-site solar generation), 2) electricity pricing with volume charge and demand charge, and 3) unknown tenant response to reward. Their solution addresses each of these challenges. Time varying PUE calculations based on outside ambient temperature and autoregressive moving average (ARMA) model to predict on-site solar energy generation are used to capture the dynamics of MTDC operation. Peak power demand of an MTDC (a significant factor in electric bill) can only be known at the end of the billing cycle. The authors propose tracking the peak power of MTDC and using it as feedback while setting the reward rate, for example, to offer higher reward if the power consumption in the next time slot is predicted to be greater than the tracked peak power. Similarly, they develop a parameterized response function that is updated online in order to learn tenant's response to reward. Finally, they design an online heuristic algorithm to set the reward rate and validate their solution with a prototype as well as simulations. Zhang et al. [2015a] also focus on minimizing the electricity cost of the MTDC taking into account the non-trivial utility pricing which has volume charge and demand charge for electricity. They design online algorithms to solve the minimization problem through two different approaches, 1) pricing – operator sets a reward rate, tenants let the operator know about their energy reduction interest, operator selects tenants to reduce energy and 2) auction – tenants bid (energy reduction and payment for it) for energy reduction, operator decides the winning bids.

Inefficiencies occur in MTDC due to lack of coordination between the tenants as well as between a tenant and operator. Guo et al. [2015] looked into this problem and designed a reward program to coordinate energy management for improved energy efficiency and sustainability of MTDC. Tenants are rewarded by the operator for energy reduction such that the MTDC electricity bill is minimized. The authors formulate the problem as a convex optimization problem in which the sum of operator cost and tenant cost is minimized such that, tenant's performance SLA is not violated (delay is below a threshold). This optimization problem can be solved easily by the operator in a central fashion. But it requires the operator to be aware of all tenant's cost function (not feasible in an MTDC) and also, the solution would not be scalable. Using alternating direction method of multipliers (ADMM), the authors solved the optimization problem in a decentralized fashion. They argue that ADMM based solution is better than the commonly used approach of dual decomposition in that, it is faster to converge, does not require the objective function to be strictly convex, and does not have the problem of choosing an appropriate step size as in the subgradient update method. The authors extend their work in [Guo et al. 2017] to take into consideration the uncertainties in a tenant's workload as well as the electricity cost from the utility. They also extend the formulation to account for tenants with delay-tolerant workloads that spans multiple time slots. The problem is restated as a stochastic optimization problem, where the objective is to minimize the expected total cost over a long period of time. They develop an online centralized solution based on Lyapunov optimization and also a practical distributed solution based on ADMM based on their previous work.

Building on previous work on tenant coordination in MTDC, Guo et al. [2018] investigate the scenario where MTDCs buy power directly from a wholesale electricity market rather than the local utility to avoid added costs by the retail supplier. They argue that cooperation among tenants can lead to lower electricity cost than acting non-cooperatively. The basic idea is that tenant power consumption when aggregated, is less uncertain due to statistical multiplexing than when considering tenant power consumption individually. This decrease in uncertainty can result in a lower cost in the day-ahead wholesale market. The authors formulate the problem as a cooperative game theory problem to model cooperative electricity procurement where each tenant tries to minimize their own electricity cost. The proposed aggregation and cost allocation method is shown

Table 5. Comparison of existing works in sustainability of MTDC

Paper	Objective	Solution	Technique	Result
Guo and Pan [2015]	Min. social cost	A decentralized algorithm	Operator and tenants iteratively solve the optimization problem	27% average cost saving
Guo et al. [2017]	Min. social cost	Online decentralized algorithm	Operator and tenants iteratively solve the optimization problem	27% average cost saving
Guo et al. [2018]	Min. tenant cost	A cost allocation solution	Tenants cooperate to reduce their aggregate power cost	18.03% average cost saving
Islam et al. [2014]	Min. carbon footprint	GreenColo	Bidding	18% decrease in carbon emission
Islam et al. [2015]	Min. operator cost	RECO	Operator sets reward rate	27% reduction in operator cost
Mahmud and Iyengar [2016]	Min. carbon footprint and cost	CAGE	Distribute workload to geo-distributed data centers	36% decrease in carbon emission
Zhang et al. [2015a]	Min. operator cost	Various online algorithms	1) Operator sets reward rate 2) Bidding	Close to offline optimal cost

to be both fair and stable. A simulation based on Google cluster trace shows a saving of 18.03% for average hourly electricity cost.

All works related to sustainability of MTDC are summarized in Table 5.

## 6 POWER OVERSUBSCRIPTION

In this section we discuss works that aim to enable safe power oversubscription of MTDC. While in DR participation, MTDC reduced power consumption when signaled by the utility, for safe oversubscription, MTDC must reduce power consumption when they reach their power limit to avoid sustained power overload.

### 6.1 Tenant-Side Power Oversubscription

*6.1.1 Problem description.* MTDC typically have a power-based pricing model where they charge per kilowatt of power subscription for the tenant. Tenants can save cost for power subscription if they oversubscribe, that is, put more servers than allowed by the power capacity. Tenant-side power oversubscription happens at the lower levels of the power hierarchy, the rack level or cage (few racks) level, where the power use fluctuations (power dynamic range) and probability of power overloading is greater compared to higher levels in the power hierarchy [Fan et al. 2007; Wu et al. 2016]. If tenants exceed their power draw limit, the circuit breaker may trip leading to a power outage in their server rack, or the operator may penalize tenants for power over-utilization and violating the SLA [TekLinks 2017; Verizon 2014]. How can tenants safely oversubscribe the power capacity they have bought?

*6.1.2 Existing solutions.* There are existing control mechanisms designed for operator-owned data centers that can be used by tenants to enable power oversubscription. The basic building block is server level power capping [Gandhi et al. 2009; Lefurgy et al. 2007, 2008; Zhang and

Hoffmann 2016] through techniques like CPU throttling, DVFS, and admission control to operate a server within a power constraint. Techniques that enable power capping at higher (PDU, cluster) level of the data center [Azimi et al. 2017; Bhattacharya et al. 2013; Fu et al. 2011; Lim et al. 2011; Raghavendra et al. 2008; Wang et al. 2012a; Wu et al. 2016] coordinate server-level power capping in order to minimize performance degradation. Tenants can apply these control methods to ensure that power consumption is always below the subscribed power budget.

An alternative approach to power oversubscription, without affecting server performance or workload, is to use UPS as an energy buffer that can provide the excess power. Govindan et al. [2011] looked into tapping the stored energy in central UPS batteries to handle peak power draw in data centers. This may not be directly applicable to tenants as they have no control over the central UPS in an MTDC. However, distributed smaller UPS at rack level or even server level rather than a central UPS are becoming more popular in data centers. Distributed UPS avoids a central point of failure and can also avoid the AC-DC-AC double conversion loss at the central UPS [Kontorinis et al. 2012]. Previous work [Govindan et al. 2012; Kontorinis et al. 2012; Palasamudram et al. 2012] has looked into utilizing such distributed UPS to oversubscribe power (or equivalently, under provision the power infrastructure). Such techniques can be applied by tenants in MTDC to oversubscribe power and save cost.

Palasamudram et al. [2012] proposed using batteries for tenants to oversubscribe power and reduce power cost. Although this work focuses on content delivery networks (CDN) deployed in multiple MTDC around the world, the technique is applicable to other types of tenants as well. The idea is for tenants to put their own batteries at the server or rack level. Tenants do not use all of the subscribed power all the time. Hence, when they are using less power than the subscribed power, charge the battery and when the power use is high and exceeds the power subscription, discharge the battery to meet the excess demand rather than over drawing from the data center power source. One advantage of this approach is that tenants do not have to cap server power and hence, their performance is not affected. Total energy consumption does not change in such case. However, procuring batteries incurs capital expenditure (CapEx) although it decreases the monthly leasing fee (OpEx) of tenants due to oversubscription. Therefore, the authors look at the total cost (leasing fee + amortized cost of the batteries) and analyze the trade off between CapEx and OpEx.

Palasamudram et al. formulate two optimization problems. First, power supply minimization is for finding the minimum power to subscribe given tenants who already have a battery of certain capacity. Second, total cost minimization is for simultaneously finding the power to subscribe and battery capacity to purchase in order to minimize the total cost (CapEx + OpEx). These optimization problems are characterized as linear programming which can be solved efficiently. Given the power trace of a tenant, we solve the optimization problem in an offline fashion. The proposed techniques are evaluated on Akamai's CDN traces from 22 different cluster location in U.S. over a 25 days period. In their result the authors found that amount of safe power oversubscription is a concave function of battery size and saturates after certain battery capacity. A battery that can last for 5 minute allows 7.5% oversubscription whereas a 40-minute battery allows up to 16% oversubscription and remains flat at 16% as the battery size is increased. This is because, when we oversubscribe excessively, we are discharging the battery more often than charging it, thus there is less time to charge the battery and we run out of charge regardless of its size. The authors also do a total cost analysis of trade-off between the CapEx and OpEx. Cost savings depend on different factors of the battery (such as unit price, technology used, and expected lifetime) but they show that in a typical case, up to 13.9% cost savings is possible by using reasonably sized batteries.

## 6.2 Operator-Side Power Oversubscription


**6.2.1 Problem description.** An MTDC operator would benefit from oversubscribing the power infrastructure to increase power utilization and lower the provisioning cost. For example, an MTDC that operates on average at 40% of the peak capacity has effectively double the provisioning cost per watt than an MTDC that operates on average at 80% power utilization [Barroso et al. 2013]. Leasing out more power capacity to tenants is also more profitable to the operator. But as soon as the power hierarchy is oversubscribed, there is always a chance of power overload. Power capping solutions are proposed to avoid coincidental power peaking of servers in a cluster to prevent power overload [Azimi et al. 2017; Wang et al. 2012a; Wu et al. 2016]. However, the MTDC operator lacks control over tenant servers and there is no mechanism in place to inform tenants about the power overload. Randomly cutting off power to tenants is unacceptable [Islam et al. 2016] as it causes service disruption to tenants and may damage the reputation of the MTDC operator, causing tenants to move to another MTDC. The problem is to find indirect mechanisms for power capping in an MTDC such that the effect to tenants is both controlled and minimized in addition to avoiding power overload situations in an oversubscribed MTDC.

**6.2.2 Existing solutions.** Two solutions have been proposed to address the problem of safe oversubscription in MTDC.


**COOP:** Islam et al. [2016] proposed a market mechanism called COOrdinated Power (COOP) based on supply function bidding. The basic idea is to offer reward to tenants to incentivize them for power reduction during power overload. They consider power oversubscription at both, PDU level and central UPS level. The problem is formulated as minimization of total cost of tenants (due to power reduction) subject to the constraint that all power capping requirements are satisfied. The use of supply function bidding prevents tenants from revealing their private cost functions. This is efficient as tenants only need to bid one parameter in a predetermined supply function to specify the amount of power they are willing to reduce at different reward rates. Tenants are suppliers of power reduction and operator is trying to buy back power from tenants.

The mechanism works as follows. The operator continually monitors the power use of all the tenants. Whenever the aggregate power consumption exceeds the capacity at any oversubscribed level, operator broadcast the emergency to all tenants along with the form of supply function and the maximum power reduction possible (usually the current power usage) for the tenant. Tenants then bid for power reduction depending upon their cost for performance degradation. A key point is that tenants do not know the operator's reward rate at this point. So the tenant's bid reflects its power reduction flexibility, power reduction at different reward rates, rather than actual power to be reduced. After receiving bids from the tenants, the operator chooses the lowest reward rate that is able to satisfy power capping constraint at each level (PDU, UPS) and communicates to tenant. Tenants determine their power reduction based on the reward rate and cap power until notified by the operator to resume normal operation. Based on the power capping time interval, power reduced and the reward rate, operator calculates reward for each tenant.

The authors evaluated COOP on a prototype of 11 servers considering 5 tenants with realistic workload. They considered different oversubscription levels, aggressive (20%), moderate (15%), and conservative (10%). Tenants reduced their power consumption through DVFS technique available on the servers. They evaluated COOP for two cases, *price-taking tenants* (tenants bid without considering the effect of their bid on the operator's reward rate) and *price-anticipating tenants* (tenants bid to maximize profit considering the effect of their bid on the predicted operator's reward rate) and compared it with the optimal case (OPT) in which the operator can directly control tenant servers and knows all their cost functions (which is not practical in real MTDCs). Fig. 7a shows the total performance cost of tenants in which COOP is quite close to the optimal. Reward rates




(a) Performance cost




(b) Reward rate

Fig. 7. Cost and reward comparison of different algorithms at different oversubscription level [Islam et al. 2016].



(a) Price taking



(b) Price anticipating

Fig. 8. Economic benefit of COOP to operator and all five tenants when tenants are price taking and price anticipating [Islam et al. 2016].

offered to tenants at different oversubscription levels is shown in Fig. 7b. Reward rate increases with the increase in level of oversubscription as the amount of power reduction required increases. Moreover, the operator has to offer higher rewards when tenants are intelligent (price anticipating) and can predict how rewards are offered.

The authors argue that COOP is profitable to both tenants and the operator, as shown in Fig. 8. They show that COOP increases profit for the operator as they are able to sell more power and increase profit for tenants due to the rewards they receive. However, when an operator oversubscribes aggressively, the profit actually decreases as a higher reward rate must be offered for large power reduction from tenants during power overload.

**LOCAP:** We proposed a dynamic pricing method called LOCAP (LOCAL price for Power) in [Malla and Christensen 2017] to enable safe power oversubscription. We consider the case of energy-based pricing where the tenants are charged separate for space/power capacity as a subscription fee (dollar per kW) and for energy use as local price (dollar per kWh). In this scheme, tenants pay for power consumption and this incentivizes them to be energy efficient. The key idea to enabling oversubscription is to make the local price dynamic and vary it to reflect total power demand in the data center. During normal operation, the local price can be fixed at a certain level. In case of a power overload, the operator can increase the local price to reduce power demand. The increase in local price must be just enough to keep the power consumption below the power limit.






Fig. 9. Comparison between NOOV, COOP, and LOCAP for monthly profits, monthly energy consumption, and their reaction to power overload [Malla and Christensen 2017].

Table 6. Comparison of existing works in power oversubscription of MTDC

Paper	Objective	Solution	Technique	Result
Islam et al. [2016]	Min. tenant cost	COOP	Operator sets reward rate	Close to optimal performance cost and power reduction
Malla and Christensen [2017]	Max. tenant utility	LOCAP	Dynamic power pricing	34% energy reduction
Palasamudram et al. [2012]	Min. power supply; Min. power cost	Solve linear programming	Use batteries during power overload	40-minute battery allows 16% oversubscription; 13.9% cost savings

We consider data center level power oversubscription and formulate the problem as a utility maximization problem with the constraint that total power use does not exceed the data center critical power capacity. This problem can be solved by the operator in a central fashion but it requires the operator to know the utility function of each individual tenant. By taking a distributed approach and solving the dual problem using the iterative gradient projection method, the operator does not need to know the utility function of each tenant. We end up with a simple and intuitive algorithm to update the local price: when power demand exceeds the capacity, increase the local price (rate of increase is proportional to the amount by which the demand exceeds supply) and vice versa. Meanwhile, tenants (who are self interested) consume power to maximize their own profit. During periods of high power price, it may be more profitable for tenants to reduce power consumption and incur some performance degradation rather than pay the high price for power.

In order to evaluate our proposal, we simulated an MTDC with 3 tenants having 10,000 servers each. Using real world workload traces and identical scenarios, we compared LOCAP with COOP and a case without any oversubscription (NOOV). Fig. 9a and fig. 9b shows profit and total data center energy consumption for the three cases respectively. Again, oversubscription is profitable to operator as well as tenants. LOCAP, as it charges for power use, incentivizes tenants to be energy efficient all the time and hence, lower the total energy consumption of the MTDC as a whole. LOCAP and COOP both are able to cap the total power consumption of MTDC. As shown in Fig. 9c, during a power emergency, COOP actively offers rewards to tenants while LOCAP raises the local price to avoid power overload situation.

All works relating to power oversubscription of MTDC are summarized and compared in Table 6.

## 7 RELATED SURVEYS

Several surveys of power management techniques for data centers have been published over the past years [Beloglazov et al. 2011; Cavdar and Alagoz 2012; Ge et al. 2013; Hammadi and Mhamdi 2014; Jin et al. 2016; Kong and Liu 2014; Liu and Zhu 2010; Mastelic et al. 2014; Mittal 2014; Orgerie et al. 2014; Rahman et al. 2014; Shuja et al. 2016; Zakarya and Gillam 2017], however with only one survey focused on MTDC [Oo et al. 2017]. In this section we briefly review these published surveys and show how our survey adds to the body of knowledge by reviewing additional key works and focusing on an important topic that has been neglected in existing surveys – that of power oversubscription in MTDCs.

Beloglazov et al. [2011] surveyed works in energy and power management of computer systems at different levels. Works ranging from hardware (including dynamic component deactivation and DVFS), operating system, virtualization, and data center levels (including workload consolidation) are described in detail. The authors point out that less instantaneous power consumption may not necessarily mean less energy consumption as same amount of computation can take a longer time at a lower power level. Consuming less energy can reduce the electricity bill while consuming less power can, in addition to reducing the electricity bill, lower the cost/capacity of the required UPS, generators, power distributions units, and cooling infrastructure. Surveys by Mastelic et al. [2014], Orgerie et al. [2014], Mittal [2014], and Shuja et al. [2016] review literature that focus on improving energy efficiency of servers, storage, and networks, building blocks of a data center. Mastelic et al. [2014], in addition to hardware (servers and networks), survey various research work on energy efficiency of data center software (resource management systems and user applications) with a specific focus on cloud computing, while Mittal [2014] notes that integrating different solutions for optimal energy efficiency of the entire data center in a coordinated way is a challenge. A similar survey of works that improves the energy efficiency of large scale computing systems including, clusters, grids, and cloud data centers is conducted by Zakarya et al. [2017]. Similarly, Liu et al. [2010] survey research works having the goal of improving energy efficiency, capping power consumption, and thermal management of data centers. They also survey power management of high performance computing (HPC) systems where the jobs are mostly non-interactive.

There are also research works that focus on greening the data center. They emphasize on integration of on-site renewable energy sources and/or reducing the overall energy use of the data center. Cavdar et al. [2012] and Kong et al. [2014] survey power management techniques that focus on making the data centers green. Rahman et al. [2014] surveys the use of geographical load balancing for power management of data centers. The focus of the survey is on Internet-scale geo-distributed data centers that are connected to smart grids. Such data centers can utilize the dynamic features of smart grids like, distributed generation, integration of various renewable energy sources, and demand response, etc. to achieve different goals including minimize electricity cost, minimize carbon emission, and maximize renewable energy usage. In addition to surveying works in energy efficiency, resource management, and temperature control, Jin et al. [2016] survey different metrics used to monitor and track greenness of data centers.

Data center networks have also received considerable attention. Architectural evolution and energy efficiency of data center networks have been surveyed by Hammadi et al. [2014]. Additionally, Ge et al. [2013] surveys power saving techniques in content delivery networks used for fast and effective delivery of web contents to users around the globe.

A comprehensive survey by You et al. [2017] of more than 30 surveys from 2011 to present on energy efficiency of data centers and cloud computing environments, reflects a myriad of research studies and surveys conducted on the subject. However, multi-tenant type data centers have received less focus from the research community. Previous power management techniques

that are suitable for operator-owned data centers cannot be directly applied as the operator does not have control over a tenant's IT equipment. Survey by Oo et al. [2017] about coordination and power management in multi-tenant data centers is most closely related to our work. The authors mainly focus on market mechanisms that enable DR participation by MTDC and/or make MTDC's more sustainable while leaving out works in power oversubscription of MTDC which we think is an essential part. MTDC power oversubscription can increase utilization of resources and is profitable to the operator.

In contrast to previous work, we survey works on power management of multi-tenant data centers with a focus on power oversubscription. We formulate the problem and describe challenges in solving it. We discuss the reward based and the dynamic pricing market mechanisms that have been proposed in the literature, highlighting their strengths and weakness. Finally, we propose possible future directions and open issues still left to be addressed.

## 8 SUMMARY AND FUTURE DIRECTION

In this survey paper, we have reviewed previous research on power management of MTDC. First, we explained why power management in MTDC is equally, if not more, important than in owner-operated data centers, despite the fact that this area has been largely overlooked by the research community. Then, to have a general understanding of an MTDC, its power hierarchy and economics were described with a detailed discussion on what it means to oversubscribe the power hierarchy and the motivations for oversubscription. Next, we introduced a new taxonomy to classify previous work.

Key findings from our survey are:

- Previous work in power management of MTDC can be divided into three categories, enabling demand response participation, enabling safe power oversubscription, and making MTDCs sustainable. Most of the previous work has focused on demand response participation (and, specifically, emergency demand response participation), with less focus on sustainability and power oversubscription.
- In MTDC, there is a lack of coordination between the operator and the tenants. The recurring theme in previous work is designing an incentive mechanism to enable coordination between the tenants and the operator. Coordination can be achieved through dynamic pricing or dynamic reward incentive mechanisms.
- Most previous work develops a reward mechanism to incentivize tenants for power reduction and makes an implicit assumption that MTDCs employ power-based pricing where a tenant has to pay a flat rate for power subscribed rather than for the power actually used. However, the MTDC industry is moving towards energy-based pricing [Digital Realty 2014] which is already used in some wholesale MTDCs [Equinix 2017]. In energy-based pricing (also referred to as "metered power") a tenant pays for its energy usage and this inherently incentivizes tenants to adopt energy efficient mechanisms (such as powering down idle servers during period of light workload). Dynamic pricing, rather than rewards, is arguably a more natural (varying price to match demand and supply as generally done in practice) control mechanism in such a setting.

In order to have safe power oversubscription of MTDC, we need to better understand the power usage characteristics of tenants. In the future, it would be interesting to investigate the relationship between oversubscription level and probability of overloading, that is, the frequency of power overload events in an oversubscribed MTDC. Markov models have been used to model home electricity consumption [Ardakanian et al. 2011] and teletraffic theory have been used to size transformers in a distribution network for a given loss-of-load probability (LOLP) [Ardakanian

et al. 2012]. Similar analysis could be used to find the appropriate level of power oversubscription for a set of tenants to keep the probability of power overload event below a fixed value. This would help MTDC operators in making capacity provisioning decisions for their data centers.

Another future direction could be to see if power overload events can be predicted. Time series predictive methods such as autoregressive moving average (ARMA) and machine learning techniques such as neural networks and support vector regression have been used for predicting renewable (solar and wind) energy production [Aksanli et al. 2012], home peak load [Singh et al. 2012], and data center network load prediction [Prevost et al. 2011]. Similar predictive models could be used for short-term (minutes, hours) and long-term (days) prediction of tenant power demand. The relationship between the accuracy of such models and forecast length could be studied. Using such prediction models, we could develop predictive control methods for safe power oversubscription rather than reactive control methods which detect power overload events and reduce its duration by power capping. If successful, predictive control method would avoid power overload events preventing any transient power spikes to the power infrastructure of the data center where such power spikes compromise the reliability of the data center. One major challenge to the analysis of tenant power consumption is the availability of MTDC power use data. We are not aware of any publicly available power use data of tenants in MTDCs.

## REFERENCES

- Kishwar Ahmed, Mohammad A Islam, and Shaolei Ren. 2015. A Contract design approach for colocation data center demand response. In *2015 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. 635–640.
- Baris Aksanli, Jagannathan Venkatesh, Liuyi Zhang, and Tajana Rosing. 2012. Utilizing green energy prediction to schedule mixed batch and service jobs in data centers. *SIGOPS Oper. Syst. Rev.* 45, 3 (2012), 53–57.
- Mohamed H Albadi and Ehab F El-Saadany. 2008. A summary of demand response in electricity markets. *Electric Power Systems Research* 78, 11 (2008), 1989 – 1996.
- Apple. 2017. Environmental Responsibility Report. (2017). Retrieved October 12, 2017 from [https://images.apple.com/environment/pdf/Apple\\_Environmental\\_Responsibility\\_Report\\_2017.pdf](https://images.apple.com/environment/pdf/Apple_Environmental_Responsibility_Report_2017.pdf)
- Omid Ardakanian, Srinivasan Keshav, and Catherine Rosenberg. 2011. Markovian models for home electricity consumption. In *Proceedings of the 2Nd ACM SIGCOMM Workshop on Green Networking (GreenNets '11)*. ACM, 31–36.
- Omid Ardakanian, Srinivasan Keshav, and Catherine Rosenberg. 2012. On the use of teletraffic theory in power distribution systems. In *Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet (e-Energy '12)*. ACM, Article 21, 10 pages.
- Reza Azimi, Masoud Badiei, Xin Zhan, Na Li, and Sherief Reda. 2017. Fast decentralized power capping for server clusters. In *2017 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. 181–192.
- Luiz André Barroso, Jimmy Clidaras, and Urs Hölzle. 2013. The datacenter as a computer: An introduction to the design of warehouse-scale machines. *Synthesis lectures on computer architecture* 8, 3 (2013), 1–154.
- Anton Beloglazov, Rajkumar Buyya, Young Choon Lee, and Albert Zomaya. 2011. A taxonomy and survey of energy-efficient data centers and cloud computing systems. *Advances in Computers* 82, 2 (2011), 47–111.
- Arka A Bhattacharya, David Culler, Aman Kansal, Sriram Govindan, and Sriram Sankar. 2013. The need for speed and stability in data center power capping. *Sustainable Computing: Informatics and Systems* 3, 3 (2013), 183 – 193.
- Derya Cavdar and Fatih Alagoz. 2012. A survey of research on greening data centers. In *2012 IEEE Global Communications Conference (GLOBECOM)*. 3237–3242.
- Niangjun Chen, Xiaoqi Ren, Shaolei Ren, and Adam Wierman. 2015. Greening multi-tenant data center demand response. *Performance Evaluation* 91 (2015), 229 – 254. Special Issue: Performance 2015.
- Louis Columbus. 2017. Roundup Of Cloud Computing Forecasts, 2017. (2017). Retrieved October 12, 2017 from <https://www.forbes.com/sites/louiscolombus/2017/04/29/roundup-of-cloud-computing-forecasts-2017/>
- Data Center Map. 2017. Colocation Data Centers. (2017). Retrieved August 28, 2017 from <http://www.datacentermap.com/datacenters.html>
- Howard David, Eugene Gorbato, Ulf R Hanebutte, Rahul Khanna, and Christian Le. 2010. RAPL: Memory power estimation and capping. In *2010 ACM/IEEE International Symposium on Low-Power Electronics and Design (ISLPED)*. 189–194.
- Wei Deng, Fangming Liu, Hai Jin, Bo Li, and Dan Li. 2014. Harnessing renewable energy in cloud datacenters: opportunities and challenges. *IEEE Network* 28, 1 (2014), 48–55.

- Digital Realty. 2014. Colocation and Data Center Worlds Collide. (2014). Retrieved November 7, 2017 from <https://www.digitalrealty.com/blog/colocation-data-center-worlds-collide/>
- Digital Realty. 2017. (2017). Retrieved August 28, 2017 from <https://www.digitalrealty.com/data-centers/chicago/>
- Rachel A Dines. 2011. Build or buy? The economics of data center facilities. (2011).
- EnerNOC. 2013. Ensuring U.S. Grid Security and Reliability: U.S. EPA's Proposed Emergency Backup Generator Rule. (2013). Retrieved October 23, 2017 from <http://www.sustainableplant.com/assets/Sustainable-Plant/PDFs/EPA-Proposed-Backup-Generator-Rule-EnerNOC-Position.pdf>
- Equinix. 2017. Business Suites. (2017). Retrieved November 7, 2017 from <https://www.equinix.com/services/data-centers-colocation/colocation-support-space/business-suites/>
- Facebook. 2017. Sustainability. (2017). Retrieved November 7, 2017 from <https://sustainability.fb.com/>
- Xiaobo Fan, Wolf-Dietrich Weber, and Luiz Andre Barroso. 2007. Power provisioning for a warehouse-sized computer. In *Proceedings of the 34th Annual International Symposium on Computer Architecture (ISCA '07)*. ACM, 13–23.
- Xing Fu, Xiaorui Wang, and Charles Lefurgy. 2011. How much power oversubscription is safe and allowed in data centers. In *Proceedings of the 8th ACM International Conference on Autonomic Computing (ICAC '11)*. ACM, 21–30.
- Anshul Gandhi, Mor Harchol-Balter, Rajarshi Das, Jeffrey O Kephart, and Charles Lefurgy. 2009. Power capping via forced idleness. In *Proceedings of Workshop on Energy-Efficient Design (WEED '09)*.
- Chang Ge, Zhili Sun, and Ning Wang. 2013. A survey of power-saving techniques on data centers and content delivery networks. *IEEE Communications Surveys Tutorials* 15, 3 (2013), 1334–1354.
- Google. 2017. Environmental Report: 2017 progress update. (2017). Retrieved November 7, 2017 from <https://environment.google/projects/environmental-report-2017/>
- Sriram Govindan, Anand Sivasubramaniam, and Bhuvan Urgaonkar. 2011. Benefits and limitations of tapping into stored energy for datacenters. In *Proceedings of the 38th Annual International Symposium on Computer Architecture (ISCA '11)*. ACM, 341–352.
- Sriram Govindan, Di Wang, Anand Sivasubramaniam, and Bhuvan Urgaonkar. 2012. Leveraging stored energy for handling power emergencies in aggressively provisioned datacenters. *SIGARCH Comput. Archit. News* 40, 1 (2012), 75–86.
- Greenpeace. 2017. Clicking Clean: Who is Winning the Race to Build a Green Internet? (2017). Retrieved November 7, 2017 from <http://www.greenpeace.org/international/en/publications/Campaign-reports/Climate-Reports/clicking-clean-2017/>
- Yuanxiong Guo, Hongning Li, and Miao Pan. 2017. Colocation data center demand response using Nash bargaining theory. *IEEE Transactions on Smart Grid* PP, 99 (2017), 1–1.
- Yuanxiong Guo and Miao Pan. 2015. Coordinated energy management for colocation data centers in smart grids. In *2015 IEEE International Conference on Smart Grid Communications (SmartGridComm)*. 840–845.
- Yuanxiong Guo, Miao Pan, and Yanmin Gong. 2018. Aggregation-based colocation datacenter energy management in wholesale markets. *IEEE Transactions on Cloud Computing* PP, 99 (2018), 1–1.
- Yuanxiong Guo, Miao Pan, Yanmin Gong, and Yuguang Fang. 2017. Dynamic multi-tenant coordination for sustainable colocation data centers. *IEEE Transactions on Cloud Computing* PP, 99 (2017), 1–1.
- Ali Hammadi and Lotfi Mhamdi. 2014. A survey on architectures and energy efficiency in data center networks. *Computer Communications* 40 (2014), 1 – 21.
- Mohammad A Islam, Hasan Mahmud, Shaolei Ren, and Xiaorui Wang. 2015. Paying to save: Reducing cost of colocation data center via rewards. In *2015 IEEE 21st International Symposium on High Performance Computer Architecture (HPCA)*. 235–245.
- Mohammad A Islam, Shaolei Ren, and Xiaorui Wang. 2014. GreenColo: A novel incentive mechanism for minimizing carbon footprint in colocation data center. In *International Green Computing Conference*. 1–8.
- Mohammad A Islam, Xiaoqi Ren, Shaolei Ren, Adam Wierman, and Xiaorui Wang. 2016. A market approach for handling power emergencies in multi-tenant data center. In *2016 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. 432–443.
- Xibo Jin, Fa Zhang, Athanasios V Vasilakos, and Zhiyong Liu. 2016. Green data centers: A survey, perspectives, and future directions. *arXiv preprint arXiv:1608.00687* (2016).
- John Karidis, Jose E Moreira, and Jaime Moreno. 2009. True value: Assessing and optimizing the cost of computing at the data center level. In *Proceedings of the 6th ACM Conference on Computing Frontiers (CF '09)*. ACM, 185–192.
- Fanxin Kong and Xue Liu. 2014. A survey on green-energy-aware power management for datacenters. *ACM Comput. Surv.* 47, 2, Article 30 (2014), 30:1–30:38 pages.
- Vasileios Kontorinis, Liuyi Eric Zhang, Baris Aksanli, Jack Sampson, Houman Homayoun, Eddie Pettis, Dean M Tullsen, and Tajana Simunic Rosing. 2012. Managing distributed ups energy for effective power capping in data centers. In *2012 39th Annual International Symposium on Computer Architecture (ISCA)*. 488–499.
- Kien Le, Ozlem Bilgir, Ricardo Bianchini, Margaret Martonosi, and Thu D Nguyen. 2010. Managing the cost, energy consumption, and carbon footprint of internet services. *SIGMETRICS Perform. Eval. Rev.* 38, 1 (2010), 357–358.

- Charles Lefurgy, Xiaorui Wang, and Malcolm Ware. 2007. Server-level power control. In *Fourth International Conference on Autonomic Computing (ICAC'07)*.
- Charles Lefurgy, Xiaorui Wang, and Malcolm Ware. 2008. Power capping: A prelude to power shifting. *Cluster Computing* 11, 2 (2008), 183–195.
- Harold Lim, Aman Kansal, and Jie Liu. 2011. Power budgeting for virtualized data centers. In *2011 USENIX Annual Technical Conference (USENIX ATC'11)*, Vol. 59.
- Minghong Lin, Adam Wierman, Lachlan LH Andrew, and Eno Thereska. 2013. Dynamic right-sizing for power-proportional data centers. *IEEE/ACM Trans. Netw.* 21, 5 (2013), 1378–1391.
- Yongpeng Liu and Hong Zhu. 2010. A survey of the research on power management techniques for high-performance systems. *Software: Practice and Experience* 40, 11 (2010), 943–964.
- Zhenhua Liu, Yuan Chen, Cullen Bash, Adam Wierman, Daniel Gmach, Zhikui Wang, Manish Marwah, and Chris Hyser. 2012. Renewable and cooling aware workload management for sustainable data centers. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '12)*. ACM, 175–186.
- Zhenhua Liu, Iris Liu, Steven Low, and Adam Wierman. 2014. Pricing data center demand response. *SIGMETRICS Perform. Eval. Rev.* 42, 1 (2014), 111–123.
- Zhenhua Liu, Adam Wierman, Yuan Chen, Benjamin Razou, and Niangjun Chen. 2013. Data center demand response: Avoiding the coincident peak via workload shifting and local generation. *Performance Evaluation* 70, 10 (2013), 770 – 791.
- Jianying Luo, Lei Rao, and Xue Liu. 2014. Temporal load balancing with service delay guarantees for data center energy cost optimization. *IEEE Transactions on Parallel and Distributed Systems* 25, 3 (2014), 775–784.
- A Hasan Mahmud and SS Iyengar. 2016. A distributed framework for carbon and cost aware geographical job scheduling in a hybrid data center infrastructure. In *2016 IEEE International Conference on Autonomic Computing (ICAC)*. 75–84.
- Sulav Malla and Ken Christensen. 2017. Reducing power use and enabling oversubscription in multi-tenant data centers using local price. In *2017 IEEE International Conference on Autonomic Computing (ICAC)*. 161–166.
- Katrina Managan. 2014. Demand Response: A Market Overview. (2014). Retrieved October 23, 2017 from <https://pdfs.semanticscholar.org/c64d/a0a7857954f1a5b168c682eb8c7678659077.pdf>
- MarketsandMarkets. 2017. Data Center Colocation Market by Type (Retail and Wholesale), End-user (SMES and Large Enterprise), Industry (BFSI, IT & Telecom, Government & Defense, Healthcare, Research & Academic, Retail, Energy and Manufacturing), and Region - Global Forecast to 2022. (July 2017). Retrieved August 28, 2017 from <http://www.marketsandmarkets.com/>
- Toni Mastelic, Ariel Oleksiak, Holger Claussen, Ivona Brandic, Jean-Marc Pierson, and Athanasios V. Vasilakos. 2014. Cloud computing: Survey on energy efficiency. *ACM Comput. Surv.* 47, 2, Article 33 (Dec. 2014), 36 pages.
- Sparsh Mittal. 2014. Power management techniques for data centers: A survey. *arXiv preprint arXiv:1404.6681* (2014).
- Justin D Moore, Jeffrey S Chase, Parthasarathy Ranganathan, and Ratnesh K Sharma. 2005. Making scheduling “cool”: Temperature-aware workload placement in data centers. In *Proceedings of USENIX Annual Technical Conference*. 61–75.
- MTEMC. 2017. Rate Structure for Business and Industry. (2017). Retrieved October 22, 2017 from <https://www.mtemc.com/Business-Rate-Structure>
- Luyao Niu, Yuanxiong Guo, Hongning Li, and Miao Pan. 2016. A Nash bargaining approach to emergency demand response in colocation data centers. In *2016 IEEE Global Communications Conference (GLOBECOM)*. 1–6.
- NRDC. 2014. Scaling Up Energy Efficiency Across the Data Center Industry: Evaluating Key Drivers and Barriers. (August 2014). Retrieved August 28, 2017 from <https://www.nrdc.org/sites/default/files/data-center-efficiency-assessment-IP.pdf>
- Thant Zin Oo, Nguyen H Tran, Shaolei Ren, and Choong Seon Hong. 2017. *A Survey on Coordinated Power Management in Multi-Tenant Data Centers*. Springer.
- Optimal Networks. 2015. How Much Does Server Colocation Cost? (2015). Retrieved October 22, 2017 from <https://www.optimalnetworks.com/2015/01/24/cost-server-colocation/>
- Anne-Cecile Orgerie, Marcos Dias de Assuncao, and Laurent Lefevre. 2014. A survey on techniques for improving the energy efficiency of large-scale distributed systems. *ACM Comput. Surv.* 46, 4, Article 47 (March 2014), 31 pages.
- Darshan S Palasamudram, Ramesh K Sitaraman, Bhuvan Urgaonkar, and Rahul Urgaonkar. 2012. Using batteries to reduce the power costs of internet-scale distributed networks. In *Proceedings of the Third ACM Symposium on Cloud Computing (SoCC '12)*. ACM, Article 11, 14 pages.
- Ponemon Institute. 2016. 2016 Cost of Data Center Outages. (2016). Retrieved October 31, 2017 from <https://www.ponemon.org/blog/2016-cost-of-data-center-outages>
- John J Prevost, KranthiManoj Nagothu, Brian Kelley, and Mo Jamshidi. 2011. Prediction of cloud data center networks loads using stochastic and neural models. In *2011 6th International Conference on System of Systems Engineering*. 276–281.
- Asfandyar Qureshi, Rick Weber, Hari Balakrishnan, John Guttag, and Bruce Maggs. 2009. Cutting the electric bill for internet-scale systems. In *Proceedings of the ACM SIGCOMM 2009 Conference on Data Communication (SIGCOMM '09)*. ACM, 123–134.

- Ramya Raghavendra, Parthasarathy Ranganathan, Vanish Talwar, Zhikui Wang, and Xiaoyun Zhu. 2008. No “power” struggles: Coordinated multi-level power management for the data center. *SIGARCH Comput. Archit. News* 36, 1 (2008), 48–59.
- Ashikur Rahman, Xue Liu, and Fanxin Kong. 2014. A survey on geographic load balancing based data center power management in the smart grid environment. *IEEE Communications Surveys Tutorials* 16, 1 (2014), 214–233.
- Shaolei Ren. 2017. Managing power capacity as a first-class resource in multitenant data centers. *IEEE Internet Computing* 21, 4, 8–14.
- Shaolei Ren and Mohammad A Islam. 2014. Colocation demand response: Why do I turn off my servers?. In *11th International Conference on Autonomic Computing (ICAC 14)*. USENIX Association, 201–208.
- Barry Rountree, Dong H Ahn, Bronis R De Supinski, David K Lowenthal, and Martin Schulz. 2012. Beyond DVFS: A first look at performance under a hardware-enforced power bound. In *2012 IEEE 26th International Parallel and Distributed Processing Symposium Workshops PhD Forum*. 947–953.
- Arman Shehabi, Sarah Smith, Dale Sartor, Richard Brown, Magnus Herrlin, Jonathan Koomey, Eric Masanet, Nathaniel Horner, Inês Azevedo, and William Lintner. 2016. United states data center energy usage report. (2016).
- Junaid Shuja, Kashif Bilal, Sajjad A Madani, Mazliza Othman, Rajiv Ranjan, Pavan Balaji, and Samee U Khan. 2016. Survey of techniques and architectures for designing energy-efficient data centers. *IEEE Systems Journal* 10, 2 (2016), 507–519.
- Pierluigi Siano. 2014. Demand response and smart grids—A survey. *Renewable and Sustainable Energy Reviews* 30 (2014), 461–478.
- Rayman Preet Singh, Peter Xiang Gao, and Daniel J Lizotte. 2012. On hourly home peak load prediction. In *2012 IEEE Third International Conference on Smart Grid Communications (SmartGridComm)*. 163–168.
- Qihang Sun, Shaolei Ren, Chuan Wu, and Zongpeng Li. 2016a. An online incentive mechanism for emergency demand response in geo-distributed colocation data centers. In *Proceedings of the Seventh International Conference on Future Energy Systems (e-Energy '16)*. ACM, Article 3, 13 pages.
- Qihang Sun, Chuan Wu, Zongpeng Li, and Shaolei Ren. 2016b. Colocation demand response: Joint online mechanisms for individual utility and social welfare maximization. *IEEE Journal on Selected Areas in Communications* 34, 12 (2016), 3978–3992.
- Qihang Sun, Chuan Wu, Shaolei Ren, and Zongpeng Li. 2015. Fair rewarding in colocation data centers: Truthful mechanism for emergency demand response. In *2015 IEEE 23rd International Symposium on Quality of Service (IWQoS)*. 359–368.
- TekLinks. 2017. Colocation Service Level Agreement. (2017). Retrieved October 20, 2017 from <http://www.teklink.com/wp-content/uploads/2017/06/Colocation-SLA-20170621.pdf>
- Nguyen Tran, Chuan Pham, Minh Nguyen, Shaolei Ren, and Choong Seon Hong. 2016b. Incentivizing energy reduction for emergency demand response in multi-tenant mixed-use buildings. *IEEE Transactions on Smart Grid* PP, 99 (2016), 1–1.
- Nguyen H Tran, Cuong T Do, Choong Seon Hong, Shaolei Ren, and Zhu Han. 2015a. Coordinated colocation datacenters for economic demand response. *SIGMETRICS Perform. Eval. Rev.* 43, 3 (2015), 34–37.
- Nguyen H Tran, Cuong T Do, Shaolei Ren, Zhu Han, and Choong Seon Hong. 2015b. Incentive mechanisms for economic and emergency demand responses of colocation datacenters. *IEEE Journal on Selected Areas in Communications* 33, 12 (2015), 2892–2905.
- Nguyen H Tran, Thant Zin Oo, Shaolei Ren, Zhu Han, Eui-Nam Huh, and Choong Seon Hong. 2016a. Reward-to-Reduce: An incentive mechanism for economic demand response of colocation datacenters. *IEEE Journal on Selected Areas in Communications* 34, 12 (2016), 3941–3953.
- Nguyen H Tran, Chuan Pham, Minh NH Nguyen, Shaolei Ren, and Choong Seon Hong. 2016c. Incentive-aligned mechanism for emergency demand response in multi-tenant mixed-use buildings. In *2016 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. 112–117.
- Nguyen H Tran, Chuan Pham, Shaolei Ren, Zhu Han, and Choong Seon Hong. 2016d. Coordinated power reduction in multi-tenant colocation datacenter: An emergency demand response study. In *2016 IEEE International Conference on Communications (ICC)*. 1–6.
- Uptime Institute. 2014. 2014 Data Center Industry Survey. (2014). Retrieved August 28, 2017 from <https://journal.uptimeinstitute.com/2014-data-center-industry-survey/>
- USGBC. 2017. Leadership in Energy and Environmental Design. (2017). Retrieved November 7, 2017 from <https://new.usgbc.org/leed>
- Verizon. 2014. Colocation Services +. (2014). Retrieved October 20, 2017 from [http://www.verizonenterprise.com/external/service\\_guide/reg/cp\\_colo\\_plus\\_colocation\\_services.pdf](http://www.verizonenterprise.com/external/service_guide/reg/cp_colo_plus_colocation_services.pdf)
- Cheng Wang, Bhuvan Urgaonkar, Qian Wang, and George Kesidis. 2014. In *2014 IEEE 22nd International Symposium on Modelling, Analysis Simulation of Computer and Telecommunication Systems*. 305–314.
- Di Wang, Chuangang Ren, Anand Sivasubramaniam, Bhuvan Urgaonkar, and Hosam Fathy. 2012b. Energy storage in datacenters: what, where, and how much? *SIGMETRICS Perform. Eval. Rev.* 40, 1 (2012), 187–198.

- Xiaorui Wang, Ming Chen, Charles Lefurgy, and Tom W Keller. 2012a. SHIP: A scalable hierarchical power control architecture for large-scale data centers. *IEEE Transactions on Parallel and Distributed Systems* 23, 1 (2012), 168–176.
- Youshi Wang, Fa Zhang, and Zhiyong Liu. 2015. *Truthful strategy and resource integration for multi-tenant data center demand response*. Springer International Publishing, 259–270.
- Youshi Wang, Fa Zhang, Shaolei Ren, Fangming Liu, Rui Wang, and Zhiyong Liu. 2017. Energy efficiency in colocation data centers: A joint incentive mechanism approach. In *2017 Seventh International Green and Sustainable Computing Conference (IGSC)*.
- Adam Wierman, Zhenhua Liu, Iris Liu, and Hamed Mohsenian-Rad. 2014. Opportunities and challenges for data center demand response. In *International Green Computing Conference*. 1–10.
- Chia-Ming Wu, Ruay-Shiung Chang, and Hsin-Yu Chan. 2014. A green energy-efficient scheduling algorithm using the DVFS technique for cloud datacenters. *Future Generation Computer Systems* 37, Supplement C (2014), 141 – 147.
- Qiang Wu, Qingyuan Deng, Lakshmi Ganesh, Chang-Hong Hsu, Yun Jin, Sanjeev Kumar, Bin Li, Justin Meza, and Yee Jiun Song. 2016. Dynamo: Facebook’s data center-wide power management system. In *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*. 469–480.
- Huiting Xu, Xi Jin, and Qingxu Deng. 2017. Hierarchical demand response for colocation data centers. In *2017 IEEE International Conference on Smart Computing (SMARTCOMP)*. 1–8.
- Xindong You, Yeli Li, Meilian Zheng, Chuan Zhu, and Lifeng Yu. 2017. A survey and taxonomy of energy efficiency relevant surveys in cloud-related environments. *IEEE Access* 5 (2017), 14066–14078.
- Muhammad Zakarya and Lee Gillam. 2017. Energy efficient computing, clusters, grids and clouds: A taxonomy and survey. *Sustainable Computing: Informatics and Systems* 14, Supplement C (2017), 13 – 33.
- Huazhe Zhang and Henry Hoffmann. 2016. Maximizing performance under a power cap: A comparison of hardware, software, and hybrid techniques. *SIGPLAN Not.* 51, 4 (2016), 545–559.
- Linquan Zhang, Zongpeng Li, Chuan Wu, and Shaolei Ren. 2015a. Online electricity cost saving algorithms for co-location data centers. *IEEE Journal on Selected Areas in Communications* 33, 12 (2015), 2906–2919.
- Linquan Zhang, Shaolei Ren, Chuan Wu, and Zongpeng Li. 2015b. A truthful incentive mechanism for emergency demand response in colocation data centers. In *2015 IEEE Conference on Computer Communications (INFOCOM)*. 2632–2640.
- Linquan Zhang, Shaolei Ren, Chuan Wu, and Zongpeng Li. 2016. A truthful incentive mechanism for emergency demand response in geo-distributed colocation data centers. *ACM Trans. Model. Perform. Eval. Comput. Syst.* 1, 4, Article 18 (2016), 18:1–18:23 pages.
- Yanwei Zhang, Yefu Wang, and Xiaorui Wang. 2013. TESStore: Exploiting thermal and energy storage to cut the electricity bill for datacenter cooling. In *Proceedings of the 8th International Conference on Network and Service Management (CNSM ’12)*. International Federation for Information Processing, 19–27.
- Ziqi Zhao, Fan Wu, Shaolei Ren, Xiaofeng Gao, Guihai Chen, and Yong Cui. 2016. TECH: A thermal-aware and cost efficient mechanism for colocation demand response. In *2016 45th International Conference on Parallel Processing (ICPP)*. 464–473.