Ohm's Law in Data Centers: A Voltage Side Channel for Timing Power Attacks

Mohammad A. Islam
UC Riverside

Shaolei Ren
UC Riverside

ABSTRACT

Maliciously-injected power load, a.k.a. power attack, has recently surfaced as a new egregious attack vector for dangerously compromising the data center availability. This paper focuses on the emerging threat of power attacks in a multi-tenant colocation data center, an important type of data center where multiple tenants house their own servers and share the power distribution system. Concretely, we discover a novel physical side channel — a voltage side channel — which leaks the benign tenants' power usage information at runtime and helps an attacker precisely time its power attacks. The key idea we exploit is that, due to the Ohm's Law, the high-frequency switching operation (40 ∼ 100kHz) of the power factor correction circuit universally built in today's server power supply units creates voltage ripples in the data center power lines. Importantly, without overlapping the grid voltage in the frequency domain, the voltage ripple signals can be easily sensed by the attacker to track the benign tenants' runtime power usage and precisely time its power attacks. We evaluate the timing accuracy of the voltage side channel in a real data center prototype, demonstrating that the attacker can extract benign tenants' power pattern with a great accuracy (correlation coefficient = 0.90+) and utilize 64% of all the attack opportunities without launching attacks randomly or consecutively. Finally, we highlight a few possible defense strategies and extend our study to more complex three-phase power distribution systems used in large multi-tenant data centers.

CCS CONCEPTS

• **Security and privacy → Side-channel analysis and countermeasures**;

KEYWORDS

Data center; power attack; voltage side channel

ACM Reference Format:
Mohammad A. Islam and Shaolei Ren. 2018. Ohm's Law in Data Centers: A Voltage Side Channel for Timing Power Attacks. In *2018 ACM SIGSAC Conference on Computer and Communications Security (CCS '18), October 15–19, 2018, Toronto, ON, Canada.* ACM, New York, NY, USA, 17 pages. https://doi.org/10.1145/3243734.3243744

1 INTRODUCTION

In the age of cloud computing and Internet of Things, data centers have experienced an exponential growth at all scales and undeniably become mission-critical infrastructures without which our society cannot function. In fact, even a single data center outage can egregiously affect our day-to-day life. For example, an outage in Delta Airlines' data center in 2016 stranded tens of thousands of passengers in transit, costing more than 150 million U.S. dollars [1]. Moreover, a recent survey shows that unplanned data center-wide outages caused by malicious attacks have increased by 11 times from 2010 to 2016 [2]. Thus, securing data centers against attacks has been of paramount importance.

While data center's cyber security has been extensively investigated [3–5], a much less studied security aspect — power infrastructure security — has also emerged as an equally, if not more, important concern. Even worse, besides being afflicted with random system failures, data center power infrastructures are also increasingly becoming a target for malicious attacks due to the criticality of their hosted services [2, 6]. Concretely, recent studies [7–11] have found and successfully demonstrated that an attacker can inject malicious power loads (referred to as *power attacks*) to overload the data center power infrastructure capacity, thus creating more frequent data center outages. Such power attacks are achieved by increasing the attacker's own server power usage [8, 11] and/or sending more workloads to the target data center [9, 10].

The primary reason for data centers' vulnerability to power attacks stems from the common practice of power capacity oversubscription. Data center power infrastructures are very expensive (and sometimes impossible because of local grid capacity or other constraints) to build to accommodate the growing demand, costing 10 ∼ 25 dollars per watt of power capacity delivered to servers and taking up 25 ∼ 60% of an operator's total cost of ownership over a 15-year lifespan [12–15]. As a consequence, to maximize utilization of existing power infrastructures, data centers (even Facebook and Google data centers [14, 16]) commonly oversubscribe their power capacity by housing more servers than can be supported. The current industry average is to oversubscribe the infrastructure by 120% (i.e., provisioning 100kW power capacity to servers whose total power can reach 120kW) [11, 17], and recent research [9, 13, 14] has suggested even more aggressive oversubscription. The rationale for oversubscription is statistical multiplexing: not all servers peak their power usage simultaneously. Additionally, various techniques (e.g., throttling CPU and halting services [14, 15, 18, 19]) have been proposed to handle the very rare, albeit possible, power capacity overload resulting from oversubscription.

Nonetheless, power attacks, especially maliciously timed attacks [7, 8, 11], can alter the servers' total power usage and create frequent power capacity overloads. Despite safeguards (e.g., infrastructure redundancy), power attacks at best invoke power capping more often

CCS '18, October 15–19, 2018, Toronto, ON, Canada

ACM ISBN 978-1-4503-5693-0/18/10…$15.00
https://doi.org/10.1145/3243734.3243744

than otherwise, significantly degrading application performances (due to, e.g., CPU throttling [9, 14]). More importantly, they significantly compromise the data center availability and create more frequent outages, which can lead to catastrophic consequences (see Delta Airlines' example [1]).

In this paper, we focus on the emerging threat of power attacks in a multi-tenant colocation data center (also called colocation or multi-tenant data center), an important but less studied type of data centers [20]. A multi-tenant data center is a shared data center facility, in which multiple companies/organizations (each as a *tenant*) houses their own physical servers and the data center operator is responsible for providing reliable power and cooling to tenants' servers. Even large companies, like Google and Apple [21, 22], lease multi-tenant data center capacities to complement their own data centers.

Compared to an owner-operated data center whose operator can perform power capping/throttling to mitigate power attacks, a multi-tenant data center is more vulnerable to power attacks, because the data center operator has no control over tenants' power usage. Alternatively, the operator of a multi-tenant data center sets contractual constraints: each tenant can continuously use a certain fraction (usually 80%) of its subscribed power capacity, but can only use its full subscribed power capacity on an occasional basis; non-compliance can result in forcible power cuts [11, 23]. Therefore, to launch successful power attacks while meeting the contractual constraint in a multi-tenant data center, a malicious tenant (attacker) must precisely time its power attacks: it needs to increase its server power to the full capacity only at moments when the benign tenants are also using a high power [8, 11]. Nonetheless, *a key challenge for the attacker to precisely time its power attacks is that it does not know benign tenants' power usage at runtime.* Importantly, attack opportunities (i.e., benign tenants' high power usage moments) are highly intermittent, making random attacks unlikely to be successful (Fig. 16 in Section 5.2).

In order to achieve a good timing of power attacks, we discover a novel physical side channel — voltage side channel — which leaks information about benign tenants' power usage at runtime. Concretely, we find that a power factor correction (PFC) circuit is almost *universally* built in today's server power supply units to shape server's current draw following the sinusoidal voltage signal wave and hence improve the power factor (i.e., reducing reactive power that performs no real work) [24]. The PFC circuit includes a pulse-width modulation (PWM) that switches on and off at a high frequency (40 ~ 100kHz) to regulate the current. This switching operation creates high-frequency current ripples which, due to the Ohm's Law (i.e., voltage is proportional to current given a resistance) [25], generate voltage ripples along the power line from which the server draws current. Importantly, the high-frequency voltage ripple becomes more prominent as a server consumes more power and can be transmitted over the data center power line network without interferences from the nominal grid voltage frequency (50/60Hz). As a consequence, the attacker can easily sense its supplied voltage signal and extract benign tenants' power usage information from the voltage ripples.

We build a prototype that represents an edge multi-tenant data center [26] to demonstrate the effectiveness of our discovered voltage side channel in terms of timing attacks. Our results show even though the attacker restricts itself from launching continuous attacks to meet the data center operator's contractual limit, it can still successfully utilize more than 64% of the available attack opportunities with a precision rate of 50%. If attacks can be launched consecutively, the attacker can even detect 80+% of attack opportunities. Importantly, the attacker's total cost is just a small fraction (3% ~ 16% in our study) of the resulting financial loss. Next, we extend our study to a three-phase power distribution system used in large multi-tenant data centers. Finally, we highlight a few defense strategies (including direct current power distribution, jamming signals, power infrastructure resilience, and attacker identification) and discuss their limitations in practice.

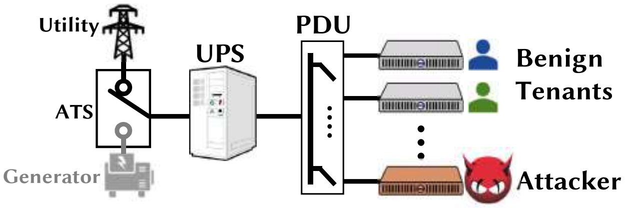

Figure 1: Data center power infrastructure with an attacker.

2 PRELIMINARIES ON POWER ATTACKS

In this section, we provide preliminaries on power attacks, highlighting the importance of multi-tenant data center, the vulnerability and impact of power attacks, and limitations of the prior work.

2.1 Overview of Multi-Tenant Data Centers

2.1.1 Importance of multi-tenant data centers. Multi-tenant colocation data centers, also commonly called *multi-tenant data centers* or *colocations*, are a critical segment of the data center industry, accounting for as much as *five times* the energy consumption by Google-type owner-operated data centers combined altogether [20].

A multi-tenant data center significantly differs from a multi-tenant cloud: in a multi-tenant cloud (e.g., Amazon), the cloud operator owns the physical servers while renting out virtualized resources (e.g., virtual machines) to cloud users; in a multi-tenant data center, the data center operator only owns the data center facility and physical power/cooling infrastructures, whereas tenants manage their own physical servers in shared spaces.

There are more than 2,000 large multi-tenant data centers in the U.S. alone, serving almost all industry sectors that even include large IT companies (e.g., Apple, which houses 25% of its servers in leased multi-tenant data centers) [22, 27]. Importantly, the multi-tenant data center industry is experiencing a double-digit growth to meet the surging demand [28].

Moreover, many emerging Internet of Things workloads, such as augmented reality and assisted driving, are hosted in geo-distributed *edge* multi-tenant data centers in proximity of the data sources for latency minimization. For example, Vapor IO, a data center operator, plans to build thousands of edge multi-tenant data centers in wireless towers [29, 30].

2.1.2 Data center power infrastructure. Typically, data centers employ a tiered power infrastructure as illustrated in Fig. 1. An uninterrupted power supply (UPS) takes the grid voltage as input and outputs voltage to the lower-tier power distribution unit (PDU). The PDU acts as a local power distribution hub and delivers power to server/racks. Each infrastructure has a power capacity protected by a circuit breaker. An automatic transfer switch (ATS) will switch to the backup generator (if any) during grid outages.

The power infrastructure shown in Fig. 1 represents an edge multi-tenant data center where the total power capacity is small (usually in the order of 10+kW or less) and each tenant houses a few servers in a shared server rack. In Section 6 and Appendix H, we also show (three-phase) power infrastructures used in large multi-tenant data centers where an individual tenant houses at least one dedicated server rack and the data center operator oversubscribes its more expensive upper-level PDUs each with 40 ∼ 200kW capacity.

2.2 Vulnerability and Impact of Power Attacks

As stated in Section 1, the common practice of power capacity oversubscription improves the utilization of power infrastructures, but meanwhile also leaves data centers vulnerable to malicious power attacks that result in more frequent and costly power outages.

2.2.1 Current safeguards. First of all, a data center operator leverages infrastructure redundancy to handle random system failures [31]. Depending on the level of redundancy, data centers are classified into four tiers, from Tier-I that has no redundancy to Tier-IV that duplicates all power/cooling systems (so-called "2N" redundancy) [31, 32]. While a power attack may not lead to an actual power outage due to infrastructure redundancies, such redundancy protection is compromised due to malicious power loads, exposing the impacted data center in a very dangerous situation. For instance, with power attacks, the outage risk for a fully-redundant Tier-IV data center increases by 280+ times than otherwise [11, 31].

Moreover, since multi-tenant data center operators cannot control or arbitrarily cut tenants' power usage (unless a tenant is found in violation of the contract), they typically impose contractual constraints on each tenant's continuous power usage (limited to 80% of a tenant's subscribed power capacity), while only allowing tenants to use up their full power capacity occasionally [11, 23]. By doing so, the tenants' aggregate power usage can stay below the actual power infrastructure capacity at almost all times. Nonetheless, these safeguards are highly vulnerable to well-*timed* power attacks that are launched at moments when tenants' power usage is also high (see Fig. 13 for illustration) [8, 11].

In addition, tenants themselves may employ software-based fault tolerance to withstand power outages. Nonetheless, power outages can induce correlated server failures that are challenging to survive through [33]. For example, even a power outage in a single facility cost Delta Airlines over 150 million U.S. dollars [1].

The above discussion highlights that, despite several safeguards, multi-tenant data centers are still highly vulnerable to (well-timed) power attacks [8, 11].

2.2.2 Cost impact of power attacks. While not every power attack can lead to an actual outage, power attacks result in more frequent capacity overloads and hence significantly compromise the data center availability over a long term. For example, the outage risk for a fully-redundant Tier-IV data center increases by 280+ times than otherwise [11, 31]. Based on a conservative estimate [8], even for a medium-size 1MW data center experiencing power attacks for only 3.5% of the time, a total financial loss of 3.5 ∼ 15.6 million U.S. dollars can be incurred per year. The financial loss is incurred not only by tenants which experience service outages, but also by the data center operator which loses its capital expense in strengthening the infrastructure resilience (let alone the reputation damage and high customer churn rate).

More importantly, the attacker only needs to spend a tiny fraction (as low as 3%) of the total loss, thus providing strong motivations for malicious tenants (e.g., organized crime groups that try to bring down services and create societal chaos, the victim data center's competitor, etc.) [8]. Interested readers are referred to [8, 11] for a detailed cost analysis of power attacks.

2.3 Recent Works on Timing Power Attacks

In a multi-tenant data center, a key challenge for an attacker is that the actual attack opportunity lasts intermittently (Fig. 13 in Section 5.2). For timing power attacks in a multi-tenant data center, the prior research has considered a thermal side channel (due to the heat recirculation resulting from servers' heat) [11] and an acoustic side channel (due to noise propagation from servers' cooling fan noise) [8]. Nonetheless, as confirmed by our discussion with a large data center operator, they suffer from the following limitations.

First, both the thermal and acoustic side channels utilize *air* as the medium. Hence, they have only a limited range (e.g., 5 ∼ 10 meters) and are highly sensitive to disturbances (e.g., supply cold air temperature and human voice disturbances). Moreover, because it takes time (1 minute or even longer) for server heat to reach the attacker's temperature sensor and for server's cooling fans to react to server power changes, these side channels cannot provide real-time information about benign tenants' power usage. In addition, exploiting a thermal side channel requires an accurate modeling of heat recirculation, whereas the acoustic side channel needs complex signal processing techniques to mitigate near-far effects (i.e., the attacker's received noise level is dominate by its neighbors) [8]. Last but not least, the thermal side channel requires a raised-floor data center layout without heat containment, whereas the acoustic side channel requires servers have conventional fan speed controls.

In sharp contrast, a distinguishing feature of our discovered voltage side channel (Section 4.3) is that it is insensitive to external disturbances (because of the *wired* power line transmission) and can carry benign tenants' power usage information throughout the power network. The voltage side channel also provides real-time information about benign tenants' power usage (with a delay of 1 second for frequency analysis). More importantly, the voltage side channel is based on the high-frequency voltage ripples generated by PFC circuits that are universally built in servers' power supply units, and can be exploited without any specific models about the data center power network. Finally, while the settings for our experiments and [8, 11] are different, our results show that given 10% attack time, the voltage side channel can achieve 80+% true positive for detecting attack opportunities (Fig. 17(a)) whereas [8, 11] only achieve around or below 50%. This translates into ∼2x successful attacks by using our voltage side channel. Therefore, our voltage side channel presents a more significant threat in real systems.

3 THREAT MODEL

As illustrated in Fig. 1, we consider a malicious tenant (i.e., attacker) that houses its own physical servers in a multi-tenant data center, sharing the power infrastructure with benign tenants.

Attacker's capability. The attacker subscribes a fixed amount of power capacity from the data center operator. It behaves normally as benign tenants, except for its malicious intention to overload the shared power infrastructure and create more power outages. Thus, for stealthiness, the attacker only occasionally uses power up to its capacity, which is allowed by the operator's power usage contract [23]. Physically tampering with the shared power infrastructure or placing explosive devices can be easily found/detected and is orthogonal to our work.

The attacker launches power attacks by running power-hungry applications (e.g., computation to maximize CPU utilization) at moments when the aggregate power usage of benign tenants is sufficiently high. Note that the attacker may also remotely send additional requests to benign tenants' servers during power attacks if benign tenants offer public services (e.g., video streaming). This complementary attack method can further increase the aggregate server power consumption. In this paper, we focus on attacks by using the attacker's own server power as in [8, 11].

To exploit a voltage side channel for timing power attacks, the attacker acquires the supplied voltage by placing an analog-to-digital converter (ADC) circuit inside the power supply unit (Fig. 4) in one of its servers.[1] The ADC samples the incoming continuous-time voltage signals at a rate of 200kHz or above, and the sampled voltage signals are stored for further processing (Section 4). Note that in a multi-tenant data center, the attacker owns its physical servers, instead of renting them from the data center operator. Furthermore, while a multi-tenant data center has a more rigorous inspection for tenants than a public cloud platform, the data center operator will not disassemble tenant servers' power supply units during the routine inspection due to intrusiveness. Thus, a coin-size or even smaller voltage ADC can be easily placed inside the power supply unit before the attacker moves its servers into the target multi-tenant data center. In modern power supply units, a high-speed voltage ADC is already in place as part of the PFC design, and in this case, the attacker can simply read the ADC's output without placing an additional ADC circuit.

Successful attack. Power attacks compromise the data center availability over a long term. Thus, we consider a power attack *successful* as long as the combined power usage of the attacker and benign tenants continuously exceeds the power infrastructure capacity for at least L minutes (e.g., $L = 5$ is enough to trip a circuit [8, 34]), even though an actual outage does not occur due to infrastructure redundancy. Instead of targeting a specific tenant, power attacks compromise the availability of shared power infrastructures and hence significantly affect the normal operation of both data center operator and benign tenants.

Other threat models for power attacks. Next, we highlight the differences between our threat model and other relevant models for power attacks.

• *Power attacks in public clouds.* Some studies [7, 9, 10] propose to use malicious virtual machines (VMs) to create power overloads in public clouds like Amazon. For tripping the circuit breaker and successful attacks, the attacker needs to launch a large number of VMs co-residing in the same PDU. Nonetheless, this is nontrivial and can be difficult to accomplish in practice, because the cloud operator frequently randomizes its VM placement (to prevent co-residency side channel attacks [5, 35]). In addition, the cloud operator has numerous knobs (e.g., CPU scaling) to throttle the VM power consumption for defending its power infrastructure against a power attack. More recently, [36] considers a related attack model but aims at using VMs to generate excessive heat for overloading the cooling capacity.

In contrast, our model focuses on a multi-tenant data center where an attacker can house its own *physical* servers to inject large power loads, tripping the circuit breaker of a shared PDU more easily. The data center operator, as discussed in Section 2.2.1, cannot control or forcibly cut a tenant's power usage unless a tenant violates its power contract.

Compared to using VMs for power attacks in a public cloud [9, 10], an attacker in a multi-tenant data center can incur more costs (e.g., for purchasing servers). At the same time, however, power attacks in our model are also more devastating due to the attacker's capability of injecting large power loads on a single PDU. Importantly, in our model, the attacker's total cost is just a small fraction (3% ∼ 16% in Section 5.2) of the resulting financial loss. Moreover, while VMs can be launched remotely without revealing the attacker's identity [9, 10], it is also difficult to identify and/or prosecute the attacker in our attack scenario, because: (1) data center outages are caused by the operator's capacity oversubscription as well as the aggregate high power of multiple tenants; and (2) the attacker does not violate any contractual constraints. Even though its servers are detected, the attacker's loss is minimum (e.g., only a few servers) because it likely uses fake/counterfeit identities when moving into the target data center. Finally, we focus on precise timing of power attacks for stealthiness, while the crucial timing issue is neglected in [9, 10].

• *Power attacks in multi-tenant data centers.* Our model builds upon those considered in two recent studies [8, 11]. In these studies, the attacker needs to install temperature sensors and microphones in order to exploit a thermal side channel [11] and an acoustic side channel [8], respectively, for timing power attacks. Both thermal sensors and microphones are exposed to the outside of the servers and hence may be detected more easily. In contrast, our model is more stealthy as the attacker only places a small ADC circuit (if not available yet) *inside* its server's power supply unit, without exposing any hardware to the outside. More comparisons (e.g., practical limitations and timing accuracy) are provided in Section 2.3.

4 EXPLOITING A VOLTAGE SIDE CHANNEL

The key novelty of our work is to exploit a voltage side channel to track benign tenants' aggregate power usage at runtime for timing power attacks in a multi-tenant data center.

Concretely, we discover that the PFC circuit inside each server's power supply unit is controlled by a switch to regulate the server's current draw for improving power factor. Because of the Ohm's

[1] ADC circuits often operate over a low voltage range (e.g., 5V) and hence, a voltage divider may be necessary to scale down the incoming voltage to an appropriate range.

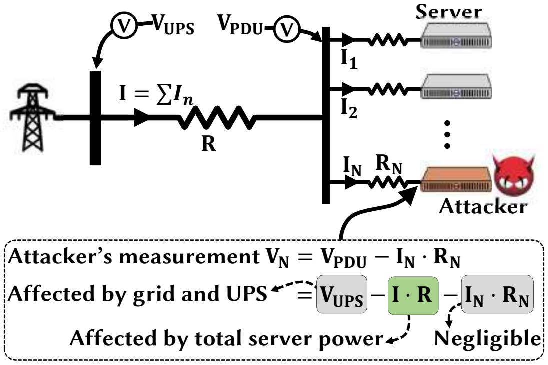
V_{UPS}
V_{PDU}
Server
I_1
I_2
$I = \sum I_n$
R
I_N R_N
Attacker
Attacker's measurement $V_N = V_{PDU} - I_N \cdot R_N$
Affected by grid and UPS $= V_{UPS} - I \cdot R - I_N \cdot R_N$
Affected by total server power
Negligible

Figure 2: Circuit of data center power distribution.

Law, this design creates high-frequency voltage ripples which, without interference from the nominal 50/60Hz frequency of the gird voltage, exist along the power lines supplying voltage to servers. Thus, by sensing the supplied voltage and extracting the frequency components associated with the ripples, the attacker can track benign tenants' power usage and launch well-timed power attacks.

4.1 Overview of the Power Network

Before presenting our voltage side channel, we show in Fig. 2 an overview of the equivalent electrical circuit of a data center power network, where one PDU delivers power to N servers. For better understanding, we focus on a single-phase system that each serves a few tens of servers and best represents an edge multi-tenant data center (hosting workloads such as augmented reality and assisted driving) [29, 37]. In Section 6, we will extend to more complex three-phase systems used in large multi-tenant data centers.

As shown in Fig. 2, the PDU distributes alternating current (AC) to servers using a parallel circuit. We denote the UPS output voltage and the PDU voltage by V_{UPS} and V_{PDU}, respectively. The power line connecting the UPS to the PDU has a resistance of R, and the total current flowing from UPS to PDU is denoted by $I = \sum_{n=1}^{N} I_n$, where I_n is the current of server n. Without loss of generality, we let server N be the one with attacker's ADC circuit, while the attacker can own multiple other servers. Thus, the voltage measured by the attacker is denoted by V_N.

Constraint on current measurement. Power is the product of voltage and current, and servers operate at a relatively stable voltage. Thus, had the attacker been able to sense the total current $I = \sum_{n=1}^{N} I_n$, it would know the aggregate power usage of all tenants and easily time its power attacks. Due to the power line constraint, however, the attacker can only measure the current flowing into its own servers.

Line voltage drop. We observe that the voltage supplied to each individual server is affected by all the servers. Concretely, the current flowing from the UPS to PDU results in a voltage drop ΔV along the power line. The phenomenon of voltage drop is also common in our daily life, e.g., dimming of a light when starting a power-consuming appliance in the same house. Then, following the Ohm's Law, the voltage measured by the attacker is expressed as $V_N = V_{UPS} - I \cdot R - I_N \cdot R_N \approx V_{UPS} - I \cdot R = V_{PDU}$, which can be rewritten as

$$V_N = V_{PDU} = V_{UPS} - R \cdot \sum_{n=1}^{N} I_n, \tag{1}$$

where, for better presentation, we replace the approximation with equality given the fact that the voltage drop $I_N \cdot R_N$ between the PDU and attacker's server is negligible due to the much smaller current I_N compared to $I = \sum_{n=1}^{N} I_n$ and small line resistance R_N. Even when $I_N \cdot R_N$ is non-negligible, the attacker can lower its server's power (i.e., decrease I_N) to make $I_N \cdot R_N$ sufficiently small.

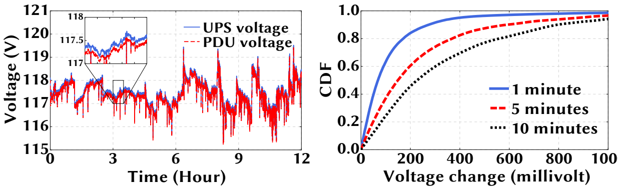

(a) (b)

Figure 3: (a) 12-hour voltage traces at the UPS (grid) and PDU. (b) Probability of temporal variation of the UPS voltage.

4.2 ΔV-based attack

We now present an intuitive strategy — ΔV-based attack — that times power attacks directly based on the attacker's voltage measurement V_N. Importantly, we will show that this seemingly effective strategy results in a rather poor timing of power attacks.

The tenants' aggregate power usage is proportional to the total current $I = \sum_{n=1}^{N}$ and hence also to the voltage drop $\Delta V = |I \cdot R| = |V_{UPS} - V_N|$ between the UPS and PDU, where $| \cdot |$ denotes the absolute value operation and the resistance R is a constant (unknown to the attacker) due to the well-conditioned data center temperature [23].

One may think that V_{UPS} is equal to the nominal voltage $V_{Nominal}$ (e.g., 120V in North America) since it is the UPS output [38]. Consequently, the attacker can simply check its own voltage measurement V_N to time power attacks: a low V_N means a high voltage drop ΔV between the UPS and the PDU, which indicates a high aggregate power usage and hence a good opportunity for power attacks. We refer to this timing strategy as ΔV-based attack.

Nonetheless, the voltage V_{UPS} output by the UPS can vary considerably over time, e.g., up and down by as much as 5V (Fig. 3(a)). The reason is that even state-of-the-art UPS can only regulate its output voltage within 3% of its nominal voltage [38]. The large temporal variation in V_{UPS} is also driven by external factors, such as the grid generator and other loads sharing the same grid. More importantly, the attacker cannot measure V_{UPS} to calculate $\Delta V = |I \cdot R| = |V_{UPS} - V_N|$ because it cannot place its ADC circuit at the output of the UPS which is owned by the data center operator. On the other hand, compared to the V_{UPS} variation, the variation in voltage drop $\Delta V = |I \cdot R|$ caused by tenants' power usage is much smaller (in the order of a few millivolts) because of the small line resistance.

In Fig. 3(a), we show a 12-hour trace of the voltage output by our CyberPower UPS and PDU voltage supplied to servers. The voltage drop between the UPS an PDU is negligible compared to the UPS

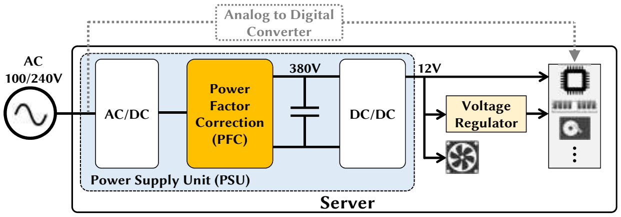

Figure 4: A server with an AC power supply unit [12]. An attacker uses an analog-to-digital converter to acquire the voltage signal.

voltage variation itself. In Fig. 3(b), we show the cumulative distribution function of the UPS output voltage at different timescales, demonstrating that the UPS output voltage can vary much more significantly than the line voltage drop due to server load.

To conclude, the change of V_N is predominantly driven by the variation in the UPS voltage V_{UPS}, rather than the actual line voltage drop $\Delta V = |I \cdot R|$ caused by the tenants' power usage. Thus, *without knowing time-varying V_{UPS}, the ΔV-based strategy cannot precisely time power attacks* (Fig. 16 in Section 5.2).

4.3 Exploiting High-Frequency Voltage Ripples

Given the ineffectiveness of the ΔV-based attack, we present our key idea: the PFC circuit inside each server's power supply unit generates high-frequency voltage ripples that have a strong correlation with the servers' power, which can reveal the aggregate power usage information at runtime. Next, we will first show the root cause of why the PFC generates high-frequency voltage ripples, and then validate the ripples through experiments.

4.3.1 Overview of server's power supply unit. We first provide an overview of server's power supply unit to facilitate the readers' understanding. All the internal components of a server/computer, such as CPU, run on DC power at a regulated voltage (usually 12V), provided by an internal power supply unit. Fig. 4 shows a block diagram of a server.

In the first step, the sinusoidal AC voltage supplied by the PDU is passed through an AC to DC full-bridge rectifier which inverts the negative part of the sine wave and outputs a pulsating DC (half-sine waves). Then, a power factor correction (PFC) circuit outputs an elevated voltage at 380V which is then fed to a DC to DC converter to lower it to 12V supplied to server's internal components. An important concept is power factor, which is a scalar value between 0 and 1 measuring the fraction of total delivered power that is actually used. The power factor goes down when the voltage or current becomes non-sinusoidal, which creates power waste and other detrimental effects [24]. In Appendix A, we show the implications of not using a PFC circuit. Thus, to improve server's power factor, a PFC circuit is required and also mandated by international regulations [39, 40].

4.3.2 Voltage ripples generated by PFC circuit. The purpose of the PFC circuit is to improve power factor by shaping the current drawn by the power supply unit to match the sinusoidal source voltage.[2] The working principle is to draw more current when the input voltage (pulsating AC at the rectifier output) is high and draw less current when the input voltage is low. Fig. 5 shows the basic block diagram of the most commonly-used boost-type PFC with an inductor, a diode, a switch (MOSFET), and the pulse width modulation (PWM) control circuit [24, 41, 42]. The PWM control circuit repeatedly closes and opens the switch at a high frequency to control the current through the inductor to shape it like a sine wave while also maintaining a stable DC voltage at the output. The current wave shapes of the inductor controlled by a server's PFC circuit are illustrated in Fig. 6(a).

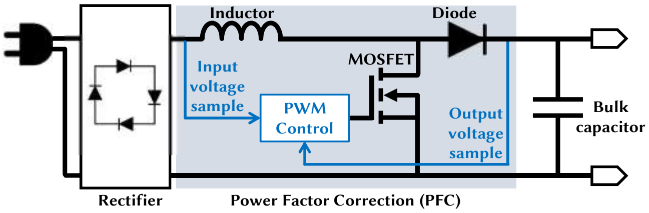

Figure 5: Building blocks of PFC circuit in server's power supply unit.

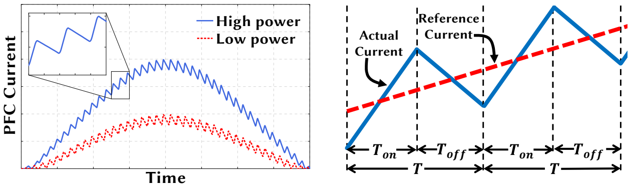

(a) (b)

Figure 6: (a) Wave shape of PFC current at different power levels. (b) Current ripples from the PFC switching.

A prominent side effect of the PWM circuit's rapid switching is the rapid change in the current (i.e., high-frequency ripple) drawn from the source. Hence, the PFC circuit in the power supply unit creates high-frequency current ripples flowing through the power line between the UPS and PDU, which in turn result in voltage ripples along the line due to the Ohm's Law.

Importantly, a key observation is that the voltage ripples are at a much higher frequency (40 ~ 100kHz) than the 50/60Hz nominal grid frequency as well as the UPS output voltage frequency. Thus, the voltage ripple signal and UPS output voltage V_{UPS} signal are orthogonal in the frequency domain. In fact, this is also the fundamental principle for power line communications that leverage power networks as the transmission medium (e.g., recent studies [43] have proposed to install special transmitters and leverage data center power lines to send control command signals for server network management).

In summary, while the PFC circuit is mandated for improving the power factor [39], *its usage of PWM-based switching design creates high-frequency ripple voltage signal that is transmitted over the data center power lines without interference from the UPS output voltage.*

[2] The voltage signal coming from the PDU is not perfectly sinusoidal; instead, it has voltage ripples due to current ripples along the UPS-PDU line (Fig. 7).

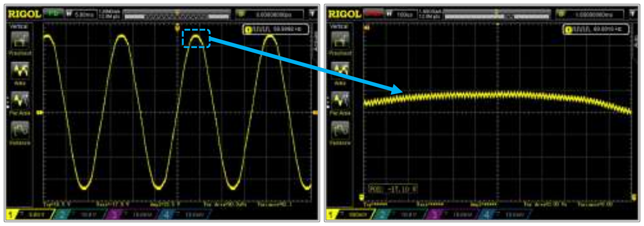

Figure 7: High-frequency voltage ripples at the PDU caused by switching in the server power supply unit.

4.3.3 *Impact of server's power usage on voltage ripples.*

A natural follow-up question is: *do the voltage ripples carry information about server's power usage?*

Note first that if we apply a band-pass filter to keep frequency components within a certain range around the PFC switching frequency (e.g., ~70kHz), the UPS output voltage V_{UPS} signal becomes approximately zero and the voltage relation in Eqn. (1) reduces to

$$\tilde{V}_N = \tilde{V}_{PDU} \approx -R \cdot \sum_{n=1}^{N} \tilde{I}_n \tag{2}$$

where $\tilde{x}$ represents a filtered version of x that only keeps frequency components around the PFC switching frequency. Thus, the attacker's filtered voltage measurement $\tilde{V}_N$ essentially only contains the voltage ripple signal. It is possible that the UPS output voltage V_{UPS} itself also has some high-frequency components (due to, e.g., grid input), but these frequency components are rather weak because of fading over a long distance and hence can be viewed as background noise (Fig. 8(a)).

There are three basic conduction modes for PFC designs: continuous, discontinuous and critical conduction modes [24]. In both discontinuous and critical conduction modes, the current ripple decreases to zero during each switching cycle and the hence peak current can be exceedingly high (twice as much as the average current). Thus, they are mostly designed for low-power devices. In today's servers, power supply units are most commonly designed with a fixed-frequency PFC operating under the continuous conduction mode where the current ripple does not decrease to zero during each PWM-controlled switching cycle (as shown in Fig. 6(a)). We take a closer look at the PFC current ramps in Fig. 6(b). The current goes up when the switch is "ON" (i.e., the MOSFET is turned on), and goes down when the switch is "OFF". The "ON" and "OFF" times are designated as T_{on} and T_{off} in Fig. 6(b), where the period is $T = T_{on} + T_{off}$ and the duty cycle is $D = \frac{T_{on}}{T}$. The duty cycle is regulated within each cycle to ensure that the average current follows the reference current shown in dashed line in Fig. 6(b). The reference current is set based on the sampled input voltage to make the resulting current follow the voltage shape (i.e., improve the power factor to close to 1).

To accommodate the server power change, the current changes and there is a multiplier applied to the current reference sampled from the input voltage. Consequently, as shown in Fig. 6(a), we have a taller current wave when the power is higher and vice versa. Intuitively, the current waves we show in Fig. 6(b) need to rise faster when the server consumes more power, as the current ramp needs to reach higher values. It also needs to drop faster from a higher current to follow the sinusoidal voltage wave. On the other hand, the PFC switching frequency is relatively fixed (with a small temporal variation shown in Fig. 25 in Appendix F). Therefore, when a server consumes more power, the current ripple needs to change faster (i.e., increasing/decreasing faster) within one switching period, and vice versa. Correspondingly, based on the Ohm's Law in Eqn. (2), we expect to see a more prominent high-frequency voltage ripple.

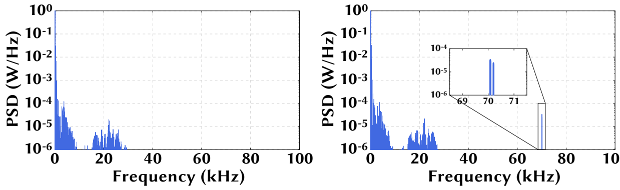

(a) Without a server (b) With a server

Figure 8: High-frequency PSD spikes in PDU voltage caused by the server power supply unit.

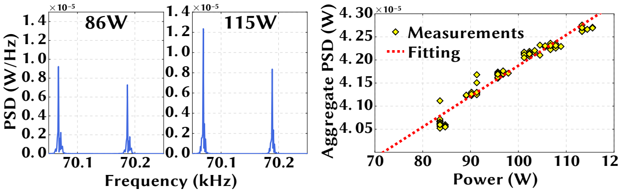

(a) (b)

Figure 9: (a) PSD at different server powers. (b) Server power vs. PSD aggregated within the bandwidth of $69.5 \sim 70.5$kHz for the 495W power supply unit.

To quantify the intensity of the voltage ripple, we use aggregate power spectral density (PSD), i.e., the sum of PSD components over a 1kHz band centered around the PFC switching frequency. We choose 1kHz as our default frequency band, and we will later vary the frequency band (Fig. 12(d)).

In summary, the high-frequency voltage ripple created by a server is expected to be more significant when it consumes more power.

4.4 Experimental validation

We now seek experimental validation on real servers to corroborate our discussion in Section 4.3.3. Here, we only present the results, while the experimental setup is described in Section 5.1.

Single server. We first run only one server with a 495W-rating power supply unit. Fig. 7 shows two zoom-in oscilloscope screenshots that reveal the voltage ripples caused by the server's power supply unit. We further run frequency analysis on the collected voltage signals over each one-second segment. We show the resulting power spectral density (PSD) with and without the server, in Figs. 8(a) and 8(b), respectively. We see that the server produces a PSD spike around 70kHz, presenting a concrete proof of the voltage ripples observable in the power line network.

We then run the server at different power levels and show the PSD around the server's PFC switching frequency in Fig. 9(a). We see that the PSD is higher when the server power consumption is higher, matching with our expectation. We next show the server power vs. aggregate PSD in Fig. 9(b), by summing up all the frequency components within a 1kHz band (69.8 ∼ 70.8kHz). In Section 4.5, we provide an algorithm to identify the frequency band over which the PSD is aggregated. We see that the aggregate PSD

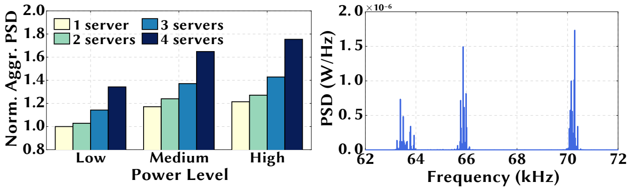

(a) (b)

Figure 10: (a) The aggregate PSD for different numbers of servers. The aggregate PSDs are normalized to that of the single server at low power. (b) Power spectral density of all servers in our testbed showing three distinct PSD groups, each corresponding to a certain type of power supply unit.

Algorithm 1 Calculating Group-wise Aggregate PSD

```
Input: PSD data $P(f)$, frequency band $F$ (e.g., 1kHz), frequency scanning lower/upper bounds $F_{lb}/F_{ub}$
Output: Group-wise aggregate PSD $P_1, P_2, \cdots, P_M$.
Find grid frequency $F_o \leftarrow \max_{45Hz \le f \le 65Hz} P(f)$
for $f$ from $F_{lb}$ to $F_{ub}$ do
    $C_f \leftarrow \frac{P(f-F_o)+P(f+F_o)}{2}$
Keep $C_f$ spikes and discard others (i.e., power line noise)
Generate bands $B[i] = [f - \frac{F}{2}, f + \frac{F}{2}]$ for each $C_f$ spike
Merge $B$ with overlapping frequency bands
Number of groups $M \leftarrow$ number of separate bands in $B$
for each item $B[i] \in B$ do
    $P_i \leftarrow \sum_{f \in B[i]} P(f)$
Return group-wise aggregate PSD $P_i$ for $i = 1, 2, \cdots, M$.
```

monotonically increases with the server power. We also conduct similar experiments with a 350W power supply unit and show the results in Appendix B. We note that given a certain server power, the resulting aggregate PSD varies little, because of the high-frequency ripple signal transmission over power lines without much interference. Finally, we also identify that the switching frequency remains relatively fixed as shown in Fig. 25 in Appendix F.

Multiple servers with identical power supply units. Next, we run four servers, each with a 495W power supply unit. We turn on the servers one by one, record the voltage readings for three different power levels, and calculate the aggregate PSD. Instead of using the absolute value, we show in Fig. 10(a) the relative aggregate PSD normalized with respect to the lowest value when only one server is running. It can be seen that the aggregate PSD becomes greater as we run more servers and/or increase their power usage, which is consistent with our discussion in Section 4.3.3.

Multiple servers with different power supply units. We run all of our 13 servers that have different configurations and power supply units. Table 1 shows the server configuration. We show the PSD of the resulting voltage in Fig. 10(b). We observe three distinct groups of PSD spikes, each corresponding to one type of power supply unit. Based on our individual server experiments, we identify that the PSD spikes around 64kHz are caused by 350W power supply units. The PSD spikes around 66kHz and 70kHz are both created by servers with the 495W power supply units. Despite the same capacity, the two types of 495W power supply units are from different generations (Appendix E) and hence have distinct switching frequencies. In addition, each group consists of several spikes because different power supply units of the same model and generation may still have slightly different switching frequencies.

In summary, our experiments have found and validated that: **(1)** the power supply unit designs of today's servers create high-frequency voltage ripples in the data center power line network; and **(2)** these ripples carry information about servers' power usage at runtime. Note that, like in today's mainstream systems [24], the PFC circuits in our servers operate under the continuous conduction mode which, as shown in Fig. 10(b), causes line voltage ripples with narrow spikes in the frequency domain. For certain high-power servers (close to 1kW or even higher), two PFC circuits may be connected in tandem (a.k.a. interleaving PFC) to meet the demand of large current flows while reducing voltage ripples. Compared to a single-stage PFC, the two individual PFC circuits in an interleaving PFC design operate with a certain time delay between each other and result in line voltage ripples with shorter but wider spikes in the frequency domain [24, 44]. Albeit over a wider range, the high-frequency components in line voltage ripples resulting from PWM switching still become more prominent as a server consumes more power [44]. Therefore, our finding holds broadly regardless of a single-stage or interleaving PFC design.

4.5 Tracking Aggregate Power Usage

Now, we study how the attacker can track the tenants' aggregate power usage based on its measured voltage signal.

4.5.1 *Calculating group-wise aggregate PSD.*

In a multi-tenant data center with servers from different manufacturers, we shall expect to see several *groups* of PSD spikes in the voltage signal, each group consisting of the PSD spikes from similar power supply units (and likely from servers owned by the same tenant). Likewise, we can also divide servers into different groups according to their PFC switching frequencies.

In general, within each group, the aggregate PSD increases when the servers in that group consume more power (Fig. 10(a)). Nonetheless, *even given the same aggregate PSD, servers in one group may have very different power usage than those in the other group (Fig. 22 in Appendix C)* because they have different power supply units and are also likely to have different configurations (e.g., different CPUs). Thus, the total PSD over all the groups may not be a good indicator of the servers' total power consumption; instead, we should consider group-wise aggregate PSD.

We leverage the frequency domain segregation and use Algorithm 1 to identify PSD groups. We use the insight from our experiment that each server creates a pair of PSD spikes separated by twice the nominal power line frequency (i.e., 60Hz in the U.S.) and centered around its PFC switching frequency. Further, the spikes are significantly greater than the power line background noise.

4.5.2 *Tracking tenants' aggregate power usage.*

To launch successful power attacks, the attacker only needs to identify the moments when the tenants' aggregate power is sufficiently high. Thus, knowing the shape of the aggregate power usage is enough.

Given the group-wise aggregate PSD $P_1, P_2, \cdots, P_M$ at runtime, the attacker can track the total power usage of servers in each

Algorithm 2 Timing Power Attacks Using Voltage Side Channel

```
Input: attack threshold $P_{th}$, timer thresholds for $T_{wait}$, $T_{attack}$, and $T_{hold}$
Initiation: current state $S_c \leftarrow idle$, next state $S_n \leftarrow idle$ $T_{wait} \leftarrow 0$, $T_{attack} \leftarrow 0$, and $T_{hold} \leftarrow 0$
loop at regular intervals (e.g., every 10 seconds)
    Use Algorithm 1 to get the aggregate PSDs
    Use historical data to get normalized PSD, $\tilde{P}_1, \tilde{P}_2, \cdots, \tilde{P}_M$
    $\tilde{P} \leftarrow \sum_{m=1}^{M} \tilde{P}_m$
    if $S_c = idle$ then
        if $\tilde{P} < P_{th}$ then
            $S_n \leftarrow idle$
        else $S_n \leftarrow wait$, start $T_{wait}$
    else if $S_c = wait$ then
        if $\tilde{P} \geq P_{th}$ then
            if $T_{wait}$ is expired then
                $S_n \leftarrow attack$, start $T_{attack}$, stop and reset $T_{wait}$
            else $S_n \leftarrow wait$
        else $S_n \leftarrow idle$, stop and reset $T_{wait}$
    else if $S_c = attack$ then
        if $T_{attack}$ is expired then
            $S_n \leftarrow hold$, start $T_{hold}$, stop and reset $T_{attack}$
        else $S_n \leftarrow attack$
    else
        if $T_{hold}$ is expired then
            if $\tilde{P} \geq P_{th}$ then
                $S_n \leftarrow attack$, start $T_{attack}$, stop and reset $T_{hold}$
            else $S_n \leftarrow idle$, stop and reset $T_{hold}$
        else $S_n \leftarrow hold$
    $S_c \leftarrow S_n$
```

group (i.e., a high aggregate PSD means a high power in that group). Nonetheless, the attacker does not know the corresponding absolute server power given a certain aggregate PSD value. Intuitively, however, if all or most group-wise aggregate PSDs are sufficiently high, then it is likely that the tenants' aggregate power usage is also high (i.e., an attack opportunity). Thus, based on this intuition, we first normalize each group-wise aggregate PSD (with respect to its own maximum over a long window size, e.g., 24 hours) and denote the normalized values by $\tilde{P}_1, \tilde{P}_2, \cdots, \tilde{P}_M$. Then, we sum them up $\tilde{P} = \sum_{m=1}^{M} P_m$ and use it as an approximate indicator of the tenants' aggregate power usage.

4.6 Timing Power Attacks

To time power attacks based on the voltage side channel, we propose a threshold-based strategy based on the sum of normalized group-wise aggregate PSDs $\tilde{P}$. Specifically, we set four different states — **Idle**, **Wait**, **Attack**, and **Hold** — and the attacker transitions from one state to another by periodically (e.g., every 10 seconds) comparing $\tilde{P}$ against a threshold P_{th}.

•**Idle** : This is the default state. If $\tilde{P} \geq P_{th}$ is met, the attacker moves to **Wait** and starts a T_{wait} timer.

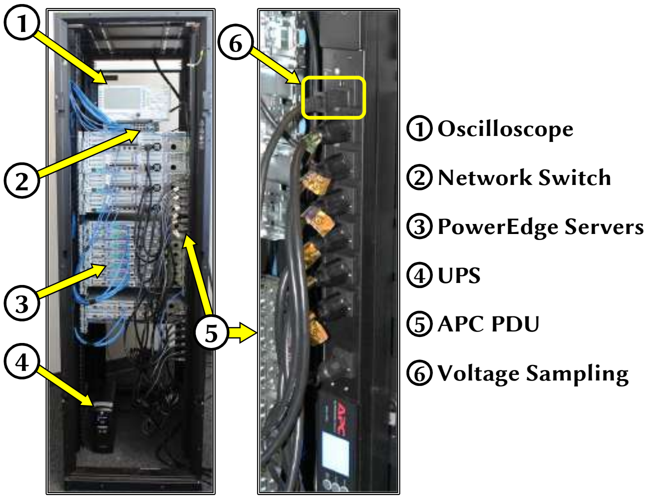

Figure 11: A prototype of edge multi-tenant data center.

•**Wait** : To avoid attacking during transient spikes of $\tilde{P}$, the attacker stays in **Wait** until T_{wait} expires. Then, if $\tilde{P} \geq P_{th}$ still holds, the attacker moves to **Attack** and, otherwise, back to **Idle**.

•**Attack** : In this state, the attacker uses its maximum power consumption for attacks. The attacker stays in **Attack** for T_{attack} time, after which it starts a T_{hold} timer and moves to **Hold**.

•**Hold** : To avoid suspiciously consecutive attacks, the attacker stays in this state until T_{hold} expires. Then, if $\tilde{P} \geq P_{th}$ is still met, it moves back to **Attack** and otherwise to **Idle**.

Finally, we present the formal algorithm description in Algorithm 2.

5 EVALUATION

This section presents our evaluation results of exploiting the voltage side channel for timing power attacks in a scaled-down multi-tenant data center. We focus on how well the attacker can track tenants' aggregate power usage at runtime and how well it can time its power attacks. Our experimental results demonstrate that, by launching non-consecutive attacks no more than 10% of the time, the attacker can successfully detect 64% of all attack opportunities (i.e., true positive rate) with a precision of 50%.

5.1 Methodology

As shown in Fig. 11, we set up 13 Dell PowerEdge servers connected to a single-phase 120V APC8632 rack PDU for our experiments. This setup represents an edge colocation data center, an emerging type of data center located in distributed locations (e.g., wireless towers) [29].

The server configuration is shown in Table 1. The PDU is powered by a CyberPower UPS that is plugged into a power outlet in our university server room. We use a Rigol 1074Z oscilloscope to measure the voltage at a sampling rate of 200kHz. The oscilloscope probe is connected to the PDU through a power cable with polarized NEMA 1-15 plug. While we use an oscilloscope to collect the voltage signal (as we cannot open the power supply unit for lab safety), the attacker can place a small ADC circuit inside its power supply unit to achieve the same purpose in practice.

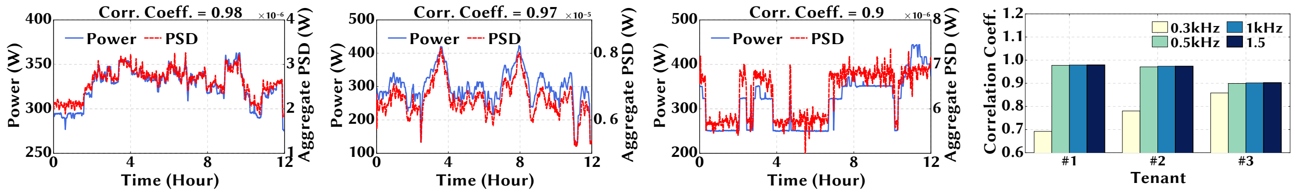

(a) Tenant #1 (63 ∼ 64 kHz)

(b) Tenant #2, (65.5 ∼ 66.5 kHz)

(c) Tenant #3, (69.5 ∼ 70.5 kHz)

(d) Impact of frequency band size (F)

Figure 12: Detection of power shape of different server groups.

Table 1: Server configuration of our experiments.

Tenant	CPU/Memory	Power Supply Rating	PFC Switching Frequency	Number of Servers	Subscribed Power
#1	Xeon/32GB	350W	∼63kHz	4	360W
#2	Dual Xeon/32GB	495W	∼66kHz	2	460W
#3	Xeon/32GB	495W	∼70kHz	4	480W
#4	Pentium/32GB	350W	∼63kHz	3	200W

Tenants. As shown in Table 1, we divide the 13 servers among the four tenants. There are three benign tenants and one attacker (tenant #4). The total capacity of the three benign tenants is 1,300W and the capacity of the attacker is 200W (i.e., 13% of the total capacity subscription). The data center power capacity is 1,250W with 120% oversubscription (i.e., sold as 1,500W) [11, 17].

Workload traces. Like in the prior research [8, 11], the four tenants' server power usage follows four different power traces. Traces for two of the benign tenants are collected from Facebook and Baidu workloads [14, 45], while the other two power traces are generated offline using workload traces (SHARCNET and RICC logs) from [46, 47]. These power traces are scaled to have 75% utilization for the benign tenants and 65% utilization for the attacker. We assign a real workload trace to the attacker so that it behaves normally as benign tenants and stays stealthy. The tenants' total power consumption is shown in Fig. 13. While we use these power traces to reflect the temporal variation in tenants' power usage for our evaluation, *our approach to timing power attacks also applies for any other power traces.*

Duration. We run the experiment for 12 hours and record the power consumption and voltage readings. We also run simulation studies by extending our 12-hour experiment into one year. Specifically, we split the 12-hour data into 10-minute pieces and randomly order them into yearly voltage signals and corresponding power readings. In our yearly trace, the attack opportunities take up 7.5% of the time, consistent with the settings in [8, 11]. Note that, because they are transmitted over power lines, the voltage ripples do not vary significantly over time given a certain power level. Thus, our extended trace still preserves the voltage signal patterns that we would otherwise see in real experiments and hence suffices for our purpose of evaluating the timing accuracy.

Others. By default, in Algorithm 2, we set the frequency band as F = 1kHz, and the scanning lower and upper bounds as F_{lb} =55kHz and F_{ub} =80kHz, respectively. We set T_{wait}=2 minutes, T_{attack}=10 minutes, and T_{hold}=10 minutes. We perform frequency analysis of the voltage signal over each one-second segment.

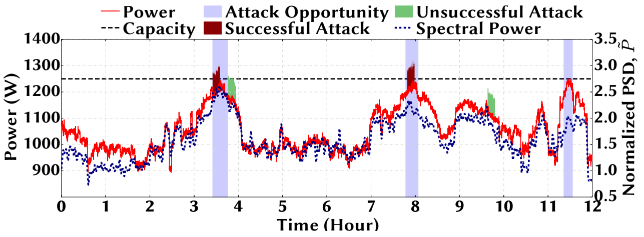

Figure 13: Illustration of power attack.

5.2 Results

We first focus on how well the attacker can track tenants' aggregate power usage at runtime based on the voltage side channel, and then turn to how well the attacker can time its power attacks.

Tracking tenants' power usage. We apply Algorithm 1 to detect groups of PSD spikes (i.e., different tenants in our case) and see how well the per-group aggregate PSD represents the corresponding servers' power consumption. As the attacker knows its own power usage, we separate its own spikes from the PSD scanning process. Fig. 12 shows the three benign tenants' power usage and the corresponding group-wise aggregate PSD from our 12-hour experiment. We see that the aggregate PSDs and the power usage have a strong correlation for all tenants (with correlation coefficient ranging from 0.9 to 0.98), demonstrating the effectiveness of the voltage side channel in extracting the power usage of benign tenants. Note that, for a similar power level, tenant #1 has a lower aggregate PSD than the other two tenants. This further confirms that we need to look at group-wise aggregate PSD to track the power of each server group. We show the power vs. PSD for the three traces in Appendix C.

Next, we study the impact of the choice of frequency band F in Algorithm 1 on the power-PSD relation. We see from Fig. 12(d) that the correlation between the power usage and the resulting aggregate PSD is not quite sensitive to the choice of F, provided that it is higher than 0.5kHz (to include all the PSD spikes).

Illustration of power attacks. Section 4.6 and Algorithm 2 guide the attacker to time power attacks based on the sum of group-wise aggregate PSD of its measured voltage signal. We now illustrate in Fig. 13 a 12-hour trace in our experiment. Concretely, Fig. 13 includes the aggregate power usage (without power attacks), sum of group-wise aggregate PSD, malicious power loads injected by the attacker, and attack opportunities. There are three attack opportunities in 12 hours, each lasting less than 30 minutes and emphasizing the need for precisely timing attacks. We see that two successful attacks are launched during the attack opportunity windows around the 4th and 8th hour. The attacker also launches unsuccessful attacks around the 4th hour and 10th hour. Note that Fig. 13 is to

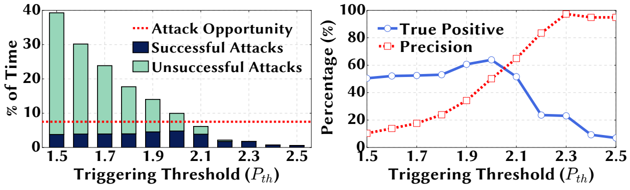

(a)

(b)

Figure 14: Impact of attack triggering threshold P_{th}. The legend "Attack Opportunity" means the percentage of times an attack opportunity exists.

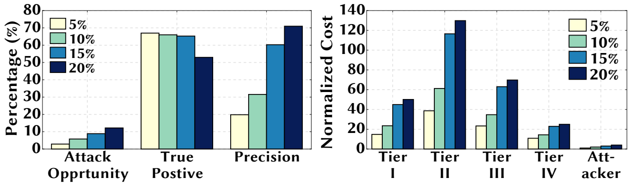

(a)

(b)

Figure 15: Cost and impact of attacker size. x% in the legend indicates the "%" capacity subscribed by the attacker. The tiers specify the infrastructure redundancies, from Tier-I with no redundancy up to Tier-IV with 2N full redundancy.

illustrate what would happen if there are power attacks; if an actual outage occurs due to a power attack, then the power trace following the outage would be changed as servers are out of power.

Timing statistics. We now look into the timing accuracy of our proposed threshold-based attack strategy described in Algorithm 2. Fig. 14(a) shows the impact of the threshold value P_{th}. Naturally, with a lower threshold, the attacker will attack more frequently, but there will be more unsuccessful attacks because the total available attack opportunities remain unchanged.

We consider two metrics for timing accuracy: *True positive rate*: the percentage of attack opportunity captured by the attacker to launch successful attacks. *Precision*: the percentage of power attacks that are successful.

Fig. 14(b) shows the evaluation results under different attack thresholds P_{th}. We see that the true positive rate is high when the attacker sets a lower threshold and launches more attacks, consequently capturing more attack opportunities. Nonetheless, the true positive rate may not always increase by lowering the threshold. This is because of the attack strategy in Algorithm 2: with a low triggering threshold, the attacker sometimes launches an attack prematurely and hence misses on actual attack opportunity that follows immediately, due to the holding time T_{hold} (to meet contractual terms and stay stealthy) before launching another attack. When the triggering threshold is higher, the attacker is more conservative and launches attacks only when it anticipates a sufficiently high aggregate power by benign tenants. Thus, the precision rate increases as the threshold increases.

Impact of attacker size. For stealthiness, the attacker behaves as a benign tenant and launches attacks by increasing its power only to its subscribed capacity (i.e., allowed power limit). Now, we show the impact of the attacker's size (i.e., its subscribed power capacity) on the detection statistics in Fig. 15(a). For this, we keep the benign tenants' capacity fixed and increase both the attacker and data center capacity while we also limit the total attacks under 10% of the time. Naturally, there are more attack opportunities if the attacker has a larger power capacity as it can more easily elevate the aggregate power by itself to create capacity overloads. We also see that true positive rate goes down while the precision goes up when the attacker's power capacity becomes larger. This is because we keep the attack time percentage fixed at 10%. As a result, even with more attack opportunities, the attacker cannot launch more frequent attacks and hence misses more attack opportunities (i.e., lower true positive rate) while its chance of capturing an actual attack opportunity increases (i.e., higher precision).

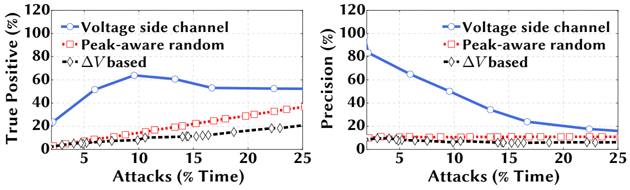

(a)

(b)

Figure 16: Detection statistics for different attack strategies.

Although increasing the attacker's power capacity allows the attacker to launch successful power attacks more easily, the attacker also needs to spend more money for its power capacity subscription and equipment. We now study the cost impact of power attacks. All the costs are normalized with respect to the attacker's own cost when it subscribes 5% of the total subscribed power capacity. We estimate the cost based on the method provided in [11]. The results are shown in Fig. 15(b), demonstrating that the attacker only needs to spend a tiny fraction (3% ~ 16% in our study) of the total resulting losses for the data center operator and other benign tenants. Our findings are similar to those in [8, 11]. In practice, these normalized values correspond to tens of million dollars even for a relatively small data center with only 1MW power capacity [11].

Comparison with other attack strategies. We now compare the timing the power attacks with two other attack strategies.

- *Peak-aware random attack*: This strategy is an improved version of purely random attacks and assumes that the attacker knows the probability of when attack opportunities arise per hour and allocate its total attack times to maximize its overall success rate.
- *ΔV-based attack*: As described in Section 4.2, the attacker simply checks its voltage reading (in RMS) for attacks.

We compare these different attack strategies in terms of their true positive rates and precisions and show the results in Fig. 16. We see that our proposed approach to timing attacks significantly outperforms the peak-aware random attack and ΔV-based attack under our default total attack time of 10%, demonstrating the need of a precise timing for attacks. Importantly, the voltage reading in RMS can be misleading for indicating attack opportunities, since it is predominantly affected by the UPS output voltage V_{UPS} rather than the line voltage drop. Note that if the attacker attacks more frequently, the peak-aware random attack and our approach come closer to each other in terms of the timing accuracy. Nonetheless, frequent attacks are not only prohibited by contracts [23], but also will likely be detected as suspicious behavior.

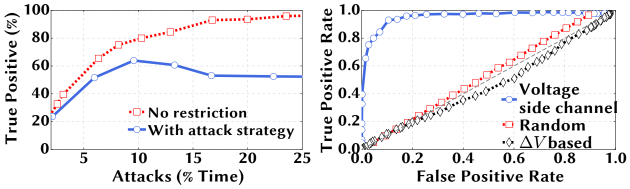

(a) (b)

Figure 17: (a) Impact of the attack strategy (e.g., T_{hold}) on true positive rate. (b) ROC curves showing the accuracy of detection of attack opportunities.

Detection accuracy. Finally, we show the effectiveness of our voltage side channel in detecting attack opportunities when the attacker can attack consecutively without any restriction (e.g., T_{hold}). Fig. 17(a) shows the true positive rates for the cases with and without consecutive attack restrictions. The gap between the two lines indicates that although the voltage side channel can identify attack opportunities, the holding time before launching a new attack for stealthiness and contract constraints can result in a few missing opportunities. Fig. 17(b) shows that our voltage side channel can identify most of the attack opportunities with a low false positive rate. By comparison, the random attack strategy performs rather poorly, and the Δ-based attack is even worse because the measured voltage V_N is mostly affected by the grid and UPS operations rather than tenants' aggregate power (Section 4.2).

6 EXTENSION TO THREE-PHASE SYSTEM

Our previous sections focus on a single-phase power distribution that is mostly used in edge multi-tenant data centers. Next, we extend our study to a three-phase system that is commonly used in large-scale multi-tenant data centers [48].

6.1 Three-Phase Power Distribution System

All large data centers use three-phase AC distributions to deliver power at high efficiency [49]. Each PDU supports 40 ∼ 200kW of server power (10 ∼ 50 server racks) and is oversubscribed by the data center operator, and each tenant typically houses at least one full dedicated server rack. Here, we consider an attacker with multiple server racks sharing one oversubscribed PDU with benign tenants.

There are a few different ways to connect servers in a three-phase system. We show in Fig. 18 the most widely-used three-phase systems, where the servers are connected to two of the phases with a supply voltage at 208V [49]. This is also the most complicated case since each server/server rack is connected to and hence also affected by two different phases. We show another two types of three-phase systems in Appendix H.

6.1.1 Voltage equations in a three-phase system.

As illustrated in Fig. 18, all the server racks connected to the same two phases are considered as one cluster. We represent the total load of each server cluster using their combined current I_{ab}, I_{bc}, and I_{ca}, respectively. Like in Section 4, because the voltage levels are relatively fixed (apart from some temporal variations around the nominal levels), the current flowing into each server cluster are a good indicator of the cluster's power usage.

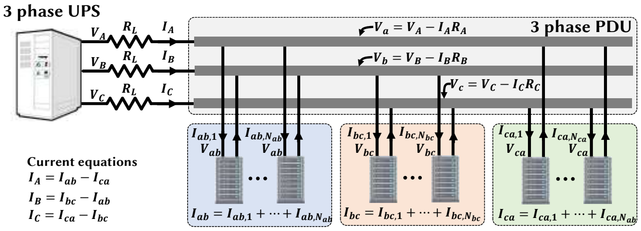

Figure 18: 3-phase power distribution with 2-phase racks.

A distinguishing feature of the three-phase connection is that each server rack is connected to two phases. For each phase, the line voltage drop is affected by the current flowing from the UPS output to the PDU. As shown in the current flow equations in Fig. 18, the line current for each phase jointly depends on two server clusters.

Next, by ignoring the practically negligible line voltage drop between the PDU and servers, we write the voltage V_{ab}, which is supplied by the PDU to the server cluster connected to phase A and phase B, as follows:

$$\begin{aligned} V_{ab} \quad = V_a - V_b \quad &= V_A - I_A \cdot R_L - (V_B - I_B \cdot R_L) \\ &= V_{AB} - R_L \cdot (I_A - I_B) \\ &= V_{AB} - R_L \cdot (2I_{ab} - I_{bc} - I_{ca}), \end{aligned}$$

where the last step follows from $I_A = I_{ab} - I_{ca}$ and $I_B = I_{bc} - I_{ab}$. Similarly we can also write

$$\begin{aligned} V_{bc} &= V_{BC} - R_L \cdot (-I_{ab} + 2I_{bc} - I_{ca}) \\ V_{ca} &= V_{CA} - R_L \cdot (-I_{ab} - I_{bc} + 2I_{ca}). \end{aligned}$$

6.1.2 Exploiting the voltage side channel in a three-phase system.

Like in the single-phase system (Section 4.3), we apply a high-pass filter to keep the high-frequency voltage ripples introduced by servers' PFC circuits, while removing the nominal UPS output voltage frequencies and harmonics. Thus, with a high-pass filter, the line voltage components V_{AB}, V_{BC}, and V_{CA} becomes almost zero. Next, by using $\tilde{x}$ to represent the filtered version of x that only keeps frequency components around the servers' PFC switching frequencies, we get the following relations:

$$\begin{aligned} \tilde{V}_{ab} &\approx -R_L \cdot (2\tilde{I}_{ab} - \tilde{I}_{bc} - \tilde{I}_{ca}) \\ \tilde{V}_{bc} &\approx -R_L \cdot (-\tilde{I}_{ab} + 2\tilde{I}_{bc} - \tilde{I}_{ca}) \\ \tilde{V}_{ca} &\approx -R_L \cdot (-\tilde{I}_{ab} - \tilde{I}_{bc} + 2\tilde{I}_{ca}). \end{aligned}$$

Thus, by collecting the $\tilde{V}_{ab}$, $\tilde{V}_{bc}$ and $\tilde{V}_{ca}$ signals using its voltage probes, the attacker can easily solve the above equation set and extract the high-frequency voltage ripple signals (i.e., $R_L \cdot \tilde{I}_{ab}$, $R_L \cdot \tilde{I}_{bc}$, and $R_L \cdot \tilde{I}_{ca}$) resulting from the server clusters' power usage.

Consequently, based on the approach proposed in Section 4, the total power usage of each server cluster at runtime can be tracked and, when combined together, provides the attacker with an estimate of the total PDU-level power usage for timing its attacks.

In summary, even in the most complicated three-phase power distribution system, *the benign tenants' aggregate power usage can be extracted by the attacker through our discovered voltage side channel for precisely timing its attacks.*[3]

[3] To exploit the voltage side channel in the three-phase system illustrated in Fig. 18, the attacker needs to house at least one server rack in each of the three server clusters (e.g., by pretending to be three tenants) to measure cluster-wise voltage signals.

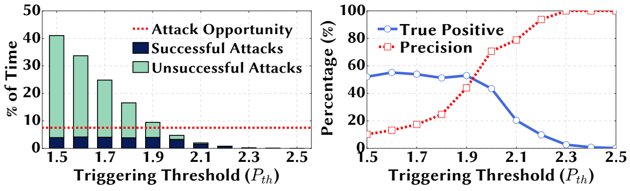

(a) (b)

Figure 19: Performance of voltage side channel for a three-phase 180kW system.

6.2 Evaluation Results

In Section 6.1.2, we have provided a theoretical foundation for timing power attacks based on a voltage side channel in three-phase data centers. Next, we evaluate the timing accuracy of the voltage side channel.

6.2.1 Methodology. We only have a limited access to a large multi-tenant data center with three-phase power distribution and cannot perform experiments due to the destructing nature of our research. Hence, we re-use the experimental results from our servers that have three different types of power supply units. Concretely, we generate three different sets of server power and voltage signal traces based on experiments done on our single-phase server setup with 13 servers. To simulate a large three-phase system, we make 50 copies for each set of trace and add 10% randomness in the power load and PFC switching frequency for each copy. The randomness accounts for the heterogeneity in servers' power supply units and PFC switching frequencies in large systems. Hence, each set of server power and voltage signal traces obtained through our experiments are essentially scaled up by 50 times, and represent the power loads and voltage signals of one server cluster in the three-phase system. Therefore, the three-phase system under consideration has 650 servers (50 times of our single-phase experiment) in each of the three clusters.

In our simulation, the attacker has at least one server rack in each cluster and can measure the phase-to-phase voltages (V_{ab}, V_{bc}, and V_{ca}). Since each server rack is connected to two different phases and the phase voltages are affected by multiple server clusters (hence, multiple power-voltage traces), we use the three-phase voltage equations in Section 6.1.1 to generate the attacker's voltage measurements (V_{ab}, V_{bc}, and V_{ca}). Note that, while we consider the UPS supplies a balanced three-phase voltage (i.e., $V_{AB} = V_{BC} = V_{CA}$, with a 120° phase difference), the supplied voltage is eliminated from the filtered voltages ($\tilde{V}_{ab}$, $\tilde{V}_{bc}$, and $\tilde{V}_{ca}$) which the attacker uses for extracting the server clusters' power usage. The benign tenants and attacker are scaled proportionally according to the composition in Table 1. The resulting attack opportunities take up 7.5% of the time.

6.2.2 Results. Due to the space limitation, we only show the most important results — timing accuracy. Specifically, Fig. 19 shows the true positive and precision rates under different triggering thresholds. We see that, compared to the results in Fig. 14, the timing accuracy is still reasonably good although it becomes a bit worse in the three-phase system. This is mainly due to the randomness added in the power and PSD data when scaling up our edge data center to a large multi-tenant data center, and also due to the fact the attacker needs to track the power usage of three server clusters rather in a three-phase system.

Our results demonstrate the effectiveness of the voltage side channel in terms of timing power attacks in a three-phase system. This matches with our expectation, because the high-frequency voltage ripples generated by servers' PFC circuits exist both in single-phase and three-phase systems and these voltage ripple signals can be transmitted over the data center power line network.

7 DEFENSE STRATEGY

To mitigate the threat of well-timed power attacks, we highlight a few possible defense strategies to degrade the voltage side channel and, more generally, against power attacks.

DC power distribution. The PFC circuit universally built in today's server power supply units is the root cause for high-frequency voltage ripple signals that leak server power usage information through the power lines (i.e., voltage side channel). Thus, the voltage side channel may be eliminated by adopting DC power distributions where the AC to DC conversion is done at the UPS rather than at the server end, as illustrated in Fig. 26 in Appendix G. Naturally, given DC power distributions, the PFC circuit is no longer needed in a server power supply unit. While this is effective for containing the voltage side channel, it requires a holistic industry-wide change, including an entirely new set of power distribution system as well as new power supply units for every server. Thus, we do not anticipate this change will happen any time soon.

Modifying power supply unit. Another approach to getting rid of the voltage side channel is to modify/update the power supply unit design for removing current/voltage ripples. However, it could be challenging to find a suitable substitute for existing mature design. Further, it also requires an industry-wide swap of power supply units, which is highly unlikely in practice.

Jamming signal and filtering. Inspired by jamming attack in communications [50], an inexpensive alternative to the above DC power distribution is to add PSD noise to the PDU and UPS distribution buses around the servers' PFC switching frequency range (e.g., 40kHz to 100kHz). Also, using advanced signal processing techniques and detection, antiphase voltage signal can be injected at the PDU to cancel out the PSD spikes due to server loads. Nonetheless, this may require modification/upgrade of the existing power distribution equipment. In addition, adding jamming signals may reduce the overall power factor of the data center and incur more power losses. Another approach is to install low-pass filters to prevent high-frequency voltage ripple signals from entering the data center power network but, if improperly chosen, the filters may also block legitimate communications (e.g., for network management [43]). Moreover, in practice, filters can reduce the strengths of high-frequency voltage ripples but not completely eliminate them.

Infrastructure reinforcement. Since power attacks target to exploit the data center power infrastructure vulnerability (due to the common practice of oversubscription [14, 15, 17]), another natural approach is to strengthen the infrastructure against power attacks. Toward this end, additional investment can be made to increase the infrastructure redundancy (e.g., installing extra UPSes), but this comes at a great capital expense and can be especially challenging for existing data centers. Moreover, it is a *passive* defense:

attackers can still launch attacks to compromise the desired data center availability, though actual outages may occur less frequently.

Attacker identification. A more proactive approach is to identify attackers inside the data center and evict them in the first place. For example, high-granularity monitoring and rigorous analysis of tenants' power usage can expose a tenant's malicious intention. The main challenge here is to distinguish an attacker from benign tenants because the attacker also follows all contractual limits and can behave like a benign tenant in terms of power usage. In addition, it is even more difficult to identify an attacker if the attacker houses its servers in different racks (pretending to be multiple different benign tenants) and/or launches well-timed power attacks by increasing benign tenants' power usage through request flooding (instead of only relying on the attacker's own power capacity).

To conclude, *it is non-trivial to defend data center power infrastructures against power attacks timed through a voltage side channel.* Thus, effective and inexpensive defense strategies are one of the future research directions in data center power security [7–11].

8 RELATED WORK

Power capping. Power infrastructure oversubscription has been extensively applied for increasing capacity utilization. To handle the ensuing possible capacity overloads, power capping has been proposed, such as throttling CPU [14, 19], rerouting workloads [45], and partially offloading power demand to energy storages [51–53]. However, these techniques cannot be applied in multi-tenant data centers due to the operator's lack of control over tenants' servers. While [15] proposes a market approach for handling capacity overloads in multi-tenant data centers, the market assumes that all tenants are benign and, more crucially, broadcasts the data center's high-power periods (i.e., attack opportunities) unsuspectingly to all tenants including possibly an attacker.

Data center security. Securing data centers in the cyber domain has been widely studied to defend against attacks such as DDoS [3, 4], and data stealing and privacy breach [5, 35, 54, 55]. Meanwhile, an emerging attack vector has been malicious power loads that target the oversubscribed data center power infrastructures to create outages. Studies [7, 9, 10, 36] investigate how VMs can be used to create power overloads in cloud data centers. Another two recent works [8, 11] exploit physical side channels in multi-tenant data centers to time power attacks. In contrast, we propose a novel voltage side channel that is not sensitive to external disturbances, does not require any offline modeling, does not suffer from time lag, and can accurately track power shapes of multiple tenants. A detailed comparison between our work and these related studies is provided in Sections 2.3 and 3.

Power management in multi-tenant data centers. Finally, our work furthers the growing literature on power management in multi-tenant data centers. While the recent studies have predominantly focused on improving power/energy efficiency and reducing cost [15, 56–59], we focus on the complementary aspect of its physical security against devastating power attacks.

9 CONCLUSION

In this paper, we consider the emerging threat of power attacks in multi-tenant data centers. We discover a novel voltage side channel resulting from the high-frequency switching operation in the PFC circuit of a server's power supply unit. The voltage side channel can accurately track benign tenants' power usage and helps the attacker precisely time its power attacks. Our experiment results on a prototype edge data center show that an attacker can effectively use the voltage side channel to utilize 64% of the power attack opportunities. We also highlight a few defense strategies and extend to more complex three-phase power distribution systems.

ACKNOWLEDGEMENT

This work was supported in part by the U.S. NSF under grants CNS-1551661 and ECCS-1610471. We would also like to thank the anonymous reviewers and our shepherd, Dr. Yinqian Zhang, for their valuable comments.

REFERENCES

[1] CNN, "Delta: 5-hour computer outage cost us $150 million," Sep. 07 2016 (http://money.cnn.com/2016/09/07/technology/delta-computer-outage-cost/).

[2] Ponemon Institute, "2016 cost of data center outages," 2016, http://goo.gl/6mBFTV.

[3] J. Mirkovic and P. Reiher, "A taxonomy of DDoS attack and ddos defense mechanisms," *SIGCOMM Comput. Commun. Rev.*, vol. 34, pp. 39–53, Apr. 2004.

[4] S. Yu, Y. Tian, S. Guo, and D. O. Wu, "Can we beat ddos attacks in clouds?," *IEEE Transactions on Parallel and Distributed Systems*, vol. 25, pp. 2245–2254, September 2014.

[5] Y. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, "Cross-vm side channels and their use to extract private keys," in *CCS*, 2012.

[6] M. Korolov, "Data center infrastructure, the often-overlooked security risk," in *DataCenterKnowledge*, April 2018.

[7] X. Gao, Z. Gu, M. Kayaalp, D. Pendarakis, and H. Wang, "ContainerLeaks: Emerging security threats of information leakages in container clouds," in *DSN*, 2017.

[8] M. A. Islam, L. Yang, K. Ranganath, and S. Ren, "Why some like it loud: Timing power attacks in multi-tenant data centers using an acoustic side channel," in *SIGMETRICS*, 2018.

[9] C. Li, Z. Wang, X. Hou, H. Chen, X. Liang, and M. Guo, "Power attack defense: Securing battery-backed data centers," in *ISCA*, 2016.

[10] Z. Xu, H. Wang, Z. Xu, and X. Wang, "Power attack: An increasing threat to data centers," in *NDSS*, 2014.

[11] M. A. Islam, S. Ren, and A. Wierman, "Exploiting a thermal side channel for power attacks in multi-tenant data centers," in *CCS*, 2017.

[12] L. A. Barroso, J. Clidaras, and U. Hoelzle, *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines.* Morgan & Claypool, 2013.

[13] D. Wang, C. Ren, and A. Sivasubramaniam, "Virtualizing power distribution in datacenters," in *ISCA*, 2013.

[14] Q. Wu, Q. Deng, L. Ganesh, C.-H. R. Hsu, Y. Jin, S. Kumar, B. Li, J. Meza, and Y. J. Song, "Dynamo: Facebook's data center-wide power management system," in *ISCA*, 2016.

[15] M. A. Islam, X. Ren, S. Ren, A. Wierman, and X. Wang, "A market approach for handling power emergencies in multi-tenant data center," in *HPCA*, 2016.

[16] X. Fan, W.-D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," in *ISCA*, 2007.

[17] Leagle.com (Case No. 5:13-cv-03093-PSG), "Layton v. Terremark North America, LLC," June 2014.

[18] Intel, "Rack scale design: Architectural requirement specifications," *Document Number: 332937-003*, Jul. 2016.

[19] X. Fu, X. Wang, and C. Lefurgy, "How much power oversubscription is safe and allowed in data centers," in *ICAC*, 2011.

[20] NRDC, "Scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers," *Issue Paper*, Aug. 2014.

[21] Y. Sverdlik, "Google to build and lease data centers in big cloud expansion," in *DataCenterKnowledge*, April 2016.

[22] Apple, "Environmental responsibility report," 2016.

[23] Internap, "Colocation services and SLA," http://www.internap.com/internap/wp-content/uploads/2014/06/Attachment-3-Colocation-Services-SLA.pdf.

[24] On Semiconductor, "Power factor correction (PFC) handbook," http://www.onsemi.com/pub/Collateral/HBD853-D.PDF.

[25] H. W. Beaty and D. G. Fink, *Standard handbook for electrical engineers.* McGraw-Hill New York, 2007.

[26] EdgeConnex, http://www.edgeconnex.com/.

[27] DatacenterMap, "Colocation USA," http://www.datacentermap.com/usa/.

[28] "Colocation market by solutions, end users, verticals & region - worldwide market forecast and analysis (2013 - 2018)," http://www.researchandmarkets.com/research/pxndbm/colocation_market.
[29] DatacenterKnowledge, "Vapor IO to sell data center colocation services at cell towers," http://www.datacenterknowledge.com/archives/2017/06/21/vapor-io-to-sell-data-center-colocation-services-at-cell-towers.
[30] Vapor IO, "The edge data center," https://www.vapor.io/.
[31] Telecommunications Industry Association, "Data center standards overview," *TIA 942*, 2005 (amended in 2014).
[32] Colocation America, "Data center standards (Tiers I-IV)," 2017, https://www.colocationamerica.com/data-center/tier-standards-overview.htm.
[33] M. Sedaghat, E. Wadbro, J. Wilkes, S. D. Luna, O. Seleznjev, and E. Elmroth, "Diehard: Reliable scheduling to survive correlated failures in cloud data centers," in *CCGrid*, 2016.
[34] M. Sheppy, C. Lobato, O. V. Geet, S. Pless, K. Donovan, and C. Powers, "Reducing data center loads for a large-scale, low-energy office building: NREL's research support facility," Nov. 2011.
[35] S.-J. Moon, V. Sekar, and M. K. Reiter, "Nomad: Mitigating arbitrary cloud side channels via provider-assisted migration," in *CCS*, 2015.
[36] X. Gao, Z. Xu, H. Wang, L. Li, and X. Wang, "Reduced cooling redundancy: A new security vulnerability in a hot data center," in *NDSS*, 2018.
[37] DatacenterKnowledge, "Can edgeconnex disrupt incumbent data center providers?," http://www.datacenterknowledge.com/archives/2016/09/13/can-edgeconnex-disrupt-incumbent-data-center-providers/.
[38] Vertiv, "Flexible, efficient scalable UPS for room or row-based applications," *Liberty APS Product Brocure*, 2016.
[39] Energy Star, "Computers specification version 7.0," 2018, https://www.energystar.gov/products/spec/computers_specification_version_7_0_pd.
[40] IECEE, "IEC 61000-3-2:2018: Electromagnetic compatibility (EMC) - part 3-2," http://www.onsemi.com/pub/Collateral/HBD853-D.PDF.
[41] Microchip, "Switch mode power supply (SMPS) topologies," http://ww1.microchip.com/downloads/en/AppNotes/01114A.pdf.
[42] Infineon Technologies AG, "PFC boost converter design guide," https://goo.gl/MePNFj.
[43] L. Chen, J. Xia, B. Yi, and K. Chen, "Powerman: An out-of-band management network for datacenters using power line communication," in *NSDI*, 2018.
[44] J. Zhang, J. Shao, P. Xu, F. C. Lee, and M. M. Jovanovic, "Evaluation of input current in the critical mode boost pfc converter for distributed power systems," in *IEEE Applied Power Electronics Conference and Exposition*.
[45] G. Wang, S. Wang, B. Luo, W. Shi, Y. Zhu, W. Yang, D. Hu, L. Huang, X. Jin, and W. Xu, "Increasing large-scale data center capacity by statistical power control," in *EuroSys*, 2016.
[46] D. G. Feitelson, D. Tsafrir, and D. Krakov, "Experience with using the parallel workloads archive," *Journal of Parallel and Distributed Computing*, vol. 74, no. 10, pp. 2967–2982, 2014.
[47] Parallel Workloads Archive, http://www.cs.huji.ac.il/labs/parallel/workload/.
[48] N. Rasmussen, "High-efficiency ac power distribution for data centers," *Schneider Electric White Paper Library*, http://www.apc.com/salestools/nran-6cn8pk/nran-6cn8pk_r2_en.pdf.
[49] N. Rasmussen, "Efficiency and other benefits of 208 volt over 120 volt input for it equipment," *Schneider Electric White Paper Library*, http://www.apc.com/salestools/SADE-5TNQZ7/SADE-5TNQZ7_R3_EN.pdf.
[50] W. Xu, W. Trappe, and Y. Zhang, "Anti-jamming timing channels for wireless networks," in *WiSec*, 2008.
[51] S. Govindan, D. Wang, A. Sivasubramaniam, and B. Urgaonkar, "Leveraging stored energy for handling power emergencies in aggressively provisioned datacenters," in *ASPLOS*, 2012.
[52] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, "Energy storage in datacenters: what, where, and how much?," in *SIGMETRICS*, 2012.
[53] D. S. Palasamudram, R. K. Sitaraman, B. Urgaonkar, and R. Urgaonkar, "Using batteries to reduce the power costs of internet-scale distributed networks," in *SoCC*, 2012.
[54] M. Guri, M. Monitz, Y. Mirski, and Y. Elovici, "Bitwhisper: Covert signaling channel between air-gapped computers using thermal manipulations," in *IEEE Computer Security Foundations Symposium*, 2015.
[55] P. C. Kocher, J. Jaffe, and B. Jun, "Differential power analysis," in *CRYPTO*, 1999.
[56] M. A. Islam, H. Mahmud, S. Ren, and X. Wang, "Paying to save: Reducing cost of colocation data center via rewards," in *HPCA*, 2015.
[57] C. Wang, N. Nasiriani, G. Kesidis, B. Urgaonkar, Q. Wang, L. Y. Chen, A. Gupta, and R. Birke, "Recouping energy costs from cloud tenants: Tenant demand response aware pricing design," in *eEnergy*, 2015.
[58] N. Nasiriani, C. Wang, G. Kesidis, B. Urgaonkar, L. Y. Chen, and R. Birke, "On fair attribution of costs under peak-based pricing to cloud tenants," in *MASCOTS*, 2015.
[59] M. A. Islam, X. Ren, S. Ren, and A. Wierman, "A spot capacity market to increase power infrastructure utilization in multi-tenant data centers," in *HPCA*, 2018.
[60] On Semiconductor, "Switch-mode power supply reference manual," https://www.onsemi.com/pub/Collateral/SMPSRM-D.PDF.

APPENDIX

A THE NEED OF PFC IN SERVER'S POWER SUPPLY UNIT

Here, we discuss the implication of not using a power factor correction (PFC) circuit in a server's power supply unit. Fig. 20(a) shows a AC-to-DC converter without the PFC stage. The output of the bridge-rectifier is delivered to the load (e.g., CPU, disk, etc.) through a diode. A bulk capacitor is placed at the output to stabilize the voltage. As shown in Fig. 20(b), this circuit can provide a relatively stable DC voltage to the load. However, due to the diode, the load gets current from the source only when the load voltage is lower than the rectifier output voltage. It gets its needed current from the bulk capacitor when the rectifier output voltage is lower than the load voltage. Hence, the bridge rectifier only conducts current during a brief moment near the input voltage's peak when the input voltage is higher than the load voltage (which is also the voltage across the bulk capacitor). In addition, during the rectifier conduction, the external source line powers both the load and the bulk capacitor, creating a high current spike from the source line. Fig. 20(c) shows the source voltage at the AC-to-DC converter's input and the current supplied by the source line (i.e., drawn from the UPS through the PDU). The large difference between the voltage and the current shape reduces the power factor to a level much less than 1 (i.e., for much of the time, the UPS and PDU are not delivering any actual current or power to the server).

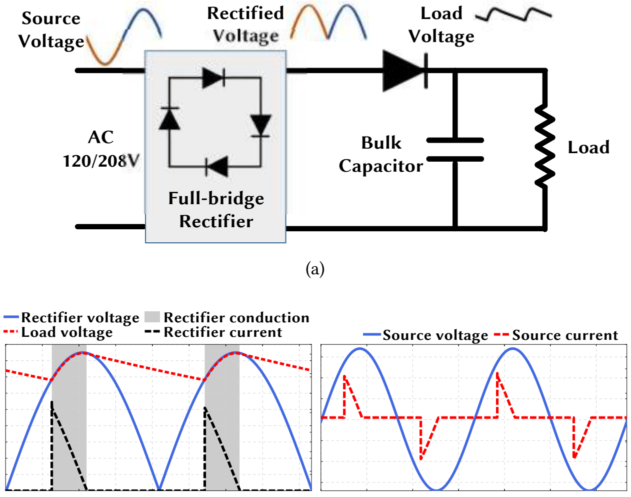

(a)

(b)

(c)

Figure 20: (a) A basic AC/DC power supply unit without a PFC circuit. (b) Rectifier delivers current in bursts. (c) The heavily-distorted current drawn by the server's power supply unit without the PFC.

Thus, for improving power factor as mandated by [39, 40], a PFC circuit (which universally uses the PWM switching design due to its high efficiency [60]) is needed for server's power supply unit.

B SERVER POWER VS. PSD FOR ALTERNATE POWER SUPPLY UNIT

Fig. 21 shows the relation between the server power and the resulting PSD for one of our servers with a 350W power supply unit. As in Fig. 9(b), we see a monotonic increasing relation between the aggregate PSD and server's power consumption.

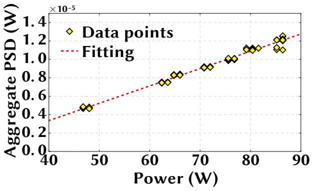

Figure 21: Server power vs. PSD aggregated within the bandwidth of 63 ∼ 64kHz for the 350W PSU.

C POWER VS. PSD FOR DIFFERENT SERVERS

Fig. 22 shows the power vs. aggregate PSD plot for the three benign tenants from our 12-hour experiment. This figure is based on the same results of Fig. 12. Instead of plotting all data points from the 12-hour trace, we randomly choose 500 data points for this figure. It reveals that for the same power level, tenant #1's servers have a smaller aggregate PSD than the other two tenants. This supports our choice of treating each group of spikes separately.

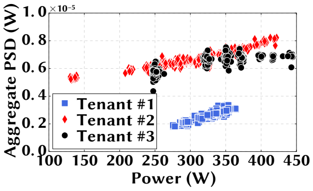

Figure 22: Power vs. PSD plot of different server groups.

D WHEN UPS IS ON BATTERY.

In our previous experiments, the UPS is working in the normal mode without performing voltage regulation. Now, we see if the voltage ripple is still prominent when the UPS is active and providing power through its battery. Fig. 23(a) shows the waveforms of the UPS output voltage in the normal and on-battery modes: the UPS passes the grid voltage to the PDU without much modification in the normal mode, whereas it produces an alternating square wave with a duty cycle of 0.5 when on battery (typically due to loss of grid power). Intuitively, when on battery, the non-sinusoidal voltage will add additional frequency components. When the UPS is on battery, we turn on one server and show in Fig. 23(b) the resulting PSD of the measured voltage. While the battery operation introduces additional frequency components (mostly at the lower frequencies), the PSD spike due to the server's PFC switching is still clearly visible at around 70kHz. Along with our previous experiments, this demonstrates that, even though the UPS modifies the incoming grid voltage, the attacker can still exploit the PSD spikes to estimate servers' aggregate power usage.

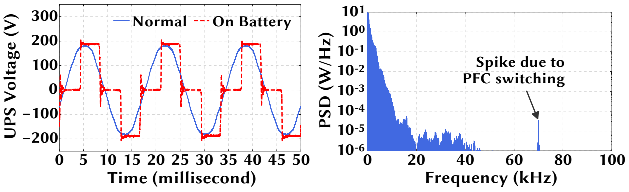

(a) Voltage

(b) On battery backup

Figure 23: (a) Voltage reading at the PDU with the UPS running in normal mode and in battery-backup mode. (b) The PSD spike from server power supply unit is still visible with the non-sinusoidal voltage during the battery-backup mode.

E POWER SUPPLY UNITS

Fig. 24 shows the three different types of power supplies we have in our servers. Even though all of these are used by the Dell PowerEdge servers, they have different model numbers and are manufactured by different companies. While the 350W power supply unit apparently has a different size and rating than the other two, and therefore cannot be swapped, the similar sized 495W power supplies are also not allowed (server does not boot) to be swapped due to their model differences.

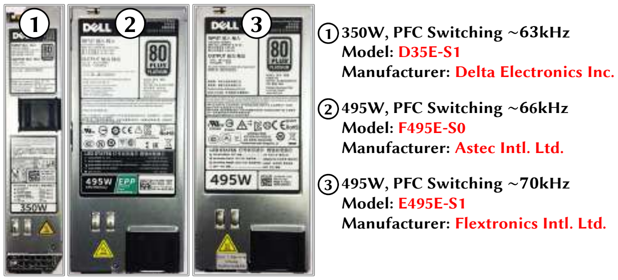
1
2
3
350W
495W
495W
① 350W, PFC Switching ∼63kHz
Model: D35E-S1
Manufacturer: Delta Electronics Inc.
② 495W, PFC Switching ∼66kHz
Model: F495E-S0
Manufacturer: Astec Intl. Ltd.
③ 495W, PFC Switching ∼70kHz
Model: E495E-S1
Manufacturer: Flextronics Intl. Ltd.

Figure 24: The three different types of server power supplies in our experiments.

F PFC SWITCHING FREQUENCY

Here we investigate the temporal variation of the switching frequency of the server power supply unit. For this, we record the PDU voltage reading for two hours with one server running. We then get the PSD for each one-second sample and extract the switching frequency from the PSD spike around 70kHz. Fig. 25 shows the probability mass function (PMF) of the switching frequency. We see that while not perfectly fixed, the switching frequency remains mostly within a narrow (<100Hz) window. Due to this small temporal variation combined with the random manufacturing imperfections, PSD spikes from multiple servers with the same (model and generation) power supply unit do not perfectly overlap with each other, and therefore generates a *group* of spikes in the voltage PSD.

G DC POWER DISTRIBUTION

We show a data center a with DC power distribution in Fig. 26. In this type of data center, a UPS converts the grid's three phase AC voltage to a 380V-DC voltage and distributes inside the data center through DC PDUs. DC distribution requires all servers to have DC power supply units which use only a DC to DC converter

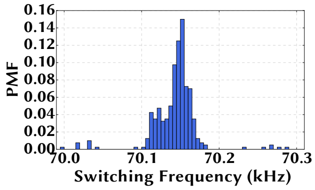

Figure 25: PMF shows the PFC switching frequency only fluctuates slightly.

to supply regulated 12V to server internal components. Unlike AC power supplies, the DC power supplies do not require a power factor correction circuit.

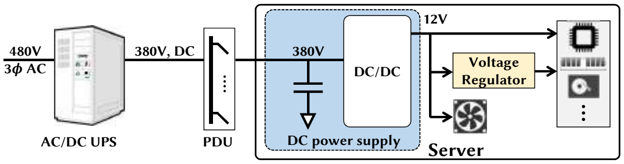

Figure 26: DC power distribution with DC server power supply unit that has no PFC circuit.

H DIFFERENT THREE-PHASE LINE CONNECTIONS

We show two alternate three-phase line connections used in large multi-tenant data centers.

Fig. 27 shows the connectivity for a three-phase line-to-neutral connection that supplies power to servers at 120V. In this type of connection, a server/server rack is connected to only one phase. Thus, the voltage drop of each phase depends on only the servers connected to that phase. As a result, this can be handled as three separate single-phase systems, and the attacker can easily track the tenants' per-phase total power usage by housing one server rack on each phase.

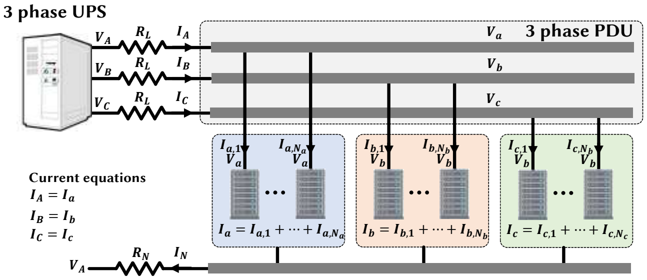

Figure 27: Three-phase distribution with single-phase, 120V server/server racks.

Fig. 28 shows the line connectivity that uses three-phase distribution all the way down to servers. This type of connection is not common because of the three-phase server power supply unit requirement. The 3-phase power supply unit equally loads its three phases, and hence this type of connection can be treated as a single-phase system replicated three times.

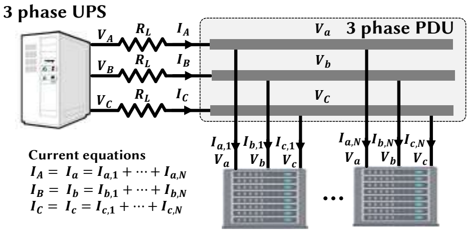

Figure 28: Three-phase distribution with three-phase server/server racks.