

Why Some Like It Loud: Timing Power Attacks in Multi-tenant Data Centers Using an Acoustic Side Channel

MOHAMMAD A. ISLAM, University of California, Riverside
LUTING YANG, University of California, Riverside
KIRAN RANGANATH, University of California, Riverside
SHAOLEI REN, University of California, Riverside

The common practice of power infrastructure oversubscription in data centers exposes dangerous vulnerabilities to well-timed power attacks (i.e., maliciously timed power loads to overload the infrastructure capacity), possibly creating outages and resulting in multimillion-dollar losses. In this paper, we focus on the emerging threat of power attacks in a multi-tenant data center, where a malicious tenant (i.e., attacker) aims at compromising the data center availability through power attacks. We discover a novel acoustic side channel resulting from servers' cooling fan noise, which can help the attacker time power attacks at the moments when benign tenants' power usage is high. Concretely, we exploit the acoustic side channel by: (1) employing a high-pass filter to filter out the air conditioner's noise; (2) applying non-negative matrix factorization with sparsity constraint to demix the received aggregate noise and detect periods of high power usage by benign tenants; and (3) designing a state machine to guide power attacks. We run experiments in a practical data center environment as well as simulation studies, and demonstrate that the acoustic side channel can assist the attacker with detecting more than 50% of all attack opportunities, representing state-of-the-art timing accuracy.

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

Additional Key Words and Phrases: Acoustic side channel; data center; power attack.

ACM Reference Format:

Mohammad A. Islam, Luting Yang, Kiran Ranganath, and Shaolei Ren. 2018. Why Some Like It Loud: Timing Power Attacks in Multi-tenant Data Centers Using an Acoustic Side Channel. *Proc. ACM Meas. Anal. Comput. Syst.* 2, 1, Article 6 (March 2018), 33 pages. <https://doi.org/10.1145/3179409>

1 INTRODUCTION

The exploding demand for cloud services and ubiquitous computing at the Internet edge has spurred a significant growth of multi-tenant data centers (also referred to as “colocation”). The U.S. alone has nearly 2,000 multi-tenant data centers, which are experiencing a double-digit annual growth rate and account for about 40% of all data center energy consumption [1–3]. Unlike a multi-tenant cloud platform where users rent virtual machines (VMs) on shared servers owned by the cloud provider, a multi-tenant data center offers shared non-IT infrastructures (e.g., power and cooling) for multiple tenants to house their own physical servers. It serves as a cost-effective data center

Authors' addresses: Mohammad A. Islam, University of California, Riverside, misla006@ucr.edu; Luting Yang, University of California, Riverside, lyang029@ucr.edu; Kiran Ranganath, University of California, Riverside, krang006@ucr.edu; Shaolei Ren, University of California, Riverside, sren@ece.ucr.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

2476-1249/2018/3-ART6 \$15.00

<https://doi.org/10.1145/3179409>

solution to almost all industry sectors, including large IT companies (e.g., 25% of Apple’s servers are housed in multi-tenant data centers [4]).

Naturally, it is extremely important to provide a highly reliable power supply to tenants’ servers in a multi-tenant data center. To accomplish this, data center operators have typically employed backup power and infrastructure redundancy (e.g., duplicating each power supply equipment, or 2N redundancy, as illustrated in Fig. 1), safeguarding multi-tenant data centers against random power equipment faults and utility power outages. For example, power availability in a multi-tenant data center with state-of-the-art 2N redundancy can exceed 99.995% [5–7].

The high availability of data center power infrastructures comes at a huge cost: the capital expense (CapEx) is around U.S.\$10-25 for delivering each watt of power capacity to the IT equipment, taking up 60+% of a data center operator’s total cost of ownership over a 10-year lifespan [8–11]. Thus, in order to reduce and/or defer the need for infrastructure expansion, a common technique is power *oversubscription*: similarly as in other industries (e.g., airline), a multi-tenant data center operator sells its available data center infrastructure capacity to more tenants than can be supported. The rationale of power oversubscription is that different tenants typically do not have peak power consumption at the same time. The current industry average is to have a 120% oversubscription (yielding 20% extra revenue without constructing new capacities) [12, 13]. Moreover, power oversubscription is also commonly found in owner-operated data centers (e.g., Facebook [10]), and more aggressive oversubscription [14, 15] has been advocated.

Despite the compelling economic benefit, power oversubscription is not risk-free and can potentially create dangerous situations. Concretely, although generally uncommon, tenants’ aggregate power demand can exceed the design power capacity (a.k.a. *power emergency*) when their power consumption peaks simultaneously. Power emergencies compromise infrastructure redundancy protection (illustrated in Fig. 1) and can increase the outage risk by 280+ times compared to a fully-redundant case [5]. Moreover, data center power infrastructures are not as reliable as desired. In fact, compared to cyber attacks, power equipment failures are even more common reasons for data center outages, for which overloading the design power capacity is a primary root cause [16, 17]. For example, despite backup power equipment and redundancy, a power outage recently occurred in British Airways’s data center and cost over U.S.\$100 million [18].

As a consequence, the significant outage risk associated with power emergencies has prompted active precautions. Concretely, due to the lack of control over tenants’ servers, a multi-tenant data center operator typically restricts tenants’ “normal” power usage to be below a fraction (usually 80%) of their subscribed capacities as stipulated by contractual terms. That is, tenants may only use their full subscribed capacities in limited occasions, and non-compliant tenants may face power cuts and/or eviction [19, 20]. Thus, this can effectively eliminate most, if not all, severe power emergencies, thus achieving the designed availability.

While power oversubscription has been regarded as *safe* due to safeguard mechanisms, recent studies [14, 21–23] have demonstrated an emerging threat – power attacks, i.e., malicious power loads that aim at overloading the shared capacity – which could create frequent power emergencies and compromise data center availability. Although there are only limited attack opportunities as illustrated in Fig. 2, the impact of power attacks is devastating. As shown in Table 2 in Appendix A, even if power attacks can create power emergencies for only 3.5% of the time, multi-million-dollar losses are incurred by both the operator and affected tenants (the calculation method is available in Appendix A).

In a multi-tenant data center, a malicious tenant (i.e., attacker) must precisely time its peak power usage in order to create successful power attacks without violating the operator’s contract: the attacker only uses its full subscribed capacity when the power demand of other benign tenants’ is high. In the existing research [22], such precise timing is achieved through the help of a thermal side

channel resulting from heat recirculation — benign tenants’ server heat, which can recirculate to the attacker’s temperature sensors, is a good indicator of their power usage. Nonetheless, exploiting the thermal side channel has several key *limitations*. First, heat containment techniques are increasingly common in modern data centers to improve cooling efficiency and thus can effectively mitigate, or even eliminate, the thermal side channel. Second, in order to time its power attacks, the attacker must be able to construct a data center heat recirculation model, which can deviate significantly from the actual environment and lower the timing accuracy. Last but not least, it may take a long time (> 1 minute) for the heat generated by distant servers to affect the attacker’s temperature sensor, rendering the estimation possibly outdated. All these factors would contribute to the limited applicability of the thermal side channel in practice.

Contributions of this paper. *This paper focuses on the emerging threat of power attacks in multi-tenant data centers and exploits a novel side channel — acoustic side channel resulting from servers’ noise generated by cooling fans — which assists an attacker with timing its power attacks.* Concretely, the key idea we exploit is that the energy of noise generated by a server’s cooling fans increases with its fan speed measured in revolutions per minute (RPM), which is well correlated with the server power (Section 4.1.2). Thus, through measurement of the received noise energy using microphones, an attacker can possibly infer the benign tenants’ power usage and launch well-timed power attacks, which significantly threaten the data center availability. Nonetheless, there are three *key challenges* to exploit the acoustic side channel.

- *How to filter out the computer room air conditioner’s (CRAC’s) fan noise?* In a data center, the volume of CRAC’s fan noise is often significantly greater than that of servers’ fan noise, thus making the servers’ fan noise undetectable.

- *How to relate the received aggregate noise energy with benign tenants’ aggregate power consumption?* There are many noise sources (e.g., servers) in a data center, all arriving at the attacker’s microphones through different attenuation paths. Thus, the mixed noise energy measured by the attacker has a rather poor correlation with benign tenants’ aggregate power usage.

- *How to detect real attack opportunities?* As various types of disturbances can create spikes in the attacker’s received noise energy, the attacker must be able to avoid these fake attack opportunities and launch attacks at the right moments.

In this paper, we address all these challenges (Section 4). First, we investigate differences between servers’ fan noise and the CRAC’s fan noise in terms of frequency characteristics, and then propose a high-pass filter that can filter out most of the CRAC’s fan noise while preserving the acoustic side channel. Second, we propose an affine non-negative matrix factorization (NMF) technique with sparsity constraint, which helps the attacker demix its received aggregate noise energy into multiple consolidated sources, each corresponding to a group of benign tenant’s server racks that tend to have correlated fan noise energy. Thus, when all or most of the consolidated sources have a relatively higher level of noise energy, it is more likely to have an attack opportunity. More importantly, noise energy demixing is achieved in a *model-free* manner: the attacker does *not* need to know any model of noise propagation. Third, we propose an attack strategy based on a finite state machine, which guides the attacker to enter the “attack” state upon detecting a prolonged high noise energy.

We run experiments in a real data center environment to evaluate the effectiveness of our proposed acoustic side channel in terms of timing accuracy. In addition, we complement the experiments with simulation studies over a longer timescale. Our results show that the attacker can successfully capture 54% of the attack opportunities with a precision rate of 48%, potentially creating a million-dollar financial loss yet spending a small fraction (between 3% and 23%) of the created loss. Moreover, our achieved timing accuracy is comparable to the best-known result reported by the existing research [22]. Finally, we discuss a possible set of common defense strategies to safeguard the

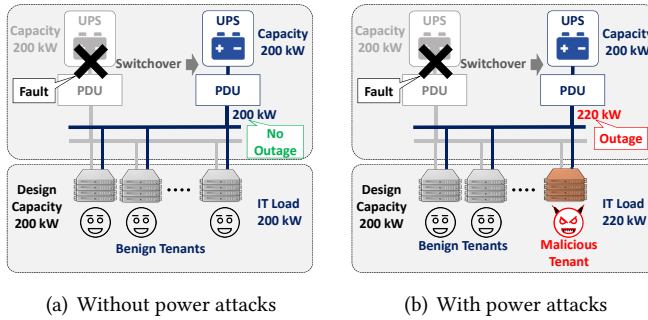


Fig. 1. Loss of redundancy protection due to power attacks in a Tier-IV data center.

data center infrastructure, such as increasing infrastructure resilience, mitigating the acoustic side channel, and early detecting malicious tenants (Section 6).

To facilitate future research on data center acoustic side channels by other researchers, we have also made our noise recordings along with server measurements, such as power and fan speeds, publicly available at [24].

2 OPPORTUNITIES FOR POWER ATTACKS

We here discuss the multi-tenant data center power infrastructure vulnerability, and show opportunities for well-timed power attacks.

2.1 Multi-tenant Power Infrastructure

As illustrated in Fig. 1, a multi-tenant data center typically has a hierarchical power infrastructure with the uninterrupted power supply (UPS) sitting at the top. The UPS acts as a buffer between the grid electricity and downstream equipment, providing conditioned power and facilitating seamless switch-over to backup generators during grid failures. Each UPS is connected to one or multiple power distribution units (PDUs) which supply power to the server racks. Each rack also has its own power strip (often called rack PDU) to connect the servers. All the power equipment have circuit breakers to protect against power surges as well as to isolate faulty equipment from the rest.

An important notion in data centers is “design capacity” (also called *critical power budget/capacity*), indicating the capacity of conditioned power supplied to IT equipment (e.g., servers). The cooling system taking away the heat from servers is also sized based on the designed power capacity. Data center capacity, therefore, is often measured based on the total designed power capacity, while it also includes the matching cooling capacity.

Most data centers have some levels of redundancy to handle random equipment failures. Specifically, *data centers are classified into four tiers* [6, 25]: a Tier-I data center does not have any redundancy, a Tier-II data center has $N+1$ redundancy only for the UPS and backup generators, while Tier-III and Tier-IV data centers have $N+1$ and $2N$ redundancy for the entire power infrastructure, respectively. Fig. 1 shows a Tier-IV data center with $2N$ redundancy.

A multi-tenant data center leases rack-wise power capacity to tenants based on its *design capacity*, and all tenants are required to meet per-rack capacity constraints. While megawatt UPSes are not uncommon, data centers often install multiple smaller UPSes (~ 200 - 300 kW), each serving one or two PDUs. For example, a large Tier-IV data center may have multiple independent sets of $2N$ redundant infrastructures. In addition, power capacity is also deployed on an *as-needed* basis: new capacity is added only when existing capacity is exhausted.




Fig. 2. Infrastructure vulnerability to attacks. An attacker injects timed malicious loads to create overloads.

2.2 Opportunities

Vulnerability to power attacks. While power oversubscription is common [13, 14, 21], multi-tenant data center operators use contractual restrictions to prohibit tenants from using their full capacities all the time (e.g., a tenant’s normal power usage cannot exceed 80% of its subscribed capacity) [19, 20]; involuntary power cuts and/or eviction may apply to non-compliant tenants. Thus, this can keep the typical aggregate power demand well below the designed capacity, achieving the designed availability.

We illustrate this point in Fig. 2, where we show aggregate power trace of four tenants subscribing a total capacity of 15.6 kW while the designed capacity is 13 kW with a 120% oversubscription.¹ In normal situations, the total power remains below the design capacity throughout our 12-hour trace. Note that our power trace includes common workloads such as data processing and web services housed in a multi-tenant data center [11, 26, 27].

The safeguards, however, are ineffective and vulnerable to *well-timed* malicious power attacks. As shown in Fig. 2, an attacker can intentionally inject malicious loads by increasing power to its maximum subscribed capacity when the other benign tenants also have a high power demand. Consequently, in contrast to the benign case, we see two power emergencies in the 12-hour trace. Here, the attacker’s peak power only lasts for 10 minutes at a time, thereby not violating the contract yet enough to trip the circuit breaker. Note that, even a benign tenant may occasionally reach its full subscribed capacity, but unlike in the malicious case, these random peaks do not necessarily coincide with the peak of other tenants.

Impact of power attacks. The immediate impact of power attacks is overloading the design capacity and compromising the infrastructure redundancy, which is extremely dangerous. We use a state-of-the-art Tier-IV data center with 2N redundancy to illustrate this point in Fig. 1. Specifically, Fig. 1(a) shows a design capacity of 200kW and, because of the 2N redundancy design, there are two independent power paths each having a capacity of 200kW. The total IT load is equally shared by the two independent paths. Without a power attack, even though one of the power paths fails, the load is switched to the alternate path without outage. Hence, random single path failures are handled by the redundancy design. Now, suppose that a power attack overloads the design capacity by 10% and that the total IT load is 220kW. As shown in Fig. 1(b), with a power attack, an actual outage occurs followed by a single power path failure.

Thus, we see that *the data center loses its redundancy protection when it is under successful power attacks*, which can increase the outage risk by 280 times compared to the redundant case [5, 25]. We can draw similar conclusions for Tier-II and Tier-III data centers with N+1 redundancy, although the degree of redundancy loss is even worse than a Tier-IV data center. For Tier-I data center without redundancy protection, a successful prolonged power attack (e.g., 10 minutes) can lead to an outage.

As shown in Table 2 in Appendix A, even if redundancy protection is compromised by power attacks for only 3.5% of the time, multi-million-dollar losses are incurred, let alone the loss of customers for the victim data center operator.

¹The capacity setting is based on our experimental setup in Section 5.1.

In summary, despite the infrastructure redundancy and contractual safeguards in place, *a multi-tenant data center with power oversubscription opens up abundant opportunities for well-timed power attacks that can result in significant financial losses.*

3 THREAT MODEL AND CHALLENGES

We now introduce the threat model and show challenges faced by an attacker for successful attacks.

3.1 Threat Model

Tenants typically sign yearly leases in multi-tenant data centers. Our threat model consists of a malicious tenant (i.e., attacker) that has its servers housed in a multi-tenant data center with oversubscribed power infrastructure. The target data center includes one or more sets of modular “UPS→PDU” power paths (possibly with redundancies). The attacker leases a certain amount of power capacity (e.g., at a monthly rate of U.S.150\$/kW) and shares one such power path with several other benign tenants. It also installs several microphones on its server covers and/or rack assemblies.

Liberties and limitations of the attacker. We now discuss what the attacker can and cannot do in our threat model. For power attacks, the attacker can peak its power usage quickly by launching CPU intensive tasks. More importantly, the attacker launches power attacks by maliciously timing its peak power usage within the operator’s contractual constraint: *the attacker poses as a normal tenant, but it intentionally creates power emergencies by peaking its power usage when benign tenants’ power usage is also high.*

There may exist other types of attacks, such as igniting explosive devices, physically tampering with the data center infrastructures, and modifying server power supply units to create power surges beyond the attacker’s leased capacity (which will first trip the attacker’s rack-level circuit breakers and isolate the attacker from other tenants). These are all beyond our scope. Moreover, attacking the (possibly shared) network infrastructures are well-studied threats [28, 29] and also orthogonal to our study.

Finally, the attacker may create multiple tenant accounts (i.e., sub-attackers), each exploiting an acoustic side channel (Section 4.1) within a local range of a few meters to infer power usage of corresponding benign tenants. Nonetheless, we do not consider multiple attackers that belong to different and possibly competing entities, which is left as interesting future work.

Successful attack. We consider a power attack successful when $p_a + p_b \geq P^{cap}$ is satisfied for a continuous time window of at least L minutes ($L = 5$ minutes in our evaluation and enough for a circuit breaker to trip [30]), where p_a is the attacker’s power, p_b is the aggregate power of the benign tenants, and P^{cap} is the capacity of the shared power infrastructure under attacks. Accordingly, an *attack opportunity* is said to arise if there could be a successful power attack (i.e., the attacker’s peak power can result in a capacity overload for $L+$ minutes), regardless of whether the attacker actually launches an attack. Fig. 2 illustrates the attack opportunities in solid bars.

Note that a successful power attack may not always cause an outage; instead, *it compromises the data center availability* and, over a long term, the outage time in a multi-tenant data center significantly increases, resulting in million-dollar losses (Table 2 in Appendix A).

Motivations for attacker. Although geo-redundancy techniques may prevent certain advanced tenants’ service dis-continuity, a successful power attack, even in a single data center, can still lead to service outages for affected tenants and cost them million-dollar losses (see the recent example of JetBlue [31]). Meanwhile, tier classification is downgraded (e.g., a Tier-IV data center becomes a Tier-II one) due to infrastructure redundancy protection loss, effectively wasting the data center operator’s huge CapEx for achieving a high availability. Additionally, power outages significantly damage the operator’s business reputation. On the other hand, the attacker can create such severe

impacts by spending only a fraction of the resulting loss (3~23%) borne by the tenants and the operator. Thus, the attacker can be a competitor of the tenant(s) and/or the data center operator, or just any criminal organization creating havoc.

3.2 Challenges for Power Attacks

While multi-tenant data centers are vulnerable to power attacks, the actual attack opportunities are intermittent due to fluctuation of benign tenants' power usage.

Naturally, attack opportunities depend on benign tenants' aggregate power demand at runtime, which is unknown to the attacker. Additionally, the attacker does not have access to the operator's power meters to monitor tenants' power usage for billing purposes. The attacker might hack into the operator's power monitoring system to gain the power usage information, but this is safeguarded in the cyber space and orthogonal to our study.

A naive attacker may try to attack the data center without any knowledge of other tenants' power usage by simply maintaining its maximum power all the time. Nonetheless, this kind of power usage violates the operator's contractual requirement, leading to involuntary power cut and/or eviction. Alternatively, the attacker may try to launch random power attacks in hopes of capturing some attack opportunities. This, however, is also not effective and has a poor success rate (Fig. 16 in Section 5.2), since attack opportunities are intermittent.

The attacker may also refine its strategy by choosing a smaller window (e.g., anticipated peak hours) to launch attacks. Nonetheless, a successful power attack needs a precise timing due to the intermittency of attack opportunities, which cannot be located by simply zooming into a smaller time window in the order of hours. Alternatively, the attacker may launch attacks whenever it sees one of the power paths is down (due to equipment fault or maintenance shut-down). Again, intermittency of attack opportunities mandates precise timing for an attack to be successful. Moreover, detecting the loss of a power path requires a dual-corded connection, which may not apply in all data centers (e.g., a Tier-II data center) [6].

Limitation of the thermal side channel. In order to achieve a *precise* timing for power attacks, a recent study [22, 23] has proposed to use a thermal side channel resulting from heat recirculation to estimate benign tenants' power. However, heat containment techniques that reduce (even eliminate) the thermal side channel are expected to be adopted widely in the modern data centers. In addition, exploiting the thermal side channel in [22] requires modeling the heat recirculation in the target data center. Although a low sensitivity to model errors is reported in [22], the attacker needs to know the data center layout to build the model and any changes in the layout (e.g., new tenants move in) will require remodeling. Last but not least, it may take >1 minute for the heat generated by distant servers to reach the attacker's temperature sensor, rendering the estimated benign tenants' power usage information possibly outdated. Thus, the thermal side channel may not be as widely applicable as desired by the attacker.

In summary, a key challenge faced by the attacker is how to precisely time its power attacks at the moments when the benign tenants' aggregate power demand is also high.

4 EXPLOITING AN ACOUSTIC SIDE CHANNEL

A key observation we make in this paper is that there exists an acoustic side channel which results from servers' cooling fan noise and carries information of benign tenants' power usage at runtime. In this section, we first show through experiments on commercial servers how the noise is generated and its relation to server power. Then, we present our approaches to address the following three challenges in order to exploit the acoustic side channel for timing power attacks in a practical multi-tenant data center environment.

- *How to filter out the air conditioner noise?*




Fig. 3. Noise tones created by rotating fan blades [32].

- How to estimate benign tenants' power from the mixed noise?
- How to detect real attack opportunities?

4.1 Discovering an Acoustic Side Channel

4.1.1 Theoretical Support. The main sources of noise in data center servers are cooling fans, hard drives with spinning disks, and electrical components such as capacitors and transformers [32]. However, the dominant noise comes from the cooling fans, which draw cold air from the data center room into servers.² The rotating blades in a server's cooling fans create pulsating variations in air pressure levels, thus generating high-pitched noise with frequency components that depends on the fan speed. The relationship between the noise major tone frequency and fan speed in RPM (revolutions per minute) is governed by: $\text{Frequency (Hz)} = \frac{1}{60} \times \text{Fan RPM} \times \text{Number of Blades}$. Fig. 3 illustrates how the rotating blades creates the noise tones [32]. The fans also generate broadband white noise due to the friction of airflow with the electrical components inside the server. Among other less significant noise sources, hard disks create low-pitched humming noise, while the transformers and capacitors create tapping noise due to mechanical stress caused by the alternating current.

More importantly, a server's fan speed increases with its power consumption, serving as a good indicator of the server power. In a server, most of the power consumption converts into heat, which needs to be removed through cold air flowing through the server to maintain the temperature of internal components below a safe operating temperature threshold. As data center rooms operate in a conditioned temperature with little to no variation [33, 34], the amount of heat carried away from a server is directly proportional to the cold air flow rate which, according to the fan law[35], is directly proportional to the fan speed for a given server. Hence, the relationship between server power consumption p and fan speed r can be approximated as follows $r \approx k_1 \cdot p$, where k_1 is a proportionality constant that depends on the server's airflow impedance, data center air density, operating temperature, among others. In addition, following the empirical formula presented in [36] for a two-dimensional passage (e.g., a server case), the noise signal energy resulting from fan rotations is proportional to the fifth power of the air-flow rate, which in turn is proportional to the fan speed as well as the server power. This gives us the following relation between the noise signal energy L and the server power (electricity) consumption p : $L \approx k_2 \cdot p^5$, where k_2 is a server-specific proportionality constant. Therefore, this relation provides us with a theoretical support that the server fan noise energy serves as a good side channel that can reveal the server power usage information.

4.1.2 Experimental Validation. We now run experiments on a set of four Dell PowerEdge 1U rack servers (a popular model used in data centers) to validate the relation between the server noise and power consumption. To minimize the disturbances from external sources, we put our servers

²Air cooling is dominant in multi-tenant data centers. Liquid cooling, i.e., using liquid inside a server to remove heat, is typically used in high-performance computing centers (a different type of data center [1]) due to their ultra-high power density.




Fig. 4. Inside of a Dell PowerEdge server and a cooling fan. The server’s cooling fan is a major noise source.




Fig. 5. The relation between a server’s cooling fan noise and its power consumption in the quiet lab environment. (a) Server power and cooling fan speed. (b) Noise spectrum. (c) Noise tones with two different server power levels.

in a quiet lab environment with the room temperature conditioned at 72°F. We vary the power consumption of the servers by running CPU intensive loads at different levels. We record the server noise using a Samson Meteor studio microphone (with a sampling rate of 8k/sec) placed in front of the server inlet. We also monitor the servers’ power consumption, fan speeds, inlet and exhaust air temperatures. Fig. 4 shows the picture of internal components of the server with a close-up picture of one cooling fan and the noise recording setup.

The first thing to notice is that there is an array of cooling fans in the server spanning the entire width. This type of fan placement facilitates cold airflow through all the components in the server and is widely used in today’s servers. By default, these fans are dynamically controlled by the server based on the temperature sensor readings from different internal components (e.g., CPU). In our servers, there are seven fans, which are regulated individually by Dell’s built-in fan control algorithm based on the need of the fan’s designated cooling zone inside the server to achieve an exhaust hot air temperature below a safety threshold [37].

Fig. 5(a) shows the server power consumption and one server’s cooling fan speed. Both individual fan speeds and the mean fan speed are shown. It can be seen that the mean fan speed closely matches the server power consumption, thereby corroborating that the server fan speed is a good side channel for server power consumption. Next, in Fig. 5(b) we show the recorded noise frequency spectrum based on FFT (Fast Fourier Transform) each having a 10-second window. We see that the effect of changes in server’s power is clearly visible: with a high power, the noise frequency components also increase and there are more high-frequency components. The lower frequency components are mainly due to the background noise in the lab. We further take two sample points in time from the frequency spectrum, representing high and low server power respectively, and show them in Fig. 5(c). We confirm that a high server power generates higher-intensity noise across the entire frequency spectrum (especially for high-frequency spectrum). Additionally, in the zoomed-in figure, we see that there are additional frequency components in the server noise




Fig. 6. (a) Sharp power change creates noise energy spike. (b) Relation between noise energy and server power.

between 400Hz and 500Hz. These frequency components clearly show the impact of changed fan speeds in the noise tone. Note that, because there are multiple fans that are separately controlled to run at different fan speeds (Fig. 5(a)), we do not see one single prominent tone in the server noise spectrum.

Relation between noise energy and server power. To quantify the volume of the server’s cooling fan noise, we use the notion of “**noise energy**” (or noise signal energy), which is the sum of the square of each frequency component after performing FFT on the recorded noise signal over a 10-second window. Equivalently, noise energy is also the same as the square of time-domain noise signal amplitudes over the same 10-second window due to the Parseval’s theorem. Note the difference between “noise energy” and “server power” that are frequently used in this paper: noise energy means the recorded noise signal energy over a certain time window (not the real energy and hence a scaler without units), whereas server power is the real power in our conventional notion.

In Fig. 6(a), we show that a sudden change in server power creates a noise energy spike, which then gradually slows down to a stable value as the server power stabilizes. This is due to the internal fan control algorithm used by the Dell servers³ for reacting to a sudden change in power and heat: the fans try to bring down the suddenly increased temperature to a safe range as quickly as possible, thus running at an exceedingly high speed and generating a noise energy spike. Note that such noise energy spikes may not represent a real attack opportunity and needs to be detected by the attacker in order to improve its timing accuracy (Section 5.2).

We also see from Fig. 6(a) that there is a time lag in the fan speed response. It is mainly because the fan control algorithm directly reacts to the server’s internal temperature change, rather than the server power change. Hence, the time lag is due to the temperature build-up time plus the fan reaction time (for increasing fan speed). The temperature build-up time depends on the magnitude of server power change, and a higher power change results in a quicker temperature increase. In our experiment, the time lag is about 60 seconds for the maximum power change from 230W to 845W, while in practice the time lag is shorter since server power often varies within a smaller range. Moreover, the generated fan noise almost instantly reaches the attacker’s servers due to the high speed of sound. In contrast, using the thermal side channel reported in [22], it can take around 100 seconds for the heat generated by the benign tenants’ servers to travel to the attacker’s server inlet. Thus, the acoustic side channel can reveal benign tenants’ server power consumption in a more timely manner than the thermal side channel.

In Fig. 6(b), we show the relationship between the noise energy and server power consumption. For this figure, we run the server at different power levels, each for 20 minutes to reach the steady state. We see that the noise energy increases exponentially with the increasing server power consumption, following the approximated relation: $\text{noise-energy} = 10^{-23} \cdot (\text{server-power})^{6.8}$. It

³Different server vendors may have different fan control algorithms, and hence the fans may react differently to power spikes, as demonstrated by our experiment with a set of SuperMicro servers (Appendix B) which do not exhibit fan speed spikes when the server power suddenly increases.




Fig. 7. Server noise and power consumption in our noisy data center. (a) Noise spectrum. (b) Cutoff frequency of high-pass filter. The ratio is based on the noise of 4kW and 2.8kW server power. (c) Noise energy and server power.

deviates slightly from the theoretical relationship because the server also has other weaker noise sources like capacitors.

To further validate the usage of a server’s cooling fan noise as an acoustic side channel, we run another experiment on a set of SuperMicro twin servers. The results are consistent with our experiments on Dell PowerEdge servers. More details are available in Appendix B.

To sum up, our experiment corroborates the theoretical investigation that a server’s cooling fan noise energy serves as a good side channel to indicate the server power consumption: a server’s cooling fan noise energy increases with its speed, which in turn increases with the server power consumption. Nonetheless, there exist multifaceted challenges to exploit the prominent acoustic side channel in a practical data center environment, thus motivating our studies in the subsequent sections.

4.2 Filtering Out CRAC’s Noise

We have yet to show the existence of our discovered acoustic side channel in a practical data center environment. For this purpose, we run the same experiment inside our school data center with our server rack consisting of 20 Dell PowerEdge servers. The details of our data center are provided in Section 5.1. We run the same stress in all the servers such that they all have the same power consumption (and hence similar fan speeds), and record the noise in front of the server inlet in the middle of the rack. In Fig. 7(a), we apply 10-second FFT on the recorded noise signal and show the frequency spectrum, from which we see that unlike in a quiet lab environment, the low-frequency background noise inside the data center overwhelms the sound spectrum and changes in servers’ cooling fan noise are hardly visible. We have also varied the blower setpoint of computer room air conditioner (CRAC), and found the same result. In fact, as shown in Fig. 7(b), if we do not filter out the low-frequency components, the recorded noise energy is nearly the same (i.e., with a ratio close to 1) even though we vary the server power significantly (2.8kW versus 4kW); nonetheless, even for fewer servers, the distinction in noise energy at two different power levels is very clear in a quiet lab environment (Fig. 6(b)).

While some low-frequency components of the background noise come from servers owned by others, their impact is relatively insignificant since most of the servers in our experimental data center are idle with minimum fan speeds; instead, most of the low-frequency noise comes from the CRAC that provides cold air to the servers through large blowers (e.g., fans). Fortunately, the CRAC fan noise has different frequency tones from servers: a majority of the CRAC noise is within a lower frequency range compared to servers’ cooling fan noise, because the blower fan speed in a CRAC is often smaller than a server’s cooling fan. Therefore, *if we apply a high-pass filter to remove low-frequency components in the recorded noise, the CRAC’s noise impact may be mitigated.*

We validate the idea of using a high-pass filter for recovering the acoustic side channel and investigate the effect of the cutoff frequency of the high-pass filter in Fig. 7(b). We specifically look at the ratio of noise energy given a high server power (4kW) to that given a low server power

(2.8kW). A higher ratio means that the noise energies between high server power and low server power are more different and hence more distinguishable (i.e., a better acoustic side channel). We also show the absolute noise energy recorded in the high server power case. We see that the ratio sharply rises up to 200Hz and remains around 1.7, while the absolute sound energy decreases significantly. While a high ratio is desirable to have a larger variation in the noise energy as the server power changes, a too low absolute noise energy is not effective since we need to have a detectable noise trace. Here, we choose 200Hz as the cutoff frequency for the high-pass filter, which gives a high ratio of noise energy given different server power levels and also a moderately high absolute noise energy. In general, the choice of “optimal” cutoff frequency that yields the best timing accuracy may vary with data centers and CRACs. However, in our experiments in Section 5.2, we find that the timing accuracy is not significantly affected over a wide range of cutoff frequencies (200Hz–600Hz), demonstrating a good robustness of our filtering approach. In Fig. 7(c), we further show the server power and filtered noise energy in the data center. We see that after filtering out the low-frequency components (mostly due to the CRAC noise), the recorded noise energy closely follows the changes in server power consumption, although the noise energy variation is not as sharp as in a quiet lab environment (shown in Fig. 6(b)) and there are more random disturbances (addressed in Section 4.4).

In addition, we collect recordings from an online source [38] which provides background noise of a data center room. Then, we mix the collected data center noise with our own server noise recordings to emulate a scenario as if we were putting our servers in another data center. Although the “optimal” cutoff frequency is around 250Hz, our results confirm that using a high-pass filter can filter out undesired CRAC noise that would otherwise become dominant in a noisy data center environment. The details are available in Appendix C.

To conclude, *in a practical data center environment, the acoustic side channel can be recovered using a high-pass filter.* In later sections, all noise signals will pass through the high-pass filter unless otherwise stated.

4.3 Demixing Received Noise Energy

Up to this point, we have demonstrated an acoustic side channel resulting from servers’ cooling fan noise in a set of servers that have the same power consumption. This can be viewed as a set of correlated noise sources. Although the attacker may create multiple tenants throughout the data center room, it is not possible for the attacker to monitor each single noise source by placing a microphone near every rack. Thus, the noise recorded by the attacker’s microphones will have multiple nearby server racks’ noises mixed together (plus the CRAC noise which, as shown in Section 4.2, can be largely filtered out via a high-pass filter).

In the time domain, amplitudes of noise signals from different tenants vary rapidly and get constructed/destroyed over time at the attacker’s microphones. Nonetheless, statistically, there is little correlation between each other (as verified in our experiment and shown in Table 3 in Appendix D). Therefore, the noise energies generated by different tenants are additive at the attacker’s microphones and can be captured using a linear mixing model. *In what follows, we directly work on the energy of noise signals (with low-frequency components filtered out by a high-pass filter).*

4.3.1 Noise Energy Mixing Model. We consider a time-slotted model where each time slot lasts for 10 seconds. The generated noise energy (after a high-pass filter) over one time slot is considered as one sample of noise energy signal in our mixing model. There are M microphones and N noise sources. Fig. 22 in the appendix illustrates the noise energy mixing process. Each server/rack can be a noise source, and the attenuation matrix $A = [a_{m,n}] \in \mathbb{R}_+^{M \times N}$ includes the attenuation coefficient




Fig. 8. Illustration: NMF converts the 15 noise sources into 3 consolidated sources.

of each path from a source n to a microphone m .⁴ The matrix $X = [x_{n,k}] \in \mathbb{R}_+^{N \times K}$ represents the noise energy generated by the sources over K time slots. $Y = [y_{m,k}] \in \mathbb{R}_+^{M \times K}$ is corresponding received noise energy in the microphones over K time slots, and $E = [e_{m,k}] \in \mathbb{R}_+^{M \times K}$ denotes random disturbing energy. Next, the noise energy mixing process can be expressed as $Y = AX + E$.

4.3.2 Noise Energy Demixing. The mixing model in Section 4.3.1 helps us understand how different noise sources impact the attacker’s microphones, but the model is *blind* to the attacker. Concretely, obtaining the attenuation matrix A is very challenging, if not impossible, because of the complex nature of acoustic transmission channels (e.g., reverberation effects) as well as equipment/obstacles in between. Moreover, even the number of noise sources (i.e., N) is unknown to the attacker.

In a *blind* environment, a naive strategy would be to simply look at the noise energy received at the attacker’s microphones and then launch attacks upon detecting a high received noise energy. We refer to this strategy as *microphone-based attack*. Nonetheless, this strategy is ineffective and would lead to a poor timing, because of the “*near-far*” effect: a noise source closer to the microphone will have a bigger impact on the noise energy received by the attacker than a more distant source, whereas under the microphone-based attack strategy, the attacker simply considers the entire data center environment as a *single* noise source without accounting for location differences of different servers/racks. Appendix F explains near-far effects in greater detail. We will further study the microphone-based attack and show its ineffectiveness in our evaluation section (Fig. 16).

To mitigate near-far effects, the attacker can *demix* its received noise energy into multiple sub-components, each representing a *consolidated* noise source (e.g., a set of servers generating correlated noise energies). While near-far effects can still exist within each consolidated noise source demixed by NMF, they tend to be less significant in general compared to those when viewing all the benign servers as a single source, which is attested to by our evaluation results under various settings. Hence, if a consolidated noise source has a high noise energy, then we see based on the acoustic side channel that the corresponding servers are likely to have a high power usage. Therefore, *if all or most of the consolidated noise sources have a high noise energy level, then it is likely that the aggregate power of benign tenants is also high and an attack opportunity arises.*

Our proposed noise energy demixing falls into the problem of *blind source separation* (BSS), which decomposes the received signals in a model-free manner [39]. Concretely, in our context, BSS can separate the mixed noise energy signals into multiple less-correlated components, each representing a consolidated noise source. We illustrate the key idea of BSS in Fig. 8, where the actual noise sources on the left side, mixed together in the data center, are demixed into several consolidated sources. Note that demixing is in general non-deterministic and hence, there is no “wrong” demixing.

⁴The attacker’s own noise energy can be excluded, as it is known to the attacker itself.

Among various BSS techniques, we choose to use affine NMF (non-negative matrix factorization) with sparsity constraint [40, 41]. NMF was introduced as a low-rank factorization technique and utilized in unsupervised learning of hidden features [42–44]. Note that the goal of our proposed NMF-based approach is *not* to group servers in such a way that matches exactly with the actual physical layout; instead, it is to *mitigate* the “near-far” effect, which would otherwise be more significant and lead to a bad timing accuracy for power attacks when viewing all the benign servers as a single source. In Appendix G, we explain in detail how NMF works and why it achieves a good timing accuracy for power attacks.

Concretely, the attacker obtains at time t the noise energy signals $y_t = [y_{1,t}, y_{2,t} \cdots y_{M,t}]^T$ through its M microphones (as before, all the noise signals have passed through a high-pass filter to filter out the CRAC noise), where T is the transpose operator. We use L to denote the number of consolidated noise energy signals $z_t = [z_{1,t}, z_{2,t} \cdots z_{L,t}]^T$, each representing the sum of a group of servers’ noise energy. The value of L is chosen by the attacker (e.g., usually $L < M$). The attenuation matrix for these L consolidated noise sources are $B = [b_{m,l}] \in \mathbb{R}^{M \times L}$. To apply NMF, the attacker has to collect enough signal samples in order to exploit the statistical attributes. Here, the attacker applies NMF over the past K samples, and we use the notations $Y_t = [y_{t-K+1}, y_{t-K} \cdots, y_t]$, $Z_t = [z_{t-K+1}, z_{t-K} \cdots, z_t]$, and $E_t = [e_{t-K+1}, e_{t-K} \cdots, e_t]$ with $e_t = [e_{1,t}, e_{2,t} \cdots e_{M,t}]^T$ being the random disturbances.

Formally, the problem at hand can be stated as: given Y_t and $Y_t = BZ_t + E_t$, the attacker *blindly* estimates Z_t without knowing the attenuation matrix B . For a better estimation, we impose a sparsity constraint on Z_t , i.e., Z_t becomes very sparse with non-zero elements only when the noise energy of consolidated sources is sufficiently high. This is good for our purpose, because the attacker only needs a good estimation for the high power (hence, high noise energy) periods. Thus, we rewrite $Y_t = BZ_t + E_t$ as $Y_t \cong B\tilde{Z}_t + B_o\mathbf{I}^T + E_t$, where $B_o \in \mathbb{R}^{M \times 1}$ is the static part in the received noise energy signals, \mathbf{I} is a $K \times 1$ unit vector, and \tilde{Z}_t is the sparse version of the consolidated noise energies Z_t . With the disturbing matrix E_t modeled as having i.i.d. white Gaussian entries, estimating \tilde{Z}_t and B can be formulated as a minimization problem with a Euclidean cost plus a sparsity target/regularization as follows:

$$F(B, \tilde{Z}_t, B_o) = \frac{1}{2} \| Y - \tilde{B}\tilde{Z}_t - B_o\mathbf{I}^T \|_F^2 + \lambda \sum_{l,k} \tilde{z}_{l,k} \quad (1)$$

subject to all entries in B , \tilde{Z}_t , B_o being non-negative. In (1), note that $\|\cdot\|_F$ is the Frobenius norm, $\tilde{B} = \left[\frac{B_1}{\|B_1\|}, \frac{B_2}{\|B_2\|} \cdots \frac{B_L}{\|B_L\|} \right]$ in which B_l is the l -th column of B , and $\lambda \geq 0$ is a weight parameter that controls the degree of sparsity. Note that the column normalization of B is to make sure the sparsity constraint does not become irrelevant in the cost function: since the sparsity part of the cost function (1) is strictly increasing, B needs to be normalized after every update; otherwise, the solution can lead to very high values of B and small values of \tilde{Z}_t [45].

The objective function in (1) is not jointly convex in B , \tilde{Z}_t , and B_o . Thus, we use alternating least squares (ALS) with gradient descent and derive the following multiplicative update rules in a compact matrix form based on [40, 41]:

$$\tilde{Z}_t \leftarrow \tilde{Z}_t \odot \frac{\tilde{B}^T(Y - B_o\mathbf{I}^T)}{\tilde{B}^T Y + \lambda + \epsilon} \quad (2)$$

$$B \leftarrow \tilde{B} \odot \frac{(Y - B_o\mathbf{I}^T)\tilde{Z}_t^T}{\tilde{B}\tilde{Z}_t\tilde{Z}_t^T + \epsilon} \quad (3)$$

$$B_o \leftarrow B_o \odot \frac{\mathbf{I}^T Y}{\mathbf{I}^T(\tilde{B}\tilde{Z}_t + B_o\mathbf{I}^T)} \quad (4)$$




Fig. 9. State machine showing the attack strategy.

where ϵ is a small positive number added to the denominator to avoid division by zero, and \odot is the Hadamard (components-wise) product. The above update rules yield fast convergence to a (possibly local) optimum [41]. Note that reaching a unique global optimum is not guaranteed for NMF (due to number of unknown variables greater than the observations) and remains an open problem, which is beyond our scope.

4.4 Detecting Attack Opportunities

To detect periods of benign tenants' high power usage based on estimates of noise energy generated by different (consolidated) sources, we propose an online estimation process as well as a state machine to guide its attacks, as explained below.

4.4.1 Online Noise Energy Dimixing. At each time t , the attacker performs NMF once over its received noise energies over the past K time slots. While the sparse version of noise energy \tilde{z}_t throughout the entire look-back window with K samples gets demixed (as shown in (1)), only the latest demixed value \tilde{z}_t is useful and employed by the attacker for detecting attack opportunities. Since the attacker does not know which set of racks correspond to which consolidated noise source and the scale of \tilde{z}_t is not preserved in NMF [40, 41], we need re-scaling to make use of \tilde{z}_t . In our study, re-scaling is done to ensure that, for each consolidated source n , the sparse version of demixed noise energy has a normalized average value of 0.5 throughout the entire look-back window: $\frac{1}{K} \sum_{k=t-K+1}^t \tilde{z}_{l,k} = 0.5$, for all $l = 1, 2 \dots L$.

Finally, we note that the NMF itself converges very quickly due to its matrix-based update rules and hence produces a fast online estimation for the attacker. In our evaluation, demixing noise energy over a 12-hour look-back window takes less than a second in Matlab run on a typical desktop computer.

4.4.2 Launching Attacks. Although demixed noise energy is re-scaled (to have the same average for all consolidated sources), an attack opportunity is more likely to arise if the latest demixed noise energy is high for all consolidated sources. Thus, we propose a threshold-based attack strategy in which the noise energy demixed online is continuously fed to the attacker for detecting attack opportunities. Specifically, at time slot t , the attacker considers there is an attack opportunity when $E_{est} > E_{th}$, where $E_{est} = \sum_{l=1}^L \tilde{z}_{l,t}$ is the aggregate estimated noise energy (i.e., sum of the latest normalized demixed noise energy of all consolidated sources, after re-scaling as discussed in Section 4.4.1) and E_{th} is the attack triggering threshold. The attacker may tune the threshold E_{th} at runtime based on how often it can launch attacks (e.g., lower E_{th} to launch more attacks and vice versa).

In addition, to avoid attacking transient noise energy spikes caused by a sudden change in server power (Fig. 6(a)), the attacker waits for T_{wait} minutes before launching an attack if the rate of change in E_{est} across two consecutive time slots is higher than a preset threshold Δ_{th} .

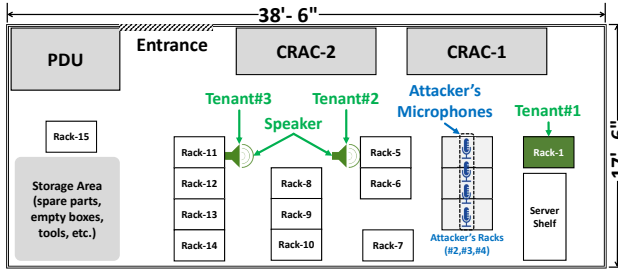


Fig. 10. Layout of our data center and experiment setup.

As per the operator’s contract, the attacker can keep its power high for only T_{attack} minutes each time. If E_{est} falls below E_{th} during an attack (i.e., before T_{attack} expires), instead of reducing power immediately, the attacker waits for T_{resume} minutes to see if E_{est} becomes high again. After each attack, the attacker waits for T_{hold} minutes before re-launching an attack.

We further illustrate the attack strategy using a state machine in Fig. 9, where the shaded boxes are the attack states (i.e., the attacker uses its full power). Using this strategy, the attacker can already achieve a reasonably high attack success rate ($\sim 50\%$, comparable to the best-known value [23]), although more sophisticated attack strategies can be interesting future work.

5 EVALUATION

We run experiments in a real data center environment and conduct simulations, in order to evaluate the effectiveness of our discovered acoustic channel and study how well it can assist an attacker with timing power attacks in a multi-tenant data center. The evaluation results show that, by using the NMF-based noise energy demixing and attack strategy proposed in Section 4.4, the attacker can detect 54% of all attack opportunities with a 48% precision, representing state-of-the-art timing accuracy.

5.1 Methodology

We conduct experiments in a real data center located on our university campus. The data center has 14 server racks mainly for archive and research purposes. These servers are owned by different research groups and idle nearly all of the time (as shown by the PDU reading). As we do not control these server racks except for our own, these racks are excluded from our experiment and they mainly generate background noises to provide us with a real data center environment (e.g., some servers housed together with the attacker may have almost no power variations). Note that the effective range of an acoustic side channel is only a few meters in practice and that the noise generated by servers far away from the attacker’s microphones is mostly viewed as background noise. Thus, if it tries to launch power attacks in a large data center room, the attacker may create multiple tenant accounts (i.e., sub-attackers), each exploiting a local acoustic side channel.

In our experiment, we consider three benign tenants and one attacker sharing the same “UPS-PDU” power distribution path as illustrated in Fig. 1. Due to limited server racks under our control, we use speakers that can reliably reproduce the server noise (see Appendix H) as two of the benign tenants’ server noise sources (#2 and #3). Our own server rack is used as another benign tenant (#1), while we consider another rack as the attacker and place four microphones to record the noise. Using the same setup as in Section 4.1.2, we record two 24-hour server noise traces in our quiet lab space and use them to emulate two benign tenants. We place the two speakers playing the two noise traces inside the data center, along with our server rack. The speaker locations mimic

tenants' location inside a real data center. We show the data center layout with the locations of the speakers, our server rack, and the attacker's microphones in Fig. 10. While scaled-down, our data center captures the acoustic environment of a real data center. Even in a large multi-tenant data center, since the acoustic side channel typically spans up to 10+ meters, the attacker needs to create multiple tenant accounts, each exploiting the side channel to detect high power usage of benign tenants within a local range. Moreover, our experimental data center has a similar size with edge multi-tenant data centers that are quickly emerging to accommodate the growing demand of Internet of Things applications [46]. Finally, we also test our timing approach in larger data centers by using online sources of data center noise (Appendix C) and simulating a different data center layout (Appendix M).

For each noise energy demixing (Section 4.4.1), we use the past 12-hour noise recording. Thus, although we run the experiment for 24 hours, there is no noise energy demixing or attack opportunity detection within the first 12 hours, and our figures only show results for the second 12 hours.

Tenant sizes. Due to equipment constraints, we perform scaled-down experiments as in prior studies [11, 27]. Our own server rack (i.e., the first tenant in our experiment) consists of 20 servers and has a maximum total power of 4.4 kW. We amplify the server noise played in the speakers to have a noise level comparable to an actual rack. With a maximum amplification of the speakers' volume, we get roughly five times the noise energy of the original recording inside our office space. With this scaling factor of five, tenants #2 and #3 each have an equivalent size of 4.5 kW (similar as our server rack or tenant #1), while the attacker's size is 2.2 kW. The total capacity of the power infrastructure with the three benign tenants and one attacker is 13kW, while the total subscribed capacity is 15.6kW due to 120% oversubscription. Thus, the three benign tenants under consideration and the attacker occupy about 86% and 14% of total subscribed capacities, respectively.

Power trace. We use four different power traces for the three benign tenants and the attacker. Two of the benign tenants' traces are taken from Facebook and Baidu production clusters [10, 26]. For the other two traces, we use the request-level logs of two batch workloads (SHARCNET and RICC logs) from [47, 48] and convert them into power consumption traces using real power models [49]. All the benign tenants' and the attacker's power traces are scaled to have 75% and 65% average power capacity utilization, respectively. Fig. 11 shows the aggregate power trace of all four tenants. Instead of considering that the attacker only consumes power during attacks, we use a real-world power trace for the attacker since an actual attacker would like to have a power consumption pattern similar to the benign tenants in order to stay stealthy.

Extended simulation. To evaluate the effectiveness of our discovered acoustic side channel and proposed attack strategies under different scenarios, we perform a year-long simulation by extending the 24-hour experiment in the data center. The key point in the extended simulation is to preserve the noise mixing process of different sources inside the data center. For this, we divide the 24-hour microphone recordings into 48 half-hour pieces. We then create the 1-year microphone trace by combining the 48 pieces in a random order. Overall, there exist attack opportunities for 6.7% of the times. We verify the extension process using a real experiment that is detailed in Appendix J.

Other settings. The recording mode of the microphones is set to 16-bit mono with a sampling rate of 8kHz. We do not use higher sampling rates for the recording since most of the frequency components of interest are well below 4kHz. The attacker uses the attack strategy described in Section 4.4 with $T_{wait} = 2$ minutes, $T_{attack} = 10$ minutes, $T_{hold} = 10$ minutes and $T_{resume} = 2$ minutes. In the default case, the attacker does not attack more than 7.5% of the time.





Fig. 11. Illustration of power attacks.

Fig. 12. Impact of attack triggering threshold E_{th} . The legend “Attack Opportunity” means the percentage of times an attack opportunity exists.

5.2 Results

We first show the results from our experiment inside the data center, and then present statistics of timing accuracy based on the acoustic side channel over a 1-year extended simulation. Our results highlight that the acoustic side channel is prominent for the attacker to launch well-timed attacks.

Demonstration of power attacks. The microphone recordings inside the data centers are converted to noise energy traces after filtering out frequencies lower than 200Hz. A snapshot of these traces is shown in Appendix I. These noise energy traces are applied in NMF to demix the noise energy traces, based on which the attacker detects the periods of benign tenant’s high power usage and launches power attacks. Fig. 11 shows the timing of power attacks by exploiting the acoustic side channel. We see that there are two attack opportunity windows, one between hours 6 and 7, and the other between hours 9 and 10. The attacker launches two successful attacks in the attack windows. However, it also launches unsuccessful attacks due to false alarms when the benign tenants’ actual power consumption is low. Further, there is one short-duration overload around hour 9, but this is deemed unsuccessful because it does not last long enough to reach our threshold of 5 minutes (to consider an attack successful).

Detection statistics. We look into two important metrics to evaluate the timing accuracy: *true positive rate* and *precision*. The true positive rate measures the percentage of attack opportunities that are successfully detected by the attacker, while precision measures the percentage of successful attacks among all the launched attacks. The attacker would seek to have both high true positive rate and high precision. Naturally, how often the attacker would launch attacks depends on the attack triggering threshold. Fig. 12(a) shows that, with a lower threshold, the attacker launches more attacks but many of them are unsuccessful, while with a higher triggering threshold, the attacker launches fewer attacks but with a better precision (since it launches attacks only when it is quite certain about an attack opportunity).

Fig. 12(b) shows the resulting true positive rate and precision with different triggering levels. We see that the precision goes up as the triggering threshold increases, while the true positive rate slightly goes down (because the attacker launches attacks less frequently). In our default setting, we use a triggering threshold $E_{th} = 2$, which results in attacks for a little over 7% of times with 54% true positive rate and 48% precision. An attacker would decide its triggering threshold mostly based on how often it is allowed to use its full capacity (i.e., launch attacks). Attacking too frequently can result in contract violation/eviction.





Fig. 13. Impact of high-pass filter cutoff frequency.

Fig. 14. Impact of attacker size. “ $x\%$ ” in the legend means the attacker subscribes $x\%$ of the total subscribed capacity.

Impact of high-pass filter cutoff frequencies. Applying a high-pass filter is crucial to filter out undesired CRAC noise while preserving the desired server noise. Thus, an important parameter to set is the cutoff frequency, below which the frequency components will be eliminated from the recorded noise. We vary the cutoff frequency and show the corresponding detection statistics in Fig. 13, while using the default settings for other parameters as specified in Section 5.1. The results show that the true positive and precision rates are relatively insensitive to the choice of cutoff frequencies within a fairly wide range of interest.

Impact of attacker size. The attacker launches power attacks by increasing its power to the maximum subscribed capacity. Hence, the subscription amount (i.e., the size of the attacker) plays an important role. Specifically, a larger attacker is more capable of launching a harmful attack, as it can cause a larger increase in the aggregate power. In Fig. 14(a), we study the impact of attacker size on available attack opportunity, true positive rate, and precision. Here, we increase the attacker’s capacity while keeping the benign tenants’ capacities fixed. We also scale the infrastructure capacity to have a default 120% oversubscription. While we vary the attacker size, we keep the attack percentage at our default 7.5%.

As expected, we see that a larger attacker results in more attack opportunities. The precision for a larger attacker is also higher, while the true positive rate goes down. This is because we keep the percentage of attacking time fixed at 7.5% while increasing the attacker size to have more attack opportunities, effectively reducing the true positive rate.

While a larger attacker can launch power attacks with a better precision, it also incurs a higher cost due to its increased footprint in the data center. Thus, we are interested in looking into the impact of different attacker sizes on the corresponding cost for data centers with different tiers. In Fig. 14(b), we show the cost impact of the power attacks with different attacker sizes in a 1MW 10,000 sqft data center, assuming the same detection statistics as in Fig. 14(a). The attacker’s cost includes the data center rent (150\$/month/kW), server purchase (\$1500/server/3-years), and electricity bill (\$0.1/kWh). On the other hand, the data center cost/loss due to power attacks and compromised availability is calculated using the method in Appendix A. We see that an attacker can create million-dollar losses, by spending only a fraction (between 3% and 23%) of the total cost borne by the data center and affected tenants.




Fig. 15. Without energy noise spike detection, the attacker launches many unsuccessful attacks.




Fig. 16. Detection statistics for different attack strategies.

Impact of noise energy spike detection. Here, we test the effectiveness of the noise energy spike detection mechanism in our attack strategy in Section 4.4.2. Concretely, we show in Fig. 15 that without the spike detection mechanism, the attacker launches power attacks (unsuccessful) upon detecting short-duration energy spikes (e.g., around hours 1, 6, 10 and 11) that possibly result from a fan RPM spike in case of a rapid server power change (Fig. 6(a)). In contrast, as shown in Fig. 11, the attacker can effectively avoid such unsuccessful attacks by waiting for the demixed energy noise to become stable.

Comparison with other attack strategies. We examine two alternatives: the simple *microphone-based power attack* (Section 4.3.2) and *peak-aware random power attacks (Random-P)*. In microphone-based power attacks, we take the average of the noise energy recorded by the four microphones (also with a high-pass filter applied) and compare it against an attack triggering threshold: if higher than the threshold, then attack. In Random-P, the attacker is assumed to have the knowledge of the probability of attack opportunities during different hours of a day, although this information is rarely available in practice. Then, the attacker distributes its total number of attacks (i.e., total amount of time it attacks) to each hour in proportion to the probability of attack opportunities during that hour. The details are presented in Appendix K. Fig. 16(a) shows the true positive rate at different percentages of attack times. We see that our proposed strategy significantly outperforms both microphone-based and random attack strategies. In Fig. 16(b), we see similar results for the precision of power attacks. Note that, while significantly worse than our proposed attack strategy, the microphone-based attack still has a reasonable true positive rate and precision. This is mainly because our experiment only has three benign tenants and hence the attenuation coefficients between different noise sources and the microphones do not differ very significantly.

Due to space limitations, more experimental results are deferred to the appendix, including detection statistics under an alternate power trace (Appendix L) and detection statistics under a different data center layout with different numbers of noise sources, microphones, and consolidated sources (Appendix M). Finally, we note that the prior research [22] discovers a thermal side channel resulting from heat recirculation in data centers with raised-floor designs where benign tenants' server heat can recirculate to the attacker's temperature sensors. It proposes a model-based estimation algorithm based on Kalman filter to infer benign tenants' runtime power usage and time power attacks, achieving a timing precision of about 50%. Nonetheless, our university data center does not use perforated tiles and has a completely different setup than the one considered in [22]; we cannot place a large number of heat sources to emulate heat recirculation in our university

data center either. Thus, we do not provide a side-by-side comparison with [22] in terms of timing accuracy, although our model-free approach achieves a comparable accuracy with the reported values in [22].

6 DEFENSE STRATEGIES

Given the danger of power attacks timed using the acoustic side channel, we briefly discuss possible defense strategies.

Power infrastructure resilience. Naturally, attack opportunities can be decreased by lessening the oversubscription ratio, but this comes at a significant revenue loss for the operator. Another approach is to increase power infrastructure resilience against power attacks. This, however, requires expensive infrastructure installation and/or upgrades, and more so for multi-tenant data center operators whose entire investment is the data center infrastructure. Additionally, tapping into stored energy (e.g., battery) during an attack may not mitigate cooling overloads, since the actual cooling load (i.e., server power) is not reduced (Section 2.1). In any case, an attacker can still launch timed attacks to compromise data center availability, albeit to a less severe extent.

Acoustic side channel. Power attacks can be potentially mitigated by weakening the acoustic side channel. Towards this end, low noise servers and/or noise-proof server racks can be used [50]. However, these are developed mainly to operate in small-scale (e.g., one/two racks) placed inside or in close proximity to office spaces, not for multi-tenant data centers. In addition, because of foam-sealed doors, special cooling arrangements are required for noise canceling server racks which are difficult to retrofit in multi-tenant data centers. Another approach is to destroy the correlation between server noise energy and server power consumption by running the cooling fans at maximum speed all the time. But, this results in a huge energy/cost wastage [51, 52], and is not decided by the data center operator due to its lack of control over tenants' servers. Finally, the data center operator may place noise-generating speakers, mimicking server noise to decrease the attacker's timing accuracy. However, it is still difficult to completely eliminate the acoustic side channel and de-correlate the attacker's received noise energy from the benign tenants' power usage by adding speaker-generated noise, because the attacker is stealthy and its microphone location is unknown to the data center operator.

Attack detection. A more proactive approach would be to find the malicious tenant(s). The data center operator can increase vigilance in its power monitoring to detect suspicious power usage patterns and pay close attention to that tenant. Then, the operator may take additional safety measures. Alternatively, the operator may revise its contract terms to prohibit certain power usage patterns that are likely malicious power attacks.

Server inspection. The attacker can conceal very small-sized or even invisible microphones (e.g., spying microphone) on its servers or racks that cannot be easily detected via visual inspection. Thus, on top of today's routine visual inspection, the data center operator may need to use advanced microphone detection devices to protect the data center facility against unauthorized listening devices.

Note that installing per-server circuit breakers (which are already available in some data centers) is not very helpful for defending against power attacks. The reason is that the attacker does not violate the operator's contractual constraint and hence, the operator cannot forcibly/arbitrarily cut power supply to any particular tenant (including the attacker) when power emergencies occur due to either coincident peaks or well-timed power attacks. While the data center operator may sign some contracts with certain tenants in advance for guaranteed power reduction in case of an emergency, it would have to pay a high reward to involved tenants.

In conclusion, the above defense strategies (or a combination) may improve data center infrastructure security in multi-tenant data centers. A more comprehensive investigation and evaluation of different defense strategies are warranted as future work.

7 RELATED WORK

Although common in data centers, oversubscription of power infrastructure requires power capping techniques to avoid outages, including CPU speed throttling [10, 53], workload migration [26], energy storage discharging [15, 54, 55], among others. These techniques, however, are inapplicable for multi-tenant data center operators due to their lack of control over tenants' servers or workloads.

While *cyber* security has long been a focus of research (e.g., mitigating distributed denial of service attacks [28, 29], and stealing private information through covert side channels [56]), power attacks to compromise data center physical security have also been recently demonstrated [14, 21, 23]. In particular, [14, 21] propose to use virtual machines (VMs) to create capacity overloads in cloud platforms. Nonetheless, VM-based attacks require co-residence of many malicious VMs to create prolonged harmful power spikes. Additionally, attackers do not directly control the power consumption of their launched VMs; instead, the cloud operator has many control knobs to migrate and/or throttle the VM power consumption across the cloud data center [57], safeguarding against VM-based power attacks. Our work, in contrast, focuses on multi-tenant data centers where a malicious attacker has its own physical servers, controls its server power consumption, and has the capacity to directly overload the shared power infrastructures.

Another recent study [23] exploits a thermal side channel for timing power attacks in a multi-tenant data center, whose limitations are discussed in Section 3.2. Compared to [23], we exploit a novel acoustic side channel which is more universally applicable, utilized using a model-free approach, and still produces a comparable (even better) timing accuracy.

Finally, our work adds to the recent literature on energy/power management in multi-tenant data centers and multi-tenant clouds. The recent works predominantly have been *efficiency*-driven, such as electricity cost reduction [2, 58, 59], improving infrastructure utilization [11], and reducing the cost of participation in utility demand response [60, 61]. In contrast, our work studies an under-explored *adversarial* setting in multi-tenant data centers.

8 CONCLUDING REMARKS

This paper studies power attacks in a multi-tenant data center. We discover a novel acoustic side channel that results from servers' cooling fans and helps the attacker precisely time its power attacks at the moments when benign tenants' power usage is high. In order to exploit the acoustic side channel, we propose to: (1) employ a high-pass filter to filter out the air conditioner's noise; (2) apply NMF to demix the received aggregate noise and detect periods of high power usage by benign tenants; and (3) design a state machine to guide power attacks. Our results show that the attacker can detect more than 50% attack opportunities, representing state-of-the-art timing accuracy.

ACKNOWLEDGEMENT

This work was supported in part by the U.S. NSF under grants CNS-1551661, CNS-1565474, and ECCS-1610471. We are grateful to Daniel Wong for providing SuperMicro servers during our experiment. We would also like to thank the anonymous reviewers and our shepherd, Sewoong Oh, for their valuable comments.

APPENDIX

A DATA CENTER AVAILABILITY AND OUTAGE COST UNDER POWER ATTACKS

Availability/Outage. We start with the data center availability for different tier specifications. Excluding cyber-related outages that are irrelevant to physical infrastructures, tier-specific data center availability is reported based on historical data [5]. Now, Tier-I availability statistics can be considered as the availability of a single set of power equipment, because Tier-I data center has no redundancy. For Tier-II and Tier-III data center, there is an outage when both the primary and redundant infrastructures fail. Then, we calculate the failure rate of redundant infrastructure for Tier-II and Tier-III as $p_f = \frac{1-p_a}{1-p_{a,I}}$, where p_a is the tier-specific availability without power attacks and $p_{a,I}$ is the availability of Tier-I data centers. Using these, we calculate the data center outage probability with 3.5% power attacks as “ $96.5\% \cdot (1 - p_a) + 3.5\% \cdot p_f$ ”. For a Tier-IV data center, as both primary and redundant infrastructures are expected to be fully independent [5, 25], we consider that they have the same failure probability calculate as $\sqrt{1 - p_a}$. Thus, with power attacks, the outage probability is “ $96.5\% \cdot (1 - p_a) + 3.5\% \cdot [2\sqrt{1 - p_a} - (1 - p_a)]$ ”. The availability statistics are shown in Table 1.

Table 1. Data center outage with power attacks.

Classification	Availability (%)	Outage (hours/Yr.)	Availability w/ Attack (%)	Outage w/ Attack (hours/Yr.)
Tier-I	99.671	28.82	96.182	334.41
Tier-II	99.741	22.69	96.995	263.26
Tier-III	99.982	1.58	99.791	18.3
Tier-IV	99.995	0.44	99.946	4.74

Outage cost. The outage cost can differ widely based on what kind of services are running in a data center. For estimation purposes, we use the cost of outage per square-foot per minute based on the Ponemon Institute’s survey report [16]. In addition to the outage cost due to increased downtime with power attacks, the intended availability target for tier classification is also lost, and hence the added CapEx to achieve a higher tier is essentially wasted. We calculate the capital loss by considering an infrastructure CapEx of \$10/W, \$11/W, \$20/W and \$22/W for Tier-I, Tier-II, Tier-III and Tier-IV data centers, respectively [7], amortized over 10 years. We downgrade a data center to a lower tier if the annual outage time becomes higher than that required for a certain tier. The outage costs induced by power attacks for different tiers are shown in Table 2.

Table 2. Cost impact of power attack 3.5% of the time on a 1MW-10,000 sqft data center.

Classification	Outage Cost (\$/hour/sqft)	Increased Outage Cost (mill. \$/Yr.)	Capital Loss (mill. \$/Yr.)	Total Cost (mill. \$/Yr.)
Tier-I	1.98	6	n/a	6
Tier-II	6.4	15.5	0.1	15.6
Tier-III	46.7	7.8	0.9	8.7
Tier-IV	55.6	2.4	1.1	3.5

B EXPERIMENT WITH A DIFFERENT VENDOR’S SERVERS

To corroborate our findings with the Dell PowerEdge servers, we borrow a set of SuperMicro twin servers from our colleague and conduct experiments to show the relation between the server power and noise signal energy. As a major server vendor, SuperMicro is ranked the fifth by NPD Group in




Fig. 17. Experiment setup with the SuperMicro twin server set. Picture #2 is taken from the manufacturer’s website, since opening the server case by ourselves would void the manufacturer warranty.

terms of market share [62]. The twin servers consist of two servers placed in parallel in a 1U server chassis. There are six cooling fans located in the front part of the twin server set, three fans for each server. We run the servers at different power levels and record the resulting server noise. The twin servers and experiment setup are shown in Fig. 17.




Fig. 18. Experimental results on SuperMicro servers. (a) Server power and noise energy. (b) Noise spectrum. (c) Noise tones with different power. (d) Power spike v.s. noise energy change.

We show the power trace and corresponding noise energy in Fig. 18(a). In particular, we can see similar results as in our Dell PowerEdge server experiments (Fig. 7(c)): the noise signal energy closely follows the server power. We also show the noise and frequency spectrum of the SuperMicro servers in Figs. 18(b) and 18(c), which match the results from our Dell servers and demonstrate that a higher server power consumption results in a higher frequency noise.

Note that SuperMicro has its proprietary algorithm for controlling fan speeds, which are different from the one used by Dell PowerEdge servers. For example, when there is a sharp change in server power, unlike Dell PowerEdge servers that start spinning fans at the maximum speed (possibly to avoid potential catastrophic consequences), we do not see a high noise energy spike for SuperMicro servers in Fig. 18(d); instead, SuperMicro servers tend to moderately increase the fan speed (hence noise energy) at first and, only after a sustained higher power, increase the fan speed to accommodate cooling needs. Nonetheless, the basic intuition still holds: with a higher server power, more cold air is needed, thus making the server run cooling fans at a higher speed and generate more noise. Although the absolute value of actual noise energy can vary with different server models, our attack strategy utilizes the relative noise levels (Section 4.4): with a higher noise, a consolidated noise source is more likely, albeit not guaranteed, to have a higher power consumption.

We also verify the high-pass filter approach in the SuperMicro servers by mixing the server noise with our data center CRAC noise. Fig. 19(a) shows the unfiltered noise energy, which does not reveal the servers’ power variation. In contrast, Fig. 19(b) shows the noise energy after a high-pass




Fig. 19. SuperMicro servers in data center. (a) Without the high-pass filter. (b) With the high-pass filter that has a cutoff frequency of 200Hz.

filter with a cutoff frequency of 200Hz is applied. We see that the high-pass filter suppresses the CRAC noise and the remaining noise energy closely matches the server power (like in the quiet lab environment shown in Fig. 18(a)).

C EXPERIMENT WITH DIFFERENT CRAC NOISE




Fig. 20. Impact of a high-pass filter on noise mixture. (a) Server power and noise energy. (b) Unfiltered noise mixture. (c) Filtered noise mixture. (d) Cutoff frequency for high-pass filter.

Since we currently do not have access to a commercial data center to test our filtering approach on different CRACs, we collect an one-hour online trace of data center noise from [38] and mix it with our own server noise recordings as an alternative to validate our filtering approach. While the online noise source [38] does not provide details of the data center from which the noise is recorded, we have compared the spectrum of its noise recordings with those from other sources such as YouTube and found that they are consistent. In particular, the collected noise is mixed with our recorded server noise at a ratio of 90% to 10%, i.e., 90% of the aggregate noise energy after mixing comes from the collected online noise while our server noise contributes to the remaining 10%. By doing so, the mixed noise emulates a scenario as if we were putting our servers in another data center (whose background noise is captured by the online noise trace) and placing a microphone nearby our servers to record their noise.

Our own server noise recording done in a lab environment and the corresponding power consumption are shown in Fig. 20(a) (these are also used as source#1 in our evaluation under default settings). Then, we repeat the collected one-hour online noise trace and make it have a length of four hours. Without filtering, the aggregate noise energy after mixing is shown in Fig. 20(b), where we can hardly see any noise pattern that changes with the server power consumption.

Next, we apply a high-pass filter with two different lower cutoff frequencies (200Hz and 600Hz) on the mixed noise and show the filtered noise energy in Fig. 20(c). We see that the filtered noise energy reveals the server power pattern exhibited in Fig. 20(a). Meanwhile, we can also observe a difference in noise energies between the two filtered signals under different cutoff frequencies. When a higher cutoff frequency is used, the absolute noise energy becomes lower since more low-frequency components are discarded. This is clearly reflected in Fig. 20(d), where the absolute noise energy is for a server power of 785W and the noise energy ratio is the ratio of the filtered

noise energy when the servers have a power consumption of 785W (high power) to that when the servers consume 620W (low power). When using the online noise source as the background noise, the noise energy ratio reduces (i.e., relatively less difference in filtered noise energy for two different server power consumptions) compared to the result in our university data center. This is because the collected online noise is recorded from a large data center that also includes many background servers in addition to CRACs. To get a good ratio while keeping a relatively higher absolute server noise energy, a cutoff frequency can be chosen between 200-500Hz, which is similar to the experiment in our university data center and hence demonstrates the practicality of our filtering approach under different CRACs.

D CORRELATION OF NOISE SIGNAL AMPLITUDES




Fig. 21. Snapshot of noise traces.

In Figs. 21(a) and 21(b), we show a 60-seconds sample of the recorded noise signals of the three different sources in the time domain, and a 12-hour sample of the noise energy, respectively. We also show the correlation of the time domain traces of the sources in Table 3. There is almost a zero correlation among the amplitudes of different noise sources in time domain, thus making their noise energy additive at the attacker's microphones.

Table 3. Correlation of Noise Signal Amplitudes.

	Source#1	Source#2	Source#3
Source#1	1.0000	0.0008	0.0001
Source#2	0.0008	1.0000	0.0009
Source#3	0.0001	0.0009	1.0000

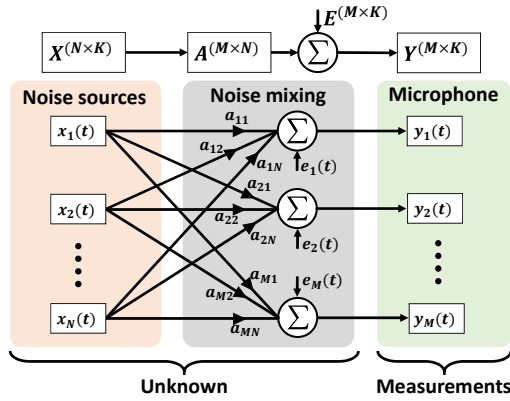


Fig. 22. Noise energy mixing process.

E NOISE MIXING MODEL IN DATA CENTER

We illustrate the server noise mixing model for N noise sources and M microphones in Fig. 22. The attenuation matrix $A = [a_{m,n}] \in \mathbb{R}_+^{M \times N}$ includes the attenuation coefficient of each path from a source n to a microphone m . The matrix $X = [x_{n,k}] \in \mathbb{R}_+^{N \times K}$ represents the noise energy generated by the sources over K time slots. $Y = [y_{m,k}] \in \mathbb{R}_+^{M \times K}$ is corresponding received noise energy in the microphones over K time slots, and $E = [e_{m,k}] \in \mathbb{R}_+^{M \times K}$ denotes random disturbing energy.

F INEFFECTIVENESS OF MICROPHONE-BASED ATTACKS




Fig. 23. Ineffectiveness of microphone-based attacks. The attacker receives a high noise energy at time slot 2, whereas the actual attack opportunity is at time slot 4.

We first illustrate near-far effects in Fig. 23 with three racks, with a normalized distance of 3, 2, and 1 to the attacker, respectively. Each rack is considered as a single source, and the racks have high power (and hence generate high noise energy) at different times. The attenuation coefficients are chosen following the inverse-square law for sound attenuation with distances [63]. We also normalize the attenuation coefficient for rack #3 as 1. Then, we see that the microphone recording has the highest noise energy at time slot 2, whereas the aggregate server power reveals that the time slot 4 has the highest power demand and hence represents an actual attack opportunity. Thus, the noise signal received by the attacker is dominated by the noise generated by nearby servers. Thus, even though the attacker receives a high noise energy, it is possibly due to that only nearby servers are consuming a high power and generating a high noise whereas the total power consumption of benign servers (including those that are located further away from the attacker) is still quite low.

As the microphone-based attack strategy considers the entire data center environment as a *single* noise source without accounting for near-far effects, it has a rather poor timing accuracy.

G INTERPRETATION OF NMF-BASED NOISE ENERGY DEMIXING

Without knowing the attenuation matrix $A = [a_{m,n}]$, we propose to use NMF to *blindly* decompose the received noise. To facilitate understanding, we consider the basic NMF without considering sparsity regularization. Specifically, NMF is aiming to produce a new attenuation matrix $B = [b_{m,l}] \in \mathbb{R}_+^{M \times L}$ as well as a new set of L *consolidated* noise sources which generate noise energy of $Z = [z_{l,k}] \in \mathbb{R}_+^{L \times K}$, such that $Y \approx B \cdot Z$, where L denotes the number of consolidated noise sources and is chosen by the attacker (typically $L < M$ [40, 41]). Generally, for satisfying $Y \approx B \cdot Z$, the matrix B and Z are not unique and, through an iterative process to solve (1) formulated in Section 4.3.2, NMF returns one such solution that yields a small square error.

The new attenuation coefficient $b_{m,l}$ returned by NMF indicates how much noise energy is received by the attacker's microphone m for each *unit* noise energy generated by the consolidated noise source l . Likewise, $z_{l,k}$ represents the *intensity* of consolidated noise source l at time k . Each column of the attenuation matrix B can be viewed as a *feature*, which represents how much noise energy is received by the attacker's microphones when the consolidated noise source l generates a unit amount of energy. By linearly combining these features, we can get $B \cdot Z$ to approximately denote the total received noise energy by the attacker's microphones. Therefore, we can see that NMF aims to describe the noise mixing process in a different way than the actual (unknown) mixing: L consolidated sources generate noise, which is then mixed at the attacker's microphones through a new attenuation process characterized by B .

As B and Z are not unique, the way that NMF identifies consolidated sources is not unique either (e.g., it is even possible that NMF considers half of an actual physical server as part of a consolidated source). As NMF involves an iterative optimization process, there is no simple way to determine which of the possible solutions B and Z will be returned by NMF [41]. Fig. 8 illustrates an example of consolidation result, where one server in the third row is grouped with the noise sources in the second row as a single consolidated source. *This is not considered as "wrong"*.

While one might think that NMF *non-deterministically* returns B and Z such that $Y \approx B \cdot Z$, NMF is still helpful for the attacker to time its power attacks. Specifically, NMF *mitigates* the near-far effects by dividing the benign servers into different consolidated sources which, although not guaranteed, tend to have correlated noise energy. Although near-far effects can still exist within each consolidated source (e.g., consolidated source #2 illustrated in Fig. 8), the effects are much less significant compared to viewing all the benign servers as a single group without differentiation. While it is possible that in some bad cases NMF considers a nearby server and the furthest server as an individual consolidated source, NMF is still better in most cases than viewing all the benign servers as a single source.

Therefore, although NMF does not (and also cannot) separate each physical noise source or completely eliminate near-far effects, NMF suffers from near-far effects to a much less extent than microphone-based attack strategy, thus improving the attacker's timing accuracy. Note that, had NMF been able to completely eliminate near-far effects by *correctly* separating each physical noise source, the attacker would be able to achieve a precision of nearly 100% (due to the good correlation between the server noise energy and power usage, demonstrated in Section 4.1.2) as if it were using a power meter. Naturally, this is not expected for any side channels in general.




Fig. 24. Server noise is reproduced by the speaker.

H NOISE PRODUCED BY SPEAKERS

We first run the workload trace in our servers and then record the noise trace. Then, we play the recorded server noise trace in a speaker and record the noise from the speaker. Next, we compare the filtered noise energy of the two recordings (original server noise and the speaker-reproduced noise) in Fig. 24. The experiment shows that our speakers can closely reproduce the actual server noise and hence validates our experimental setup.

I MICROPHONE TRACES




Fig. 25. Snapshot of recorded noise energy ($\geq 200\text{Hz}$).

The noise energy traces recorded by the microphones inside our data center is shown in Fig. 25. Frequencies lower than 200Hz are discarded in these traces.

J EXTENDED SIMULATION APPROACH




Fig. 26. Offline extended trace matches the actual experiment.

For our extended simulation, we take a 24-hour recorded noise trace and divide it into 48 pieces. The pieces are randomly added together to create a yearly extended trace. To cross validate that the approach preserves the acoustic environment of noise mixing in the data center, we also divide the actual 24-hour source server workload trace (which is used for the initial 24-hour microphone recording inside the data center) into 48 pieces, and run them on our servers in the same order as in the extended trace. We then compare our offline-generated extended noise trace with the actual noise traces generated by our servers over a 12-hour window. In Fig. 26, we see that our extended noise trace matches the actual recorded noise trace generated by servers, demonstrating that our extended trace conserves real noise mixing process in the data center.

K ATTACK DISTRIBUTION OF PEAK-AWARE RANDOM ATTACK (RANDOM-P)




Fig. 27. Probability distribution of attack opportunities, and distribution of peak-aware random attacks.

In Random-P, the attacker is assumed to have the knowledge of the probability of attack opportunities during different hours of a day, although this information is rarely available in practice. Then, the attacker distributes its total number of attacks (i.e., total amount of time it attacks) to each hour in proportion to the probability of attack opportunities during that hour, subject to the constraint of attacking for no more than 20% of the time during each hour (since only limited peak power usage is allowed for each tenant). For illustration of Random-P, we re-order the power traces since our original power traces (due to multiplexing of multiple power traces representing diverse workloads) do not exhibit a very clear peak-valley pattern during each day. Fig. 27(a) illustrates probability distribution of attack opportunity probability under the new trace, where “peak hours” are between 5:00 pm and 10:00 pm. We also show how the attacker sets its random attack probabilities throughout the day in Fig. 27(b), subject to the constraint that the attacker attacks no more than 10% on average as in our default case and no more than 20% for each hour.

L DETECTION STATISTICS UNDER AN ALTERNATE POWER TRACE




Fig. 28. Detection statistics for different attack strategies with an alternate trace.

We test NMF against a different power (and hence also noise) trace other than the one we used in our default evaluation. We divide the original power and noise traces into half-hour windows and re-order the trace windows to generate the new traces (similar to our extended experiment which we have verified using a shorter real experiment in Appendix J). All the other settings are the same as in our default case described in Section 5.1. We show the timing performance of NMF for the new power trace in Fig. 28. We see very similar results as in our default trace results where NMF significantly outperforms the random attacks and microphone-based attacks. These results indicate that our proposed NMF-based approach works for different power traces.

M DETECTION STATISTICS UNDER DIFFERENT SETTINGS

Here, we examine the robustness of our NMF-based approach against the number of consolidated sources as well as the number of available microphones. Because of the limited number of servers we have as noise sources and our experimental data center’s restrictions on the area we can access

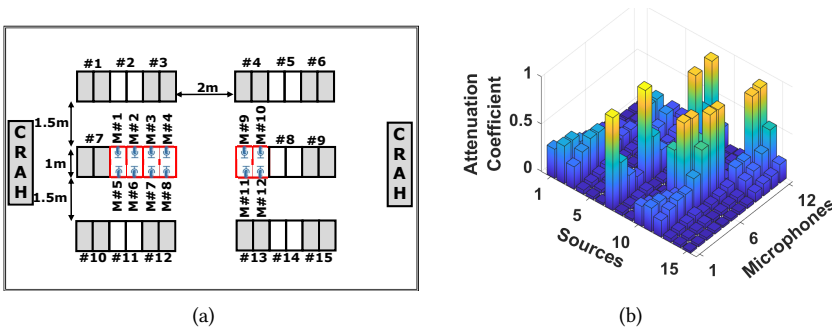


Fig. 29. Simulation settings for robustness study. (a) Simulated data center layout. (b) Attenuation modeled using the inverse square law.

to, we resort to a simulation-based approach. We show the simulated data center layout in Fig. 29(a) with 15 benign noise sources (each source consists of a pair of two adjacent racks, labeled as #1, #2, \dots , #15) and an attacker that has six racks with 2 microphones placed on each of its server racks. We follow the standard rack and row dimension for the layout. The benign noise sources are denoted by #xx, while the attacker’s microphones are denoted by M#xx, where xx is replaced by an indexing number. To determine how each source’s noise attenuates and propagates to the attacker’s microphones, we consider the inverse square law which states that noise attenuation is the inverse of the square of the distance between the noise source and the receiver [63]. We show the attenuation coefficient (i.e., the mixing matrix A) in Fig. 29(b), where the values are normalized to have a maximum of 1. With the attenuation matrix, we now can simulate the noise mixing process at the attacker’s 12 microphones from the 15 benign noise sources. We also add the background CRAC noise collected from the online source [38]. To generate the actual noise for the servers, we use the same technique we have applied in our extended simulation (Appendix J). We utilize our three 24-hour noise traces and randomly re-order half-hour windows to get the 15 source noises. Fig. 30 shows a sample power and noise trace.




Fig. 30. Power and noise energy traces of one of the offline generated sources.




Fig. 31. Detection statistics with different number of microphones. “ $L = 4/6$ ” denotes “4/6 consolidated sources when using NMF”.

We vary the number of microphone recordings by truncating the attenuation matrix and only keeping some of the 12 microphone recordings. Note that the number of consolidated sources L is

decided by the attacker, and that typically L is smaller than the number of microphones. We show the impact of available microphones and the number of consolidated sources in Fig. 31. While the actual quantitative values vary based on the specific settings, the qualitative results hold in all the examined cases: our NMF-based approach outperforms the random attacks and microphone-based attacks in terms of timing accuracy. The reason is that NMF is able to mitigate near-far effects by dividing all the benign servers into several consolidated sources with similar power/noise usage patterns. Our results demonstrate that NMF applies in a variety of settings.

REFERENCES

- [1] NRDC, “Scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers,” *Issue Paper*, Aug. 2014.
- [2] M. A. Islam, H. Mahmud, S. Ren, and X. Wang, “Paying to save: Reducing cost of colocation data center via rewards,” in *HPCA*, 2015.
- [3] “Colocation market - worldwide market forecast and analysis (2013 - 2018),” <http://www.marketsandmarkets.com/ResearchInsight/colocation.asp>.
- [4] Apple, “Environmental responsibility report,” 2016.
- [5] Colocation America, “Data center standards (Tiers I-IV),” 2017, <https://www.colocationamerica.com/data-center/tier-standards-overview.htm>.
- [6] Telecommunications Industry Association, “Data center standards overview,” *TIA 942*, 2005 (amended in 2014).
- [7] W. P. Turner, J. H. Seader, and K. G. Brill, “Tier classifications define site infrastructure performance,” *Uptime Institute White Paper 17*, 2006.
- [8] S. Pelley, D. Meisner, P. Zandevakili, T. F. Wenisch, and J. Underwood, “Power routing: Dynamic power provisioning in the data center,” in *ASPLOS*, 2010.
- [9] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel, “The cost of a cloud: Research problems in data center networks,” *SIGCOMM Comput. Commun. Rev.*, vol. 39, Dec. 2008.
- [10] 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.
- [11] 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.
- [12] Hornbaker Group, “Determining kilowatt capacity of data center space,” <http://www.hornbakergroup.com/pdf/Considerations-when-leasing-Data-Center-space-by-the-kilowatt.pdf>.
- [13] United States District Court, “Layton v. Terremark North America, LLC,” 2014.
- [14] C. Li, Z. Wang, X. Hou, H. Chen, X. Liang, and M. Guo, “Power attack defense: Securing battery-backed data centers,” in *ISCA*, 2016.
- [15] S. Govindan, D. Wang, A. Sivasubramaniam, and B. Urgaonkar, “Leveraging stored energy for handling power emergencies in aggressively provisioned datacenters,” in *ASPLOS*, 2012.
- [16] Ponemon Institute, “2016 cost of data center outages,” 2016, <http://goo.gl/6mBFTV>.
- [17] Emerson Network Power, “Addressing the leading root causes of downtime,” 2013, <http://goo.gl/b14XaF>.
- [18] Reuters, “British Airways \$100M outage was caused by worker pulling wrong plug,” Jun. 02 2017.
- [19] 365DataCenters, “Master services agreement,” <http://www.365datacenters.com/master-services-agreement/>.
- [20] Internap, “Colocation services and SLA,” <http://www.internap.com/internap/wp-content/uploads/2014/06/Attachment-3-Colocation-Services-SLA.pdf>.
- [21] Z. Xu, H. Wang, Z. Xu, and X. Wang, “Power attack: An increasing threat to data centers,” in *NDSS*, 2014.
- [22] M. A. Islam, S. Ren, and A. Wierman, “Exploiting a thermal side channel for power attacks in multi-tenant data centers,” in *CCS*, 2017.
- [23] M. A. Islam, S. Ren, and A. Wierman, “A first look at power attacks in multi-tenant data centers,” in *GreenMetrics*, 2017.
- [24] Mohammad A. Islam, “Server noise trace,” https://sites.google.com/site/mdatiqislam1985/server_noise_trace.
- [25] Uptime Institute, “Tier certifications,” <https://uptimeinstitute.com/TierCertification/>.
- [26] 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.
- [27] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, and C. Hyser, “Renewable and cooling aware workload management for sustainable data centers,” in *SIGMETRICS*, 2012.
- [28] 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.

- [29] 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.
- [30] Raritan, "Data center power overload protection," *White Paper*, 2016.
- [31] Y. Sverdlid, "Verizon data center outage delays JetBlue flights," in *DataCenterKnowledge*, January 2016.
- [32] C. E. P. Dell, "Dell enterprise acoustics," 2011, <https://www.dell.com/downloads/global/products/pedge/en/acoustical-education-dell-enterprise-white-paper.pdf>.
- [33] I. Manousakis, I. n. Goiri, S. Sankar, T. D. Nguyen, and R. Bianchini, "Coolprovision: Underprovisioning datacenter cooling," in *SoCC*, 2015.
- [34] D. L. Moss, "Dynamic control optimizes facility airflow delivery," *Dell White Paper*, March 2012.
- [35] The New York Blower Company, "Fan laws and system curves," <http://www.nyb.com/pdf/Catalog/Letters/EL-02.pdf>.
- [36] R. H. Lyon and A. E. Bergles, "Noise and cooling in electronics packages," *IEEE Transactions on Components and Packaging Technologies*, vol. 29, no. 3, pp. 535–542, 2006.
- [37] Dell Product Group - Server Engineering, "Cooling options for thermal control in dell poweredge servers," 2015, http://en.community.dell.com/techcenter/extras/m/white_papers/20441060/download.
- [38] myNoise: Custom Background Noise Machines, "Data center server room noise generator," <https://mynoise.net/NoiseMachines/dataCenterNoiseGenerator.php>.
- [39] V. Zarzoso and A. Nandi, "Blind source separation," in *Blind Estimation Using Higher-Order Statistics*, pp. 167–252, Springer, 1999.
- [40] H. Laurberg and L. K. Hansen, "On Affine Non-negative Matrix Factorization," in *ICASSP*, 2007.
- [41] J. Eggert and E. Korner, "Sparse coding and NMF," in *IJCNN*, 2004.
- [42] P. Paatero and U. Tapper, "Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values," *Environmetrics*, vol. 5, no. 2, pp. 111–126, 1994.
- [43] D. D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," in *NIPS*, 2001.
- [44] D. D. Lee and H. S. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, no. 6755, p. 788, 1999.
- [45] P. O. Hoyer, "Non-negative sparse coding," in *NNSP*, 2002.
- [46] EdgeConnex, <http://www.edgeconnex.com/>.
- [47] D. G. Feitelson, D. Tsafir, and D. Krakov, "Experience with using the parallel workloads archive," *Journal of Parallel and Distributed Computing*, vol. 74, no. 10, pp. 2967–2982, 2014.
- [48] Parallel Workloads Archive, <http://www.cs.huji.ac.il/labs/parallel/workload/>.
- [49] X. Fan, W.-D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," in *ISCA*, 2007.
- [50] NENS, "How to reduce the noise from your servers," 2017, <https://www.nens.com/reduce-noise-servers/>.
- [51] S. Li, T. Abdelzaher, and M. Yuan, "Tapa: Temperature aware power allocation in data center with map-reduce," in *IGCC*, 2011.
- [52] Z. Wang, C. Bash, N. Tolia, M. Marwah, X. Zhu, and P. Ranganathan, "Optimal fan speed control for thermal management of servers," in *InterPACK*, (Berkeley, CA, USA), 2009.
- [53] L. Li, W. Zheng, X. D. Wang, and X. Wang, "Coordinating liquid and free air cooling with workload allocation for data center power minimization," in *ICAC*, 2014.
- [54] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy, "Energy storage in datacenters: what, where, and how much?," in *SIGMETRICS*, 2012.
- [55] 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.
- [56] M. Guri, Y. A. Solewicz, A. Daidakulov, and Y. Elovici, "Fansmitter: Acoustic data exfiltration from (speakerless) air-gapped computers," *CoRR*, vol. abs/1606.05915, 2016.
- [57] S.-J. Moon, V. Sekar, and M. K. Reiter, "Nomad: Mitigating arbitrary cloud side channels via provider-assisted migration," in *CCS*, 2015.
- [58] 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.
- [59] 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.
- [60] N. Chen, X. Ren, S. Ren, and A. Wierman, "Greening multi-tenant data center demand response," in *IFIP Performance*, 2015.
- [61] Z. Liu, I. Liu, S. Low, and A. Wierman, "Pricing data center demand response," in *SIGMETRICS*, 2014.
- [62] CRN, "Npd group: Top 8 server brands of 2016 q2," <http://www.crn.com/slide-shows/storage/300081644/npd-group-top-8-server-brands-of-2016-q2.htm>.
- [63] HyperPhysics, "Inverse square law, sound," <http://hyperphysics.phy-astr.gsu.edu/hbase/Acoustic/invsqs.html>.