Enabling Low-Cost Server Level Power Monitoring in Data Centers Using Conducted EMI*

Pranjol Sen Gupta
The University of Texas at Arlington
Arlington, Texas, USA
pranjolsen.gupta@mavs.uta.edu

Zahidur Talukder
The University of Texas at Arlington
Arlington, Texas, USA
zahidurrahim.talukder@mavs.uta.edu

Tasnim Azad Abir
The University of Texas at Arlington
Arlington, Texas, USA
tasnimazad.abir@uta.edu

Phuc Nguyen
University of Massachusetts Amherst
Amherst, Massachusetts, USA
vp.nguyen@cs.umass.edu

Mohammad A. Islam
The University of Texas at Arlington
Arlington, Texas, USA
mislam@uta.edu

ABSTRACT

Server-level power monitoring in data centers can significantly contribute to its efficient management. Nevertheless, due to the cost of a dedicated power meter for each server, most data center power management only focuses on UPS or cluster-level power monitoring. In this paper, we propose a low-cost novel power monitoring approach that uses only one sensor to extract power consumption information of all servers. We utilize the conducted electromagnetic interference (EMI) of server power supplies to measure their power consumption from non-intrusive single-point voltage measurements. We present a theoretical characterization of conducted EMI generation in server power supply and its propagation through the data center power network. Using a set of ten commercial-grade servers (six Dell PowerEdge and four Lenovo ThinkSystem), we demonstrate that our approach can estimate each server's power consumption with less than ~7% mean absolute error.

CCS CONCEPTS

• **Computer systems organization** → **Sensors and actuators**; • **General and reference** → *Measurement*.

KEYWORDS

Power Monitoring, Data Center, Electromagnetic Interference

ACM Reference Format:
Pranjol Sen Gupta, Zahidur Talukder, Tasnim Azad Abir, Phuc Nguyen, and Mohammad A. Islam . 2023. Enabling Low-Cost Server Level Power Monitoring in Data Centers Using Conducted EMI. In *The 21st ACM Conference on Embedded Networked Sensor Systems (SenSys '23), November 12–17, 2023, Istanbul, Turkiye.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3625687.3625801

*A poster abstract of this paper appeared at the ACM SenSys '22 [1].

SenSys '23, November 12–17, 2023, Istanbul, Turkiye

ACM ISBN 979-8-4007-0414-7/23/11.
https://doi.org/10.1145/3625687.3625801

1 INTRODUCTION

Importance of data center power monitoring. The increasing dependence on the Internet and cloud computing has made data centers an indispensable part of our lives. Data centers, however, are huge energy hogs with massive electricity consumption [2–4]. Unsurprisingly, the efficient power management has been a focal point in data center research [5–9].

A core component of data center power management is the power monitoring system, which can provide valuable insights for exploring the optimization potential of data center operation. Toward this end, a fine-grained server-level power monitoring system can play a crucial role. By providing real-time power consumption information of servers, it can facilitate advanced server power management techniques, such as power capping and idle power saving [10–13]. Server-level power monitoring can even assist in efficient cooling control by load balancing/distribution and helping to identify possible server hot-spots [14, 15]. Moreover, fine-grained power monitoring can play a vital role in safeguarding data centers from outages/downtime due to overloading as well as identifying adversarial server behavior caused by an attacker [16–18].

Limitations of existing systems. Despite its potential, server-level power monitoring is not widely adopted in practice except for flagship IT companies such as Meta [19]. Typical data center monitoring platforms mainly focus on aggregate power monitoring at the UPS and cluster-level [20–22]. A major obstacle here is the cost barrier for fine-granular power monitoring. A server equipped with onboard power monitoring systems can increase its cost by several hundred dollars. For instance, Dell PowerEdge servers can be equipped with a proprietary monitoring system iDrac at additional equipment and licensing cost of ~$400 [23], while IBM PowerExecutive solution comes with dedicated power sensors for each server at the cost of $2000~$2800 [24, 25]. Intelligent rack PDUs with metered outlets can offer cheaper (e.g., ~$80 per server) server-level power monitoring but still cost thousands of dollars each [24, 26–28]. Moreover, adopting server-level power monitoring in existing data centers requires upgrading the rack PDUs, resulting in service disruption during server restart to switch over to the new PDU. Not to mention, such server-level power monitoring equipment requires a complex and costly network of protocol conversion and middleware to integrate with any centralized monitoring system. Alternative to hardware-based metering, software techniques have also been proposed that use system counters to

estimate each server's power consumption [29–31]. However, these require intrusive access and consume computing resources to run monitoring tools on the server.

But, the most significant limitation of existing power monitoring systems, both hardware and software-based, is their inability to monitor the server-level power consumption of a tenant's server rack in colocation data centers. Colocations are an important class of data centers that take up nearly 40% of the data center market, with a compound annual growth rate of 14.2% from 2023 to 2030 [2, 32]. Colocation data centers provide a shared data center infrastructure solution where the tenants retain full control and complete visibility of their IT resources by bringing their physical servers inside the data center while the data center operator provides power, cooling, and space. In colocations, the operator does not have any access to tenants' server racks [33–35], and therefore, is unable to deploy any existing server-level power monitoring.

Our contribution. In this paper, we propose the first-of-its-kind conducted electromagnetic interference (EMI) based server-level power monitoring in data centers. More specifically, we measure the voltage at a higher level (e.g., cluster power distribution unit (PDU)) of the data center power distribution hierarchy, and from this "single-point measurement" disaggregate the power consumption of all the servers connected to that PDU.

Fundamental working principle. Our power metering approach is built upon the following observations. We identify that the power factor correction (PFC) circuit is ubiquitously present in all server power supplies due to government-mandated EMI regulations [36–38], and these PFC circuits cause conducted EMIs at high-frequency range (e.g., 10kHz ∼ 150kHz) [39–41]. More specifically, to mitigate harmonics due to non-sinusoidal current, the PFC circuit in the power supply uses high-frequency switching to shape the current into a sinusoidal form (discussed in Section 2.2 and illustrated in Fig. 3). The high-frequency switching creates high-frequency current ripples that travel through the data center power distribution network as conducted EMIs and can be extracted from the voltage measurement at the equipment that the current flows through. More importantly, these conducted EMIs change with the server power consumption and, hence, can be used to estimate the server power.

The high-frequency conducted EMIs from different servers are independent of each other in the frequency domain and "unlikely" to overlap completely, which is observed in our experiments (Table 2) as well as in prior works [40, 42]. This is because the PFC switching frequency, and hence the frequency of the conducted EMI, is determined by the PFC circuit's inductance and capacitance. The minute differences due to process variation during the production of indicators and capacitors result in variation (in the range of a few tens of Hz) in the switching frequency, even among the power supplies from the same batch. On the other hand, the conducted EMIs from power supplies occupy a narrow band due to the prevalent choice of using continuous conduction mode (CCM) PFC circuits, which operate at a constant switching frequency [39, 43]. Two advantages of CCM dictate this choice for computer power supplies. First, for the same power level, CCM-type PFC has a lower peak current and, hence, can be designed with components with lower current ratings. Second, the constant frequency operation simplifies the choice of frequency-sensitive inductor and capacitor components [43]. The combination of non-identical and narrow-band server EMI manifests, in the frequency domain, as *approximate orthogonality* where each server's EMI remains distinguishable amidst many other servers' EMIs. This is a key enabler for our successful extraction of server-wise power from a single voltage sensor.

Merits of our approach. We overcome the aforementioned limitations in existing systems. ① We significantly reduce the cost of server-level power monitoring by eliminating the need for dedicated power sensors for every server. ② Our voltage measurement by nature is non-intrusive, making our system easily retrofittable on existing data centers without any service interruption (i.e., without any hardware/physical modification of existing power infrastructure). ③ We enable visibility into tenants' servers in colocation data centers by enabling extraction of server-level power consumption from cluster PDUs. ④ Moreover, due to low cost and retrofittable design, our approach can even be used alongside existing server-level power monitoring systems (e.g., in Facebook data centers) as a "second pair of eyes" to detect sensor malfunction and false data injection attacks [44, 45].

Advantage of a conducted EMI-based approach. Our proposed idea can be seen as an exploitation of a voltage side-channel and therefore begs the question, "Why not use other data center side-channels such as thermal or acoustic?". Thermal side channels are less stable as they are severely affected by any change in air flow, while acoustic side channels are susceptible to complex mixing due to reverberation and multi-path propagation [17, 46–50]. Moreover, the power signatures from different servers, for both thermal and acoustic side channels, completely overlap, making it extremely difficult to separate using a single sensor. In contrast, our proposed voltage side channel has a more predictable change with the data center environment (e.g., change in line resistance due to temperature change), and the conducted EMIs exhibit orthogonality that allows server-level signal extraction. Radiated EMI-based side-channels, on the other hand, suffer from a short range and have been utilized in tracking CPU instruction execution and memory access in an isolated machine [51, 52].

Experiments and evaluation. We use a set of commercial servers (six Dell PowerEdge and four Lenovo ThinkSystems) to test and evaluate our EMI-based monitoring approach. Our evaluation shows that we can estimate the server-level power consumption of our cluster of ten servers with less than 7% mean absolute error in our five-day-long experiment with real-world workload traces. To encourage further development, we also plan to publicly share our experimental data after the publication of our work. Experimental data and source code of our evaluation are available at Open Science Framework (OSF) [53].

Limitations. We would like to offer our insight into the limitations of our proposed EMI-based power monitoring on two different fronts - the fundamental limitations associated with the conducted EMI-based approach and the limitations in our current implementation of the EMI-based power monitoring.

Fundamental limitations. The conducted EMI from server power supplies is generated due to the AC-to-DC conversion using a PFC. Hence, our approach will work only in data centers with AC distribution systems and is not suitable for DC-powered data

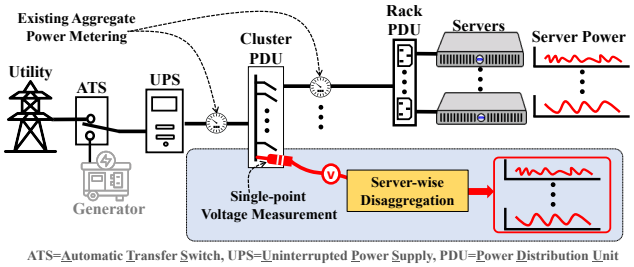

Figure 1: Overview of data center power infrastructure and our EMI-based server-level power monitoring approach.

centers. In addition, as we extract the power consumption of server power supplies, our approach cannot distinguish blade servers that are powered by the backplane of the chassis with a single unified power supply [54]. Nonetheless, neither DC distribution nor blade servers are the norms in today's data centers [55, 56], and hence our approach is still widely applicable. In our approach, the voltage measurements need to be taken within the same power network without any isolation due to transformers. Similar isolation is also present among phases of three-phase power distribution in larger data centers. Hence, we need a dedicated voltage sensor for each power network, while the data processing and extraction of server power from the voltage data can be done centrally.

Limitations of our implementation. We defer the discussion on the limitations of our implementation and our plans to overcome them towards the end of the paper in Section 6 to make it more meaningful to the reader.

2 PRELIMINARIES

2.1 Data Center

Power infrastructure. Data centers typically follow a hierarchical tree-type power infrastructure [40, 57]. As illustrated in Fig. 1, grid/utility power enters the data center through an automatic transfer switch (ATS), which switches over the main power supply to the backup generator in case of grid power failures. The automatic transfer switch feeds power to an uninterrupted power supply (UPS), which supplies "protected" power to multiple cluster-level power distribution units (PDUs) at 480/277 V. The PDUs reduce the voltage to 208/120 V and supply power to racks that house the computer servers. The server racks also have power-strip-like distribution units (also called rack-PDU) where the servers' power cords are directly connected. Transformers in the cluster PDUs act as low-pass filters, and therefore, in our context, they create isolated power networks downstream. Large data centers use three-phase power delivery, and the phases are also electrically isolated from each other. Since our power meter works within the same power network, every phase (if applicable) of every cluster PDU will require a separate voltage sensor. The architecture also varies based on the type of redundancy it provides. For example, a 2N redundant data center have two identical power infrastructure running in parallel. While the power infrastructure we have discussed here is applicable to typical commercial data centers, and the voltages at different levels are for U.S.-based equipment, specialized data centers with diverse power infrastructures exist that serve specific application needs. Nonetheless, our proposed server-level monitoring approach applies to most (if not all) data center AC power distribution architectures.

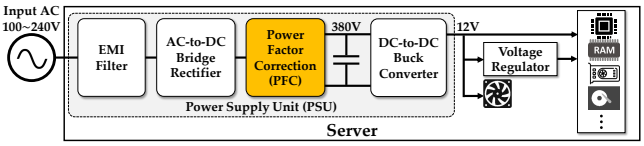

Figure 2: Server components showing the different stages of the power supply unit that delivers regulated DC voltage to the internal components. Our item of interest, the PFC stage, is highlighted.

Power metering. The UPS and cluster PDUs have circuit branches protected by circuit breakers (CB) from catastrophic impacts (e.g., fire) due to power faults such as short circuits. The CBs also protect from power overloads (e.g., due to short circuits) in one branch from propagating to others. As shown in Fig. 1, to monitor the data center power infrastructure, the branch circuits of the UPS and cluster PDUs have power meters and typically can offer server rack-level power monitoring. These power meters are also used in colocation facilities to ensure tenants' peak power compliance and energy usage billing.

2.2 Computer Power Supply

The power supply provides a regulated DC voltage to the server motherboard and other internal components (e.g., hard disks). Fig. 2 shows the basic building blocks of a server power supply [39]. In compliance with international regulation, the 100~240V input AC power first passes through an EMI filter that eliminates frequency components larger than 150kHz from coming in as well as conducted back to the data center power network [41, 58–60]. The sinusoidal AC power is then passed through an AC-to-DC bridge rectifier that converts it to a pulsating DC voltage (unipolar half-sine waves) followed by a power factor correction (PFC) stage. The PFC outputs an elevated DC voltage of 380V, which is then brought down by a DC-to-DC buck converter to 12V for the internal components.

Power Factor Correction The PFC circuit improves the power factor by avoiding the generation of harmonics due to distorted (non-sinusoidal) current wave shape. The ideal power factor for a power distribution system is "1", where the voltage and current are perfect sine waves. The PFC circuit samples the input voltage and uses it as a sinusoidal reference to shape the current consumption.As illustrated in Fig. 3, the most commonly used boost-type PFC consists of an inductor, a diode, a switch (MOSFET), and a control circuit for the pulse width modulation (PWM) [39]. The control circuit rapidly toggles the MOSFET "ON" and "OFF" at a high frequency (i.e., the PFC switching frequency). The input current goes up when the MOSFET is "ON" and goes down when the MOSFET is "OFF", resulting in ripples in the current. The control circuit shapes the input current to follow the sinusoidal reference wave shape by controlling the MOSFET's "ON" and "OFF" time durations. The right figure in Fig. 3 illustrates the sinusoidal reference and the resulting current waveform with the high-frequency ripples. Note here that, for clarity, the illustration is shown at a much lower PFC

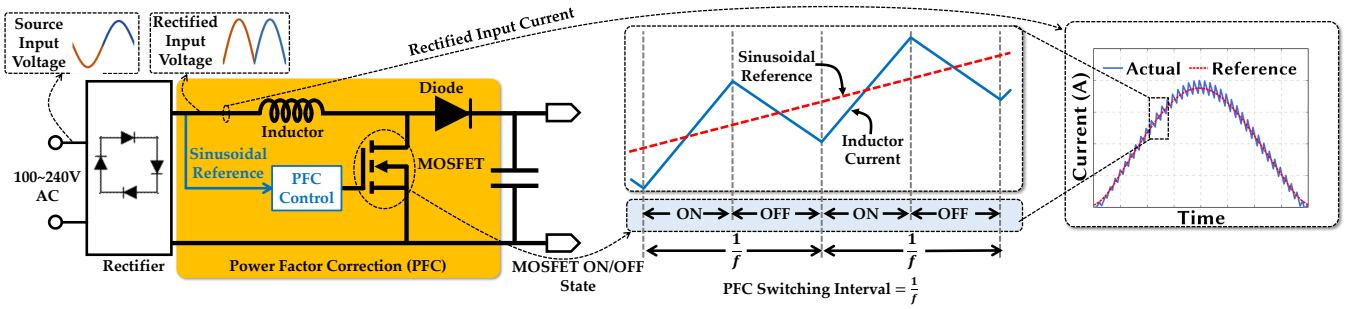

Figure 3: Illustration of current ripples generated by the PFC circuit in a computer power supply.

switching frequency of around 3kHz, whereas in commercial server PFC circuits, the switching frequency ranges between 50kHz and 150kHz.

2.3 Related Work

Server-level power monitoring. Server power monitoring using dedicated power meters has been long existing and with numerous commercially available products. There are smart rack PDUs that offer outlet-level power monitoring [24, 26–28]. Server vendors also offer add-on modules to monitor their server power consumption [23–25]. Meanwhile, [61] develops a low-profile plug-through prototype that can monitor AC outlet energy consumption. Dedicated hardware provides the most accurate power metering but requires significant investment to reach server-level metering in data centers.

Alternatively, software-based solutions also exist that build power models for server power estimation using information collected from other sources. For instance, hardware performance monitoring counters have been exploited in [29, 62, 63], while [30, 31, 64] profile the power consumption of server components using various utilization matrices. While cost-effective, software-based approaches require access to the servers' performance counters, which does not apply to multi-tenant colocation data centers. Further, software-based power monitoring requires intrusive offline training, which makes it difficult to implement in an online fashion without interrupting the service.

Utilizing voltage measurement. The most closely related to this work is [40], where a voltage side channel has been exploited to estimate the data center's total power. In sharp contrast, in this project, we go the other way and try to estimate server-level power metering. While [65, 65–67] also similarly take the voltage measurement as this project, they focus on the residential environment with fewer signal sources and tries to extract household appliance and equipment activity. We, on the other hand, focus on a commercial data center environment with many servers. More importantly, we go much further beyond identifying a server's ON/OFF status and measuring its power consumption.

3 CHARACTERIZATION OF SERVER EMI

In this section, we first offer a theoretical characterization of conducted EMI generation and propagation. We then discuss the relationship between conducted EMI and server power consumption, followed by experimental validation and design considerations for EMI sensing for power monitoring.

3.1 Generation of Conducted EMI

Conducted EMI results from server current. Since power is supplied at a constant voltage [1], the current changes with server power consumption. The PFC in the power supply shapes the server current to match the sinusoidal supply voltage at the grid's frequency (e.g., 60Hz in the US). However, as illustrated in Fig. 3, the PFC also adds high-frequency saw-tooth ripples to the current due to its rapid switching. Moreover, these high-frequency ripples are modulated by the sinusoidal reference signal. This is because the shapes of the PFC ripples change based on the reference signal's relative location within its sinusoidal shape. The saw-tooth waves have a longer rise time (MOSFET is ON) when tracking the rising half of the sinusoidal reference and vice-versa. Hence, we characterize the server current as follows

$$i(t) = \overbrace{a_1(t)\cos(2\pi f_0 t)}^{\text{Power delivery}} + \overbrace{a_2(t)\cos(2\pi f_0 t)\cos(2\pi f_p t)}^{\text{EMI noise}} \quad (1)$$

where f_0 and f_p are the frequencies of the power grid and PFC switching, respectively, and a_1 and a_2 are the amplitude of the sinusoidal signals. Here, we approximate the high-frequency saw-tooth current ripples created by PFC using a sinusoidal signal oscillating at the PFC switching frequency. This is a reasonable approximation since the fundamental frequency (i.e., frequency of the saw-tooth signal) is the dominant frequency component of a saw-tooth signal [68]. Note that the first sinusoidal term is responsible for the power delivery, and the second sinusoidal wave is "unwanted" noise/EMI. Consequently, server power supplies are designed to have $a_2 << a_1$. Unless otherwise specified, by server power consumption, we refer to its root mean square (RMS) power, not the instantaneous power, which is a sine-squared function. Now, using simple trigonometry, we can rewrite Eq.(1) as

$$\begin{aligned} i(t) &= a_1(t)\cos(2\pi f_0 t) \\ &+ \frac{1}{2}a_2(t)\left[\cos(2\pi(f_p + f_0)t) + \cos(2\pi(f_p - f_0)t)\right] \end{aligned} \quad (2)$$

Eq. 2 indicates that in frequency analysis, in addition to the fundamental component at the grid frequency (e.g., 60Hz), we should also see two other high-frequency components at frequencies $f_p + f_0$ and $f_p - f_0$. Moreover, the high-frequency components are precisely

[1] In normal operation, the data center gets power directly from the power grid, and therefore, the data center voltage perturbs in the range of a few volts [40]. However, this voltage perturbation is negligible compared to the change in current due to power change.

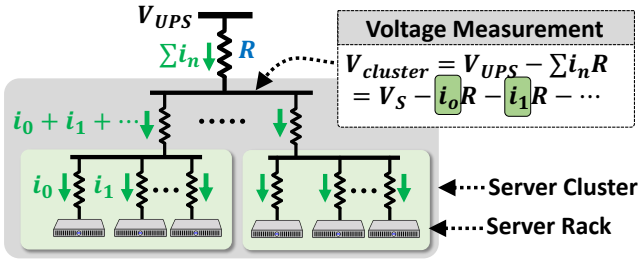

Figure 4: Server current flowing through the data center power network. Server currents, $i_0, i_1, \cdots$, are embedded in the voltage measurement at the cluster $V_{cluster}$.

$2f_0$ away from each other. We rewrite the server current in Eq. (2) by identifying their frequency components as follows

$$i(t) = i^{f_0}(t) + i^{f_p \pm f_0}(t) \tag{3}$$

where $i^{f_0}(t) = A_1(t)\cos(2\pi f_0 t)$ is fundamental power component and $i^{f_p \pm f_0}(t) = \frac{1}{2}a_2(t)\ [\cos(2\pi(f_p + f_0)t) + \cos(2\pi(f_p - f_0)t)]$ is the conducted EMI component.

3.2 Propagation of Conducted EMI

As shown in Fig. 4, the server current, and hence the conducted EMI, flows through the data center's power distribution hierarchy. The current consumption of every server flows from the UPS through the cluster PDU and rack PDU. Following Ohm's Law and dropping the time index t for clarity, the voltage measurement at the cluster level can be expressed as $V_{cluster} = V_{UPS} - R\sum_n i_n = V_{UPS} - Ri_0 - Ri_1 - \cdots$, where R is the line resistance between the cluster PDU and the UPS, and i_n is the current of the n-th server. Next, by introducing the frequency components from Eq. 3, we can express the cluster voltage as

$$\begin{aligned} V_{cluster} &= V_{UPS}^{f_0} - R\sum_n i_n \\ &= V_{UPS}^{f_0} - R\sum_n \left(i_n^{f_0} + i_n^{f_{p,n} \pm f_0}\right) \\ &= \overbrace{V_{UPS}^{f_0} - R\sum_n i_n^{f_0}}^{Grid frequency} - \overbrace{R\sum_n i_n^{f_{p,n} \pm f_0}}^{PFC frequencies} \end{aligned} \tag{4}$$

Eq. 4 offers the foundational understanding of our server power monitoring approach by showing that *each server's current/power is embedded in the cluster voltage*. We see that since the PFC frequencies of different servers are not identical, we can *approximately* separate each server's high-frequency components, $R \cdot i_n^{f_{p,n} \pm f_0}$, from the voltage reading in the frequency domain. On the other hand, Eq. 4 also reveals that extracting an individual server's current from the grid/low-frequency component f_0 is very difficult (if possible at all) because all servers' impacts overlap at f_0.

3.3 PFC EMI vs. Server Power

For the PFC, the input voltage provides the sinusoidal reference to shape the current into a sinusoid (Fig. 3). However, while the voltage remains constant (e.g., around 120V in the U.S.), the current must change (while maintaining the sinusoidal shape) in proportion to the server power. For instance, a server running at 100W will result in twice as tall PFC current sinusoid compared to the server running at 50W. Consequently, the PFC switching needs to be changed to accommodate the change required in the PFC current due to a change in the server power consumption. This change in PFC switching depends on the proprietary designs of the PFC control of the server power supply manufacturer, and hence, we do not offer any mathematical characterization here. Nevertheless, based on our experiments with commercial servers from Dell and Lenovo, we identify two fundamental ways the PFC control reacts to server power (and current) change - by changing the switching frequency (observed in our Dell servers) and by changing the amplitude of the high-frequency ripples (observed in our Lenovo servers). We denote the frequency change with power as "Frequency-Modulated PFC Control" or **FM-PFC** and amplitude change with power as "Amplitude-Modulated PFC Control" or **AM-PFC**. We offer further insight into the FM-PFC and AM-PFC in the next section, where we discuss the relationship between the EMI and server power.

3.4 Experimental Validation

EMI generation and propagation. To validate our characterization of server power supply EMI, we take one of our Dell PowerEdge servers and power it through our rack PDU. The rack PDU is connected through a 1kW Tripp-Lite isolator [69] to offer a clean EMI environment. We then run the server at high power using a CPU stressing program. We then use our Tektronix MSO54-5 series oscilloscope to monitor the PDU voltage from one of the PDU power outlets. Fig. 5 shows voltage reading and its FFT on the left and their corresponding zoomed-in versions on the right. In the FFT of the voltage, we can clearly see the frequency spikes created by the PFC. The zoomed-in version also shows the spikes' location at ~69.5kHz and a separation of 120Hz, confirming our characterization in Eqn.(3). Meanwhile, the zoomed-in figure of voltage shows the ripples from the PFC. The presence of PFC ripples in the voltage reading confirms our analysis of conducted EMI propagation in Eqn. (4) through the power network. We observe similar signals for the Lenovo server as well. In our evaluation in Section 5.2, we further study the EMI propagation and impact from voltage sensing location and different phases of the power network.

Conducted EMI vs. server power consumption. To utilize EMI for power monitoring, we need to understand how the conducted EMI changes with changes in server power consumption. For this, we separately run our Dell and Lenovo servers at different power levels and collect the voltage data at a sampling rate of 200K samples per second. We then apply Fast Fourier Transform (FFT) in one-second windows on the voltage data. From the FFT, we identify the PFC EMI frequency f_p and aggregated Power Spectral Density (PSD) of PFC EMI within 30Hz windows around the two EMI spikes at $f_p \pm f_0$. We use 30Hz windows for PSD as it offers robust EMI amplitude detection due to perturbation in the PFC frequency within our one-second FFT windows. Also, we use the sum of the square of the frequency amplitudes to approximate the PSD. From our results, we find that the Dell server's power supply has FM-PFC, and the Lenovo server's power supply has AM-PFC.

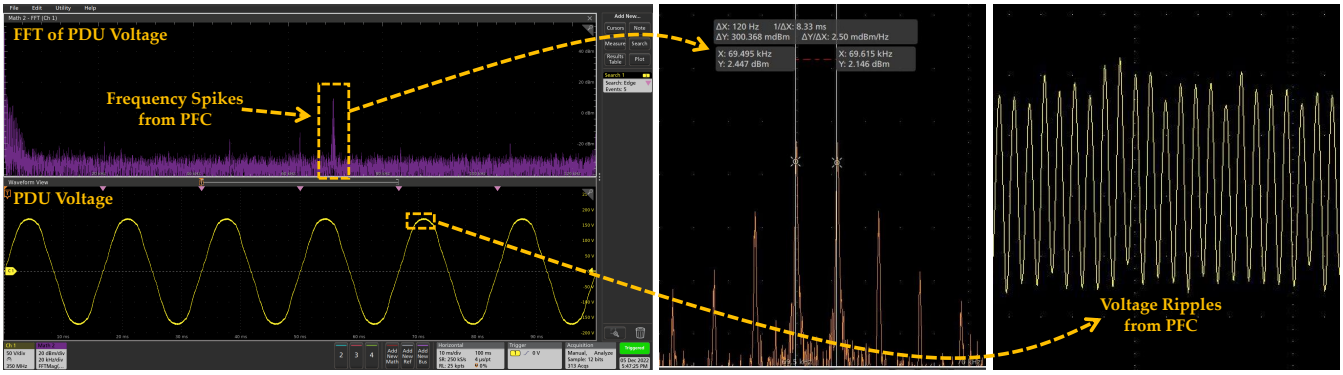

Figure 5: Experimental validation using a Dell PowerEdge server and Tektronix MSO54-5 series oscilloscope. *Left figure:* Voltage reading (bottom) of PDU and the corresponding FFT (top). The FFT shows the high-frequency components generated by the PFC. *Center figure:* Zoomed-in version of FFT at the PFC frequency shows that the spikes are at around 69.5kHz and 120Hz (i.e., $\Delta X : 120Hz$ in the figure) apart from each other. *Right figure:* Zoomed-in voltage shows the ripples created by the PFC.

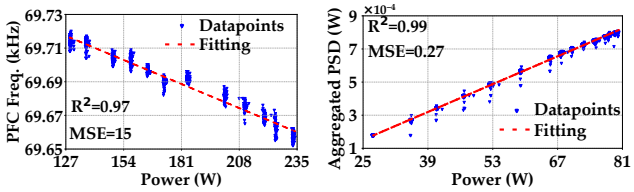

(a) Dell (b) Lenovo

Figure 6: (a) The power vs EMI frequency relation of a Dell server. (b) The power vs. PSD relation of a Lenovo server.

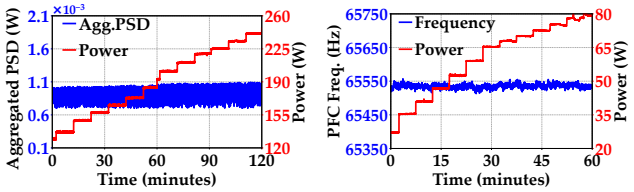

(a) Dell (b) Lenovo

Figure 7: (a) Change in aggregate PSD with power change for Dell with FM-PFC. (b) Change in PFC frequency with power change for Lenovo with AM-PFC.

FM-PFC. For FM-PFC in our Dell servers, the PFC switching frequency changes with server power consumption. We see in Fig. 6(a) that the PFC frequency goes down linearly with increased power consumption. We model server n's power and PFC switching relationship as

$$p_n(t) = A_n \cdot f_{p,n}(t) + B_n, \tag{5}$$

where A_n and B_n are model parameters. AM-PFC. For AM-PFC in our Lenovo servers, the EMI PSD changes with server power consumption. Fig. 6(b) shows that the EMI PSD goes up with increasing server power consumption. We model server n's power and EMI PSD relationship as

$$p_n(t) = C_n \cdot \text{PSD}_{f_p,n}(t) + D_n, \tag{6}$$

where C_n and D_n are model parameters.

Due to differences in the PFC switching frequency for FM-PFC and the relative distance between voltage sensing and EMI generation, the model parameters A_n, B_n, C_n, and D_n are different for different servers. Moreover, we cannot run server-wise experiments to extract these model parameters for every server we want to monitor in a data center. We discuss how we extract these model parameters without server-wise experiments in Section 4.2. We also show in Fig. 7 that the PSD for FM-PFC and the PFC frequency for AM-PFC do not change with power.

To summarize, our experiments corroborate the foundation of our approach that *servers generate high-frequency conducted EMIs due to their PFC circuit; the PFC EMI can be detected from the voltage reading at the PDU; and we can determine a server's power consumption from its PFC EMI.*

3.5 Design Considerations in EMI Sensing

Based on our experiment results, here we offer design consideration toward our EMI sensing.

Voltage acquisition from power lines. Most server power supplies use single-phase AC (e.g., 120V in the U.S.), which, as shown in Fig. 8(a), has three wires - Line (L), Neutral (N), and Ground (G) with three different configurations for voltage data collection. We find that each of these configurations carries the PFC EMI signal. Also, the voltage ranges of these configurations vary as Line-Neutral and Line-Ground are in the 120V range while the Neutral-Ground remains close to "zero." To evaluate which configuration can best capture the PFC EMI, we run experiments for 12 hours on our Dell and Lenovo servers. Fig. 8(b) shows the absolute error in the estimation of server power for the two servers, while Table 1 lists their signal-to-noise ratio (SNR). The absolute error is calculated as the difference between true power and EMI-estimated power. The box plot in Fig. 8(b) extends from the first quartile (Q1) to the third quartile (Q3) of the data, with a black line at the median and the mean represented using a diamond marker. The whiskers extend from the box by 1.5x the interquartile range (IQR). We see that L-N and N-G configurations produce better results in terms of accuracy. While the N-G configuration performs better in Lenovo power detection, the L-N configuration works better with

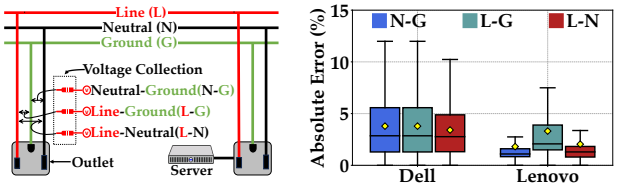

(a) Overview of Location (b) Data Collection Location vs Error

Figure 8: Effect of voltage data collection location.

Table 1: Statistics of voltage data acquisition from the power line with different configurations.

Voltage	Range	Dell		Lenovo	
		SNR	Accuracy	SNR	Accuracy
Neutral-Ground	0V-10V	11 db	High	10 db	Highest
Line-Ground	100V-140V	11 db	High	5 db	Medium
Line-Neutral	100V-140V	11.32 db	Highest	9 db	High

Dell. Out of these two configurations, we choose the L-N setup. However, as the L-N operates at high voltage and is not suitable for typical voltage sensors, we use a voltage divider to reduce the voltage to one-tenth and a second-order RC high-pass filter with a cutoff frequency of 3kHz. The application of the high-pass filter also offers better sampling resolution for EMI sensing.

Analog-to-digital converter specification. Our voltage sampling essentially is an Analog-to-Digital Converter (ADC). We run experiments with three different pieces of equipment (shown in Fig.9(a)) of various price ranges - a Tektronix MSO54-5 series oscilloscope with up to 16-bit ADC [70], a National Instruments DAQ USB-6361 with 14-bit ADC [71], and Digilent USB oscilloscope Analog Discovery-2 with 14-bit ADC [72]. We run our Dell and Lenovo servers at different power levels for 12 hours. Figs.9(b) and 9(c) show the absolute error in EMI-based power estimation for the three ADCs at different bits per second configurations. We see that the Tektronix oscilloscope performs better for the Dell server with FM-PFC, while Digilent performs better for the Lenovo server with AM-PFC. We see that even at 8-bit ADC resolution, the EMI-based power estimation performs similarly to 14-bit ADC resolution. This indicates that we can utilize lightweight ADCs to implement our EMI-based power monitoring.

Sampling rate. The PFC switching frequencies of Dell and Lenovo servers are at around 70kHz and 65kHz, respectively. Consequently, we need a sampling rate of greater than 140K samples per second to capture their EMIs. Nevertheless, to see if a higher sampling rate benefits our design, we run experiments with different levels of sampling rates from 200K samples per second to 1600K samples per second. Figs. 10(a) and 10(b) shows the absolute error in power estimation. We see that 200K samples per second are sufficient to extract useful information for both our FM-PFC and AM-PFC servers.

4 EMI-BASED POWER MONITORING

Based on our conducted EMI characterization and experiments, in this section, we develop our EMI-based server power monitoring algorithm. We first discuss how the server EMIs coexist in a multi-server environment, followed by our algorithm to extract

(a) Tektronix MSO54, NI-Daq, Digilent Analog Discovery-2

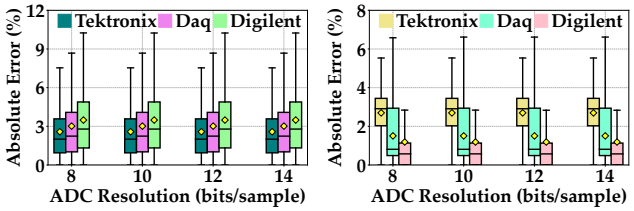

(b) Effect of Bit Rate on Dell (c) Effect of Bit Rate on Lenovo

Figure 9: Effect of ADC resolution.

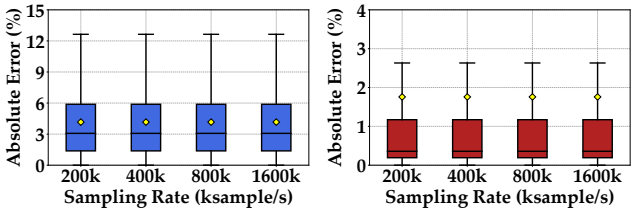

(a) Effect of Sampling Rate on Dell (b) Effect of Sampling Rate on Lenovo

Figure 10: Effect of ADC sampling rate.

server power consumption from conducted EMIs in the voltage measurement of the PDU.

4.1 EMIs in Multi-Server Environment

To understand a multi-server environment, we need to look at the EMI frequencies of different servers as well as how these EMIs mix in the voltage measurement. For this, we examine a set of six Dell and four Lenovo servers (details of the experimental setup are discussed in Section 5.1).

EMI frequencies. We run each server individually to find their corresponding conducted EMI frequencies. As listed in Table 2, we see that the servers' EMIs are generated at different frequencies. The differences in EMI frequency among the servers from Dell and Lenovo are already anticipated due to their design difference. However, to investigate the EMI frequency variation among the servers from the same manufacturer, we take out the server power supplies and examine their model and serial numbers (Fig. 11). We find that our servers from the same manufacturer acquired in a single purchase order use the same power supply model. Moreover, judging by the significant overlaps in their unique serial numbers, we presume these power supplies are also manufactured in the same production batch and, therefore, operate using the same control algorithm. This observation supports our claim that the frequency of the server EMI can vary even among the power supplies of the same model. A similar observation is also made in [40].

Table 2: Server PFC frequencies.

Server	Dell#1	Dell#2	Dell#3	Dell#4	Dell#5	Dell#6	Lenovo#1	Lenovo#2	Lenovo#3	Lenovo#4
EMI Frequency (Hz)	69648	69957	69934	70113	69756	69548	65314	65525	66335	67549

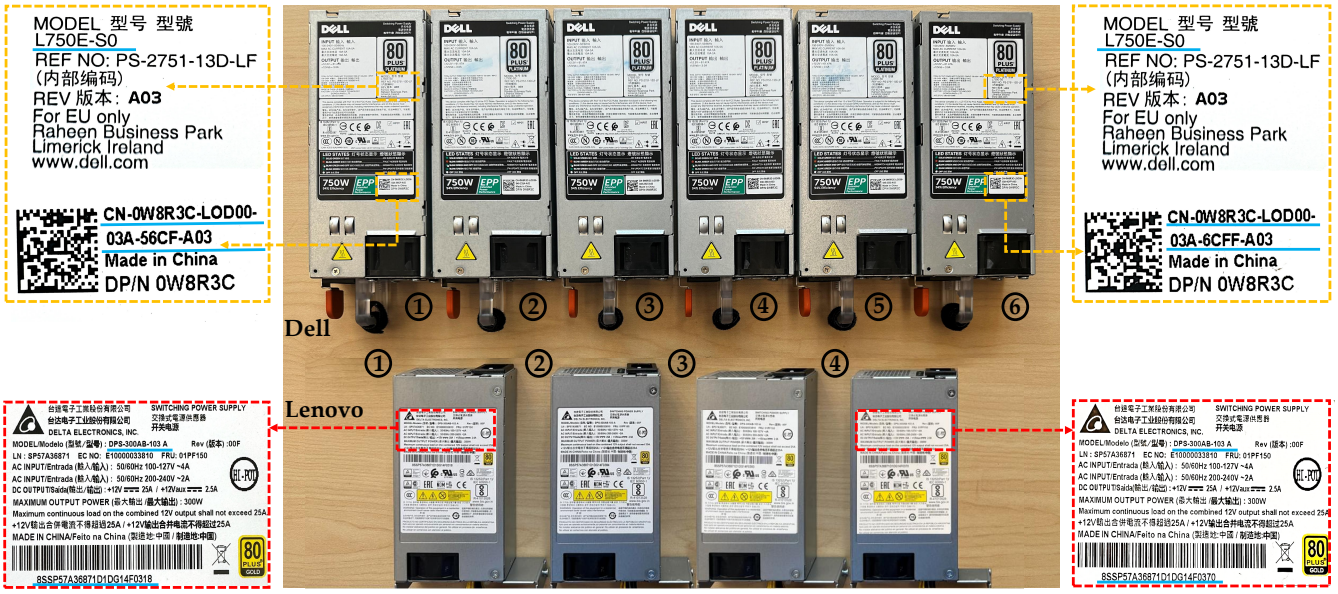

Figure 11: Power supply units of our servers. Both Dell and Lenovo are using the same power supply models in their servers of the same models, but their EMI frequencies are not identical.

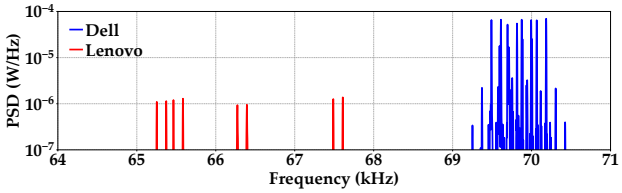

Figure 12: Frequency domain analysis of the servers.

EMI mixing. In data centers, the same power distribution system powers many servers. Consequently, the conducted EMIs from all the servers connected to the same power network appear in the voltage measurement. To investigate how conducted EMIs from multiple servers mix together, we simultaneously power all of our servers using a single PDU. We then take voltage measurement from the PDU and show its frequency analysis in Fig. 12. We identify that the *conducted EMIs are additive in the voltage measurement.* More importantly, we can also distinguish different servers' EMIs, albeit there are cases of overlaps, especially for the Dell server EMIs. Our observations on server EMI mixing match with prior works on conducted EMIs from computer power supplies [40, 42].

4.2 Algorithm for Power Monitoring

Our experiments with a Dell and a Lenovo server (Fig. 6) demonstrate that conducted EMIs reveal a server power consumption. Hence, the EMI-based power monitoring in a multi-server environment boils down to separating the EMI signatures of every server from the voltage data. Here, we rely on frequency-domain orthogonality to separate the server EMIs, and therefore, our algorithm operates primarily in the frequency domain. There are three major steps in our power monitoring algorithm. *Step 1:* We need to first detect the EMI spikes of every server in the frequency domain to determine each server's *instantaneous* PFC switching frequency $f_{p,n}(t)$. *Step 2:* For server-wise power monitoring, we need to track a server's PFC frequency over time; that is, we need to tag every PFC frequency detected in Step 1 with a server. *Step 3:* In the final step, we utilize EMI-to-power models to determine each server's power consumption.

Our goal is to achieve a power monitoring temporal resolution of one sample per second. Hence, we divide our streaming voltage data into one-second chunks and apply our algorithm. In what follows, we discuss the challenges involved in each step and how we overcome these.

***Step 1:* Finding the frequency spikes of servers' EMIs.** We first narrow down our search for frequency spikes within the few kHz band that the server EMI occupies, which in our experiments is between 64kHz and 71kHz. This band is easy to identify in the data center power network, which consists of only servers. [2] Within the target frequency band, we can find frequency spikes using standard peak-finding algorithms (e.g., Matlab's `findpeaks` [73]). For complete separation of N servers, we need to find $2N$ frequency spikes as each server creates a pair of EMI spikes around its PFC switching frequency at $f_{p,n} \pm f_0$. We pair the spikes from the same

[2] In data centers, non-IT equipment such as lights are powered by a separate power network.

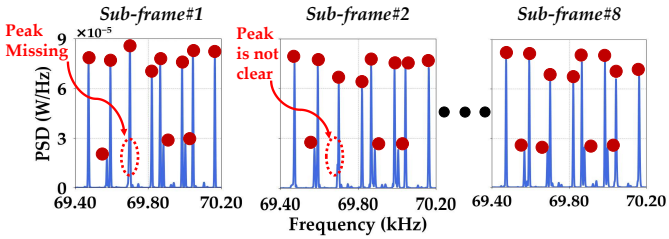

Figure 13: Peak detection using sub-framing where *Sub-frame#8* is chosen as the representative frame.

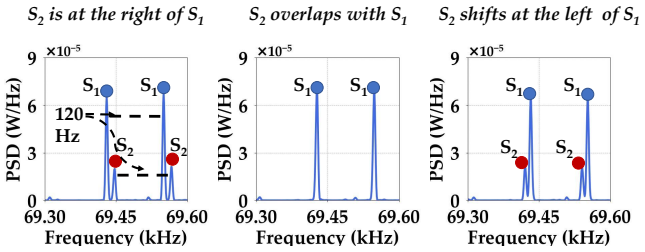

Figure 14: PFC frequency tracking across time frames for Dell servers with FM-PFC.

server together by checking if their frequency distance is exactly $2f_0$ (e.g., ~120Hz in the U.S.). Note that the additional frequency spikes at multiples of grid fundamental f_0 are negligible as they are much shorter.

As shown in Fig. 12, we have many frequency spikes close to each other, especially for the Dell servers, which also move with their power consumption. This can result in overlaps of the EMI spikes from different servers, resulting in fewer than $2N$ frequency spikes detected. Moreover, within our one-second sampling frame, the frequency of the PFC may perturb and spread among multiple frequency bins in FFT. This makes pinpointing the frequency of the EMI harder.

We can reduce the impact of the EMI frequency perturbation in FFT by reducing the FFT frame size at the expense of a decrease in the SNR. Hence, we divide our one-second frames into eight sub-frames of size 0.25 seconds with 0.125 seconds overlapping. We then apply FFT and peak detection on each sub-frame. Our choice of sub-frame size is motivated by our experiments where the SNR with 0.25 seconds FFT frames is still sufficient for EMI detection. Fig. 13 demonstrates how applying sub-framing allows us to find the frequency spike of a server buried under another. We keep one of the sub-frame as the representative frame for that particular one-second frame. The sub-framing, however, cannot completely avoid cases of missing spikes. In such rare cases, we match the detected frequency spikes of the current time frame with the previous time frame to determine the absent spikes.

With the frequency spikes of all servers detected, we know each server's PFC frequency, which is the average of the server's frequency spikes pair. Note that our EMI spike detection approach discussed here depends on frequency domain orthogonality and, therefore, may not scale up well. We offer more discussion of the limitations of the approach and possible directions for improvement in Section 6.

***Step 2:* Tracking server's EMI overt time.** For server-wise power monitoring, we need to track a server across time frames. For Lenovo servers with AM-PFC, their PFC frequencies act as the marker/identifier of a particular server, albeit there are some perturbations in PFC switching frequency even for Lenovo servers, as shown in Fig. 7(b). For the Dell servers with FM-PFC, we cannot use only the PFC switching frequency as the marker for servers. For FM-PFC, the PFC switching occupies a small frequency band - the lower limit of the band is associated with the server's peak power, and the upper limit of the band is associated with its idle power (Fig. 5). In our case of the Dell servers, the band is around 80 Hz wide. Now, the FM-PFC servers that have overlaps in this band cannot use the PFC switching frequency alone as their marker. To overcome this, we utilize the past amplitude statistics of the FM-PFC as we find in our experiments that the EMI amplitude of FM-PFCs does not vary with power consumption (Fig. 7(a)). We also impose the additional constraint that a server's power would not change drastically within a second to narrow down our server tracking across time frames. Fig. 14 illustrates how we track a server when it moves from right to left on the frequency axis.

***Step 3:* Extracting model parameters and power consumption.** Our first two steps identify every server's PFC frequency $f_{p,n}$. For an FM-PFC, we directly plugin $f_{p,n}(t)$ into its EMI-to-power model in Eqn. (5), whereas for an AM-PFC, we use $f_{p,n}(t)$ to determine the aggregation band to find the $\text{PSD}_{f_{p,n}}(t)$ to plugin into its EMI-to-power model in Eqn. (6).

To apply the EMI-to-power models, we need to extract the model parameters for every server. For this, we utilize the data center's existing cluster-level power meters (as shown in Fig. 1), which offer the measurement of the total power of all servers within the same cluster. Considering there are N_{FM} servers with FM-PFCs and N_{AM} servers with AM-PFCs, the cluster -level power $p_{cluster}(t)$ can be written as

$$p_{cluster}(t) = \sum_{n=1}^{N_{FM}} (A_n \cdot f_{p,n}(t) + B_n) + \sum_{n=1}^{N_{AM}} (C_n \cdot \text{PSD}_{f_{p,n}}(t) + D_n) \quad (7)$$

We can use Eqn. (7) to extract the model parameters for every server from online (i.e., servers are operating as normal) data samples of the server EMI frequencies, PSD, and cluster level power. For robust extraction of the model parameters, we need a number of samples $\gg 2(N_{FM} + N_{AM})$. We can estimate N_{FM} and N_{AM} by tracking the EMI frequencies over a short time (e.g., 12 hours) - if the EMI frequency randomly varies, then it is from an AM-PFC (Fig. 7(b), and if the EMI frequency follows typical server workload patterns (e.g., has a correlation with the $p_{cluster}(t)$) then it is from FM-PFC. Note here that while we collect the data samples from the live system, the extraction of model parameters can be done offline. However, at no point do we need to shut down any server to profile their EMI vs power relationship.

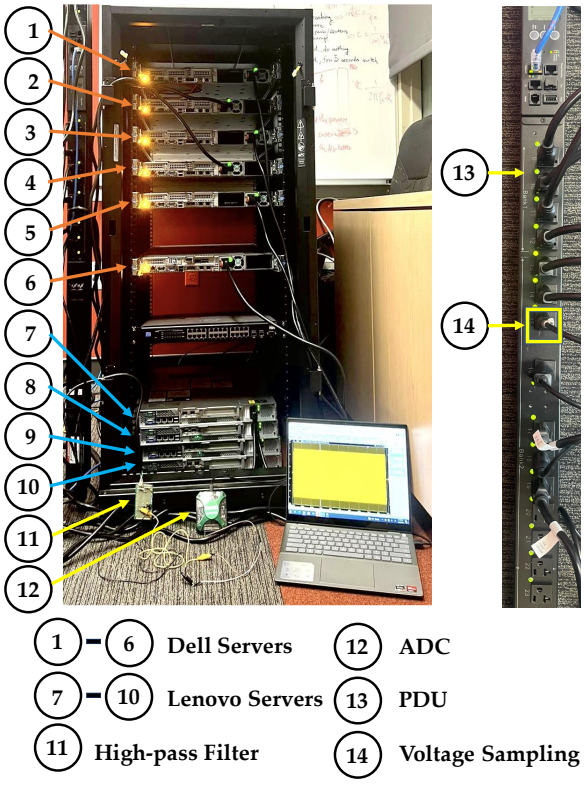

Figure 15: Experimental setup.

5 PERFORMANCE EVALUATION

5.1 Settings

Experiment setup. We utilize a cluster of six Dell PowerEdge R640 and four Lenovo ThinkSystem SR250 rack servers. The Dell servers are each equipped with two Intel Xeon CPUs, 128GB of memory, and 750W power supplies. The Lenovo servers are each equipped with one Intel Xeon CPU, 32GB memory, and 300W power supplies. The servers are powered by a 3KW APC rack PDU with metered outlets [26]. We use a Digilent USB oscilloscope with a second-order RC highpass filter to collect the voltage samples from one of the PDU outlets. We collect the server power consumption from our APC PDU using command-line queries at a rate of one sample per second. We use a laptop to collect and store our voltage and power data. We put our servers in a server rack and placed them in our lab. Fig. 15 shows our experiment setup.

Server power traces. We vary server power by stressing a varying number of CPU cores. We stress the servers following different workload traces collected from various sources, such as Google's product traffic [74], HTTP request to a university website [33], Wikipedia workload [75], and Verizon Teramark's UPS-level power consumption [33]. We scale these workload traces to each server's idle and peak power.

5.2 Results

Extended evaluation. We run our server cluster for five days, each of the ten servers with different workload traces. We then apply our EMI-to-Power algorithm to the voltage data collected from the APC rack PDU. Fig. 16 shows the original power and EMI-based power estimation of all the servers in our experiment. Fig. 17 shows the corresponding error statistics of the extended experiment. The box plot in Fig. 17 extends from the first quartile (Q1) to the third quartile (Q3) of the data, with a black line at the median and the mean represented using a diamond marker. The whiskers extend from the box by 1.5x the inter-quartile range (IQR). For error statistics, we have chosen absolute percentage error. We see that the EMI-based power monitoring can reasonably extract the power consumption of all servers with less than ~7% mean absolute error using only one voltage sensor.

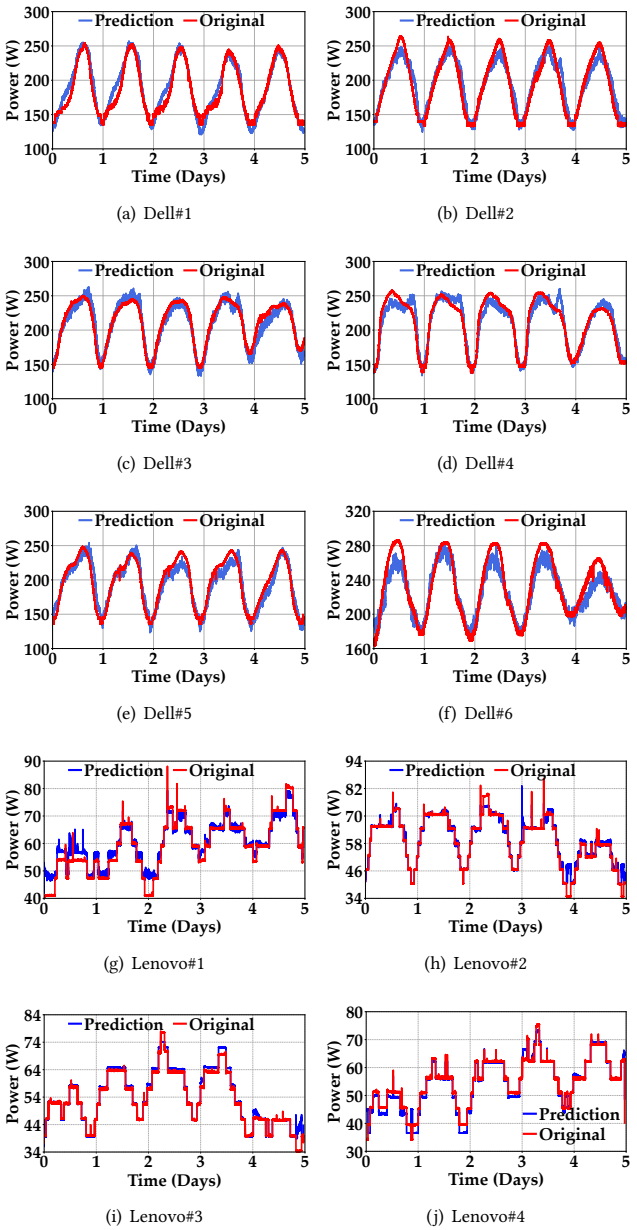

Figure 16: Original vs. EMI-based estimation of power consumption of a cluster of Dell and Lenovo servers.

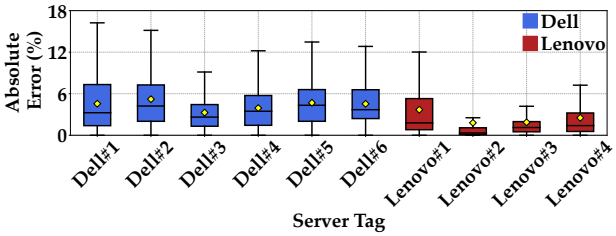

Figure 17: Error statistics of EMI-based power monitoring for an extended period.

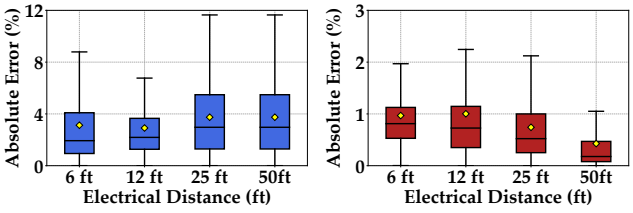

(a) Dell (b) Lenovo

Figure 18: Impact of the electrical distance between the servers and voltage sensing on power monitoring.

Impact of the electrical distance between servers and voltage sensor. In data centers with many servers, the voltage sensor location may be far away from the server power outlet. To see how the electrical distance between the voltage sensor and the server affects the EMI extraction, we run our servers at varying distances from the voltage sensor using power extension cables. For each experiment, we run the servers for 12 hours with varying power consumption and show the error statistics in Fig. 18. We see that with an increase in distance, the Dell server's power estimation suffers, while the Lenovo server performs better. While we find that with increasing distance, the EMI frequency amplitude (and hence the SNR) decreases for both Dell and Lenovo, we cannot offer an explanation here as to why the SNR loss affects the power monitoring of these two types of servers differently. We need further investigation and leave it as a future work.

Impact of outlet phase. Large data centers use a three-phase power delivery system, even when the servers are single-phase. The electric phases are EMI isolated from each other. To demonstrate that, we utilize the phase difference in the power outlets in our lab. We connect the Dell servers in two different phases and collect the voltage data from the respective outlets while running the servers. Fig. 19(a) shows the frequency domain analysis and the difference of EMIs observed in the two phases. Fig. 19(b) shows the corresponding error statistics of power estimation.

Impact of rapid power change. Data centers frequently encounter abrupt surges in workload, leading to rapid fluctuations in server power consumption. In order to assess the effectiveness of our approach in handling rapid changes in power consumption, we conduct experiments using both Dell and Lenovo servers. We create rapid power change by stressing the servers to their peak from idle power. With this, the Dell server has a 120W power change, while the Lenovo server has a 50W power change. Fig. 20 shows the predicted and original power consumption of the servers. We see in Fig. 20(a) that the predicted power lags the original power and

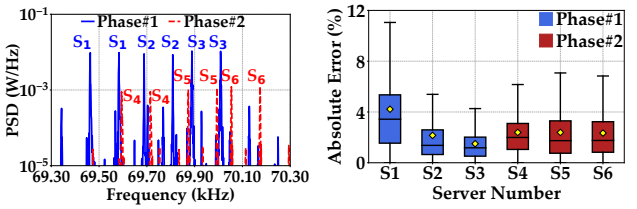

(a) Servers at different phase (b) Impact of phase on accuracy

Figure 19: Server EMIs from power outlets at different phases.

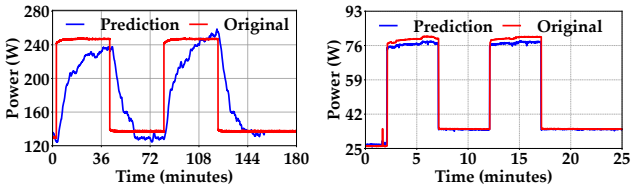

(a) Dell (b) Lenovo

Figure 20: Impact of rapid power variation

takes nearly 30 minutes to reach the peak. Meanwhile, as shown in Fig. 20(b), the predicted power changes instantaneously for the Lenovo server. This indicates that our current EMI-based power monitoring approach cannot follow rapid change for FM-PFCs. Nevertheless, we would like to point out that we applied a linear EMI-to-power model with one variable and first-order statistics. We plan to improve our EMI-to-power model for the FM-PFC in the future, incorporating second-order and time series statistics.

6 DISCUSSION

6.1 Scalability of EMI-Based Power Monitoring

Our results reveal that server EMI signatures exhibit *approximate orthogonality*, where orthogonality is *likely but not guaranteed*. Since the variation in Power Factor Correction (PFC) switching frequency is due to random process variation in the PFC component manufacturing, it is possible to have multiple servers with EMI signatures at the same frequency. The possibility of such complete overlap (i.e., non-orthogonality) increases with the number of servers. This poses scalability challenges in our server-level power monitoring approach as our proposed EMI signature separation is built on frequency domain orthogonality of server EMIs. To understand the extent of this scalability issue, we first discuss the real-life scalability requirements and the likelihood of fully overlapped cases.

Scalability requirements in data centers. The data center power infrastructure bounds the practical scalability requirement in EMI-based power monitoring. Our EMI-based server-level power monitoring captures the EMI signature from the servers within the *same power network*. Transformers in the power network act as low-pass filters and confine the high-frequency server EMIs within its secondary/load side. Hence, in our context of EMI sensing, each transformer secondary creates an isolated power network. Moreover, each phase also acts as a separate power network for three-phase transformers (Fig. 19(a)). Now, the power infrastructure of large data centers (e.g., >1MW power capacity) is divided into several clusters of server racks powered by cluster PDUs. Cluster PDUs use transformers and create isolated power networks where

the server EMIs from one network do not propagate to another. Hence, the cluster PDUs impose the upper limit on the number of servers in the same power network. Consequently, *our server monitoring approach needs to scale up only to the maximum number of servers powered by a single phase of the cluster PDU.* For instance, with a 150kVA three-phase cluster PDU, the upper limit for scaling is 200 servers, each server with a 250W power requirement.

Handling servers with overlapped EMI. Data center scale applications are typically run on multiple servers in parallel, resulting in highly correlated power consumption patterns among various servers. Consequently, we can separate a fully overlapped server signature if its power consumption is highly correlated with the power consumption of another server(s) without the EMI overlap (and can be isolated using frequency-domain orthogonality). We can identify the correlated servers by analyzing the software implementation and workload allocation of the data center applications in combination with long-term (e.g., a week) power consumption.

Remarks on the scalability concern in practice. Due to the bounded scalability requirement, we anticipate a limited number of fully overlapped EMI cases, and we do not expect fully overlapped server EMI to be a widespread issue, especially in colocation data centers where the servers are purchased by different tenants (possibly from different vendors and through different supply chains), resulting in a more diverse power supply EMIs in data centers. Meanwhile, smaller edge data centers with a few tens of servers naturally suffer less from fully overlapped server EMI cases. Hence, *we believe both the extent and impact of the scalability issue are unlikely to affect our approach's usefulness in practice significantly.*

6.2 Limitations of Implementation and Experimental Validation

Algorithm for server-wise EMI separation. Our current implementation presented in Section 4.2 is based on heuristics. Moreover, our algorithm mainly relies on frequency-domain measurements of conducted EMI. EMI-based power monitoring from voltage measurement can be further refined by applying other more sophisticated and robust signal processing techniques such as wavelet analysis, time-frequency analysis, and adaptive filtering.

Validation scale. The scale of our experiments, albeit on commercial equipment, is limited to a small cluster of ten servers. Hence, our experiment setup cannot capture the scalability concern of large data centers with hundreds, even thousands, of servers. Moreover, we run our experiments on products from two server vendors - Dell and Lenovo. While PFC circuits are mandated to be present in every server power supply, and Dell and Lenovo are the world's leading server manufacturers [76, 77], the general applicability of our approach in any data center remains unvalidated.

6.3 Cost of Implementation

In our experiments, we utilize expensive off-the-shelf equipment such as an oscilloscope and a laptop for voltage measurement and running the EMI-to-power algorithm (Fig. 15). The cost of our testing equipment in this paper, however, is not representative of the cost of real-world implementation. Hence, we here offer an estimation of the implementation cost for our solution based on a high-level design of our EMI sensing for power monitoring.

Design of the sensing prototype. Our sensor solution has three components - data acquisition, processing, and communication. For data acquisition, we mainly need a small circuit with analog filters and an ADC. For data processing, we need to run real-time frequency domain transformation followed by an EMI-to-power conversion algorithm. Since our signal of interest only occupies a frequency band of a few kilohertz, we can use frequency shifting to reduce the sampling rate and, therefore, the computation load of the data processing. Finally, we can use Bluetooth or WiFi to wirelessly send our power monitoring data.

Estimated cost of sensor components. The Raspberry Pi Pico [78] featuring a Dual-core Arm Cortex-M0+ processor with onboard ADC, 802.11n wireless LAN, Bluetooth 5.2, flexible clock speeds of up to 133 MHz, and 264kB of on-chip SRAM, is an ideal candidate for our sensor prototype with enough (estimated) processing and memory capacity. Currently, a single unit of Raspberry Pi Pico costs $7 [79]. We estimate the cost of our analog filters consisting of resistors and capacitors to be ~$3. We attribute another ~$10 for packaging and micro-USB power for the Raspberry Pi Pico [80], resulting in a total estimated cost of one sensor prototype to $20.

Cost comparison. The smart PDU from APC with 24 outlets offers one of the cheapest commercially available solutions for outlet/server-level power monitoring. But it is priced at more than $2000, resulting in a cost of ~$83 per server. In comparison, we have already demonstrated power monitoring of ten servers with one sensor with a per server cost of $2. If we extend our approach to its maximum limit of 200 servers per sensor (Section 6.1), we have an estimated cost $0.1 per server.

In the cost of power monitoring per server, our approach is a clear winner. However, we must note here that the APC smart PDU has other functionalities beyond power monitoring, such as turning outlets on and off remotely. The same applies to power monitoring using server onboard management modules [23, 25], which offer even more functionalities such as remote start and life-cycle management. Hence, we cannot offer an apple-to-apple cost comparison here. Nevertheless, we believe the sheer cost difference makes our approach appealing for server power monitoring. More importantly, for colocation, we remain the only non-intrusive solution for monitoring the tenant's server-level power.

7 CONCLUSION

In this paper, we proposed a novel low-cost EMI-based server power monitoring in a data center. Our approach utilizes the conducted EMI generated from the server power supply's PFC circuits. We presented a theoretical characterization of the EMI generation and propagation. We developed a heuristic-based algorithm for server-wise EMI separation. We offered validating results and evaluation using a set of commercial-grade servers and real workload traces. We demonstrated that our approach can estimate the server-level power with less than 7% mean absolute error.

ACKNOWLEDGMENTS

This work is supported in parts by the US National Science Foundation under grant numbers ECCS-2152357, CCF-2324915, and ECCS-2132112.

REFERENCES

[1] P. Gupta, Z. Talukder, M. A. Islam, and P. Nguyen, "Towards server-level power monitoring in data centers using single-point voltage measurement," in *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems*, SenSys '22, (New York, NY, USA), p. 855–856, Association for Computing Machinery, 2023.

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

[3] 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.

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

[5] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska, "Dynamic right-sizing for power-proportional data centers," in *INFOCOM*, 2011.

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

[7] M. Skach, M. Arora, C.-H. Hsu, Q. Li, D. Tullsen, L. Tang, and J. Mars, "Thermal time shifting: Leveraging phase change materials to reduce cooling costs in warehouse-scale computers," in *ISCA*, 2015.

[8] I. Manousakis, I. n. Goiri, S. Sankar, T. D. Nguyen, and R. Bianchini, "Coolprovision: Underprovisioning datacenter cooling," in *SoCC*, 2015.

[9] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs, "Cutting the electric bill for internet-scale systems," in *SIGCOMM*, 2009.

[10] D. Meisner, B. T. Gold, and T. F. Wenisch, "Powernap: Eliminating server idle power," in *ASPLOS*, 2009.

[11] C. Lefurgy, X. Wang, and M. Ware, "Server-level power control," in *Fourth International Conference on Autonomic Computing (ICAC'07)*, pp. 4–4, IEEE, 2007.

[12] X. Wang, M. Chen, C. Lefurgy, and T. W. Keller, "Ship: Scalable hierarchical power control for large-scale data centers," in *PACT*, 2009.

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

[14] J. Moore, J. Chase, P. Ranganathan, and R. Sharma, "Making scheduling "cool": Temperature-aware workload placement in data centers," in *USENIX ATC*, 2005.

[15] C. Bash and G. Forman, "Cool job allocation: Measuring the power savings of placing jobs at cooling-efficient locations in the data center.," in *USENIX Annual Technical Conference*, vol. 138, p. 140, 2007.

[16] 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.

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

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

[19] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels, "Dynamo: amazon's highly available key-value store," in *SOSP*, 2007.

[20] Raritan, "White paper: Data center power usage," https://www.raritan.com/landing/power-monitoring-and-metering-white-paper/thanks.

[21] J. Brown, "Efficient data center power monitoring made easy," https://www.rfcode.com/data-driven-data-center/bid/309767/efficient-data-center-power-monitoring-made-easy.

[22] Electronic Environments, "Data center power monitoring & temperature monitoring," https://www.eecnet.com/solutions/data-center-monitoring/data-center-power-monitoring.

[23] Dell, "PowerEdge R740xd Rack Server," https://www.dell.com/en-us/work/shop/povw/poweredge-r740xd.

[24] G. Tang, W. Jiang, Z. Xu, F. Liu, and K. Wu, "Zero-cost, fine-grained power monitoring of datacenters using non-intrusive power disaggregation," in *Proceedings of the 16th Annual Middleware Conference*, Middleware '15, (New York, NY, USA), p. 271–282, Association for Computing Machinery, 2015.

[25] Patsy K. Popa , "Managing server energy consumption using ibm powerexecutive," http://images.incisivemedia.com/v7_static/pdf/vnu/optit-wp-stg-power-executive.pdf.

[26] APC, "Metered-by-outlet with switching rack PDU," http://goo.gl/qvE8NV.

[27] CDW, "Raritan dominion px px3-5903v - power distribution unit - 5000 va," https://www.cdw.com/product/raritan-dominion-px-px3-5903v-power-distribution-unit-5000-va/5001493.

[28] SynapSense , "Power monitoring," https://www.vigilent.com/products-and-services/monitoring.

[29] M. Y. Lim, A. Porterfield, and R. Fowler, "Softpower: fine-grain power estimations using performance counters," in *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing*, pp. 308–311, ACM, 2010.

[30] D. Economou, S. Rivoire, C. Kozyrakis, and P. Ranganathan, "Full-system power analysis and modeling for server environments," International Symposium on Computer Architecture-IEEE, 2006.

[31] R. Ge, X. Feng, S. Song, H.-C. Chang, D. Li, and K. W. Cameron, "Powerpack: Energy profiling and analysis of high-performance systems and applications," *IEEE Trans. Parallel and Dist. Systems*, vol. 21, pp. 658–671, May 2010.

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

[33] M. A. Islam, H. Mahmud, S. Ren, and X. Wang, "Paying to save: Reducing cost of colocation data center via rewards," in *HPCA*, 2015.

[34] J. dePreaux, "Wholesale and retail data centers - North America and Europe - 2013," *IHS*, Jul. 2013, https://technology.ihs.com/api/binary/492570.

[35] CyrusOne, "Colocation: The logical home for the cloud," 2012, http://resources.idgenterprise.com/original/AST-0050996_CyrusOne_wp0926.pdf.

[36] T. Hegarty, "An overview of conducted emi specifications for power supplies," *Texas Instruments white paper SLYY136*, 2018.

[37] F. C. Commission *et al.*, "Code of federal regulations title 47 part 15—radio frequency devices," *Online Document] FCC, June*, vol. 16, 1986.

[38] B. Mammano and B. Carsten, "Understanding and optimizing electromagnetic compatibility in switchmode power supplies," in *2002 Power Supply Design Seminar Manual SEM1500*, p. 18, 2002.

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

[40] M. A. Islam and S. Ren, "Ohm's law in data centers: A voltage side channel for timing power attacks," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pp. 146–162, 2018.

[41] T. Hegarty, "An overview of radiated EMI specifications for power supplies," *Texas Instruments Whitepaper*, June 2018.

[42] Z. Shao, M. A. Islam, and S. Ren, "Your noise, my signal: Exploiting switching noise for stealthy data exfiltration from desktop computers," in *Abstracts of the 2020 SIGMETRICS/Performance Joint International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS '20, (New York, NY, USA), p. 79–80, Association for Computing Machinery, 2020.

[43] S. Abdel-Rahman, "Ccm pfc boost converter design," *Infineon, Tech. Rep. DN 2013–01*, 2013.

[44] Y. Mo and B. Sinopoli, "False data injection attacks in control systems," in *Preprints of the 1st workshop on Secure Control Systems*, pp. 1–6, 2010.

[45] P.-Y. Chen, S.-M. Cheng, and K.-C. Chen, "Smart attacks in smart grid communication networks," *IEEE Communications Magazine*, vol. 50, no. 8, pp. 24–29, 2012.

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

[47] F. Yang, M. Islam, F. Wu, and S. Ren, "Compkey: Exploiting computer's electromagnetic radiation for secret key generation," in *2021 IEEE Conference on Communications and Network Security (CNS)*, IEEE, 2021.

[48] M. A. Islam, S. Ren, N. Pissinou, A. H. Mahmud, and A. V. Vasilakos, "Distributed temperature-aware resource management in virtualized data center," *Sustainable Computing: Informatics and Systems*, vol. 6, pp. 3–16, 2015.

[49] Z. Shao, M. A. Islam, and S. Ren, "Heat behind the meter: A hidden threat of thermal attacks in edge colocation data centers," in *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 318–331, IEEE, 2021.

[50] M. A. Islam, S. Ren, N. Pissinou, H. Mahmud, and A. Vasilakos, "Distributed resource management in data centers with temperature constraint," in *IGCC*, 2013.

[51] R. Callan, N. Popovic, A. Daruna, E. Pollmann, A. Zajic, and M. Prvulovic, "Comparison of electromagnetic side-channel energy available to the attacker from different computer systems," in *2015 IEEE International Symposium on Electromagnetic Compatibility (EMC)*, pp. 219–223, IEEE, 2015.

[52] A. Nazari, N. Sehatbakhsh, M. Alam, A. Zajic, and M. Prvulovic, "Eddie: Em-based detection of deviations in program execution," in *Proceedings of the 44th Annual International Symposium on Computer Architecture*, pp. 333–346, 2017.

[53] Open Science Framework (OSF), "Server power monitoring using conducted emi," https://osf.io/qtsb7/?view_only=9f8d7351d08a4058add72331201f0724.

[54] Lenovo, "Flex System Enterprise Chassis," https://www.lenovo.com/us/en/data-center/servers/flex-blade-servers/chassis/c/chassis-rack.

[55] N. Rasmussen, "Ac vs. dc power distribution for data centers," *APC White Paper 63*.

[56] RackSolutions, "Blade server vs rack server," https://www.racksolutions.com/news/data-center-optimization/blade-server-vs-rack-server/.

[57] P. Hu, "Electrical distribution equipment in data center environments," *APC White Paper 61*.

[58] Energy Star, "Computers specification version 7.0," 2018, https://www.energystar.gov/products/spec/computers_specification_version_7_0_pd.

[59] IECEE, "IEC 61000-3-2:2018: Electromagnetic compatibility (EMC) - part 3-2," http://www.onsemi.com/pub/Collateral/HBD853-D.PDF.

[60] Electronic Code of U.S. Federal Regulations, "Unintentional radiators, section 15.107 — conducted limits," 2018.

[61] S. DeBruin, B. Ghena, Y.-S. Kuo, and P. Dutta, "Powerblade: A low-profile, true-power, plug-through energy meter," in *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, pp. 17–29, 2015.

[62] W. L. Bircher and L. K. John, "Complete system power estimation: A trickle-down approach based on performance events," in *2007 IEEE International Symposium on Performance Analysis of Systems & Software*, pp. 158–168, IEEE, 2007.
[63] S. Rivoire, P. Ranganathan, and C. Kozyrakis, "A comparison of high-level full-system power models.," *HotPower*, vol. 8, no. 2, pp. 32–39, 2008.
[64] X. Fan, W.-D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," in *ISCA*, 2007.
[65] J. Zhang, X. Ji, Y. Chi, Y.-c. Chen, B. Wang, and W. Xu, "OutletSpy: cross-outlet application inference via power factor correction signal," in *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, WiSec 2021, June 2021.
[66] S. Gupta, M. S. Reynolds, and S. N. Patel, "Electrisense: Single-point sensing using emi for electrical event detection and classification in the home," in *UbiComp*, 2010.
[67] M. Enev, S. Gupta, T. Kohno, and S. N. Patel, "Televisions, video privacy, and powerline electromagnetic interference," in *CCS*, 2011.
[68] M. Levy, "Fourier transform analysis," *Journal of the British Institution of Radio Engineers*, vol. 6, no. 6, pp. 228–246, 1946.
[69] TRIPP-LITE, "Isolator series 120v 1000w isolation transformer-based power conditioner, 4 outlets," https://tripplite.eaton.com/isolator-series-120v-1000w-isolation-transformer-based-power-conditioner-4-outlets~is1000.
[70] Tektronix, "Tektronix mso54-5 series," https://www.tek.com/en/products/oscilloscopes/5-series-mso.
[71] N. Instruments, "Ni daq usb 6361," https://www.ni.com/en-us/shop/hardware/products/multifunction-io-device.html?skuId=15665.
[72] Digilent, "Digilent analog discover-2," https://digilent.com/shop/analog-discovery-2-100ms-s-usb-oscilloscope-logic-analyzer-and-variable-power-supply.
[73] Mathworks, "Matlab - find local maxima," https://www.mathworks.com/help/signal/ref/findpeaks.html.
[74] Google, "Google transparency report," https://transparencyreport.google.com/traffic/overview?hl=en.
[75] G. Urdaneta, G. Pierre, and M. Van Steen, "Wikipedia workload analysis for decentralized hosting," *Computer Networks*, 2009.
[76] J. Collins, "The 10 largest server companies in the world, and what they do."," https://history-computer.com/largest-server-companies-in-the-world-and-what-they-do.
[77] T. P. Morgan, "No server recession at lenovo and supermicro so far."," https://www.nextplatform.com/2023/06/09/no-server-recession-at-lenovo-and-supermicro-so-far/#:~:text=Yuanqing%20added%20that%20based%20on,among%20external%20storage%20array%20makers.
[78] R. Pi, "Raspberry pi pico," https://www.raspberrypi.com/products/raspberry-pi-pico.
[79] Sparkfun, "Price of raspberry pi pico," https://www.sparkfun.com/products/20174?src=raspberrypi.
[80] Amazon, "Price of power supply," https://www.amazon.com/Bawofu-Supply-Charger-Adapter-Universal/dp/B0BSF3TB8P/ref=sr_1_8?keywords=Bawofu&qid=1696323242&sr=8-8.