

Poster Abstract: Towards Server-Level Power Monitoring in Data Centers Using Single-Point Voltage Measurement

Pranjol Gupta, Zahidur Talukder, Mohammad A. Islam, and Phuc Nguyen pranjolsen.gupta@mavs.uta.edu;zahidurrahim.talukder@mavs.uta.edu;mislam@uta.edu;vp.nguyen@uta.edu
The University of Texas at Arlington

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 of server power supplies to measure its power consumption from non-intrusive single-point voltage measurement. Using a pair of commercial grade Dell PowerEdge servers, we demonstrate that our approach can estimate each server's power consumption with ~3% mean absolute percentage error.

1 INTRODUCTION

Motivation. Power monitoring is a core component in efficient data center management, and towards that, fine-grained server-level power monitoring plays a crucial role. It can facilitate advanced server power and cooling management techniques by providing real-time power consumption information. Moreover, fine-grained power monitoring can play a vital role in safeguarding data centers from outages/downtime due to overloading as well as identifying malicious server behavior caused by an attacker.

Limitations of existing systems. Despite its compelling benefits, server-level power monitoring is not widely adopted. Meanwhile, existing power monitoring systems for data centers focus mainly on the UPS and cluster-level power monitoring. The cost of installing dedicated power meters for every server has been a critical barrier to achieving server-level power monitoring. Also, existing/traditional power monitoring systems cannot monitor the server-level power consumption of a tenant's server rack in colocation data centers. With nearly 40% market share, colocation data centers provide a shared data center solution where the tenants bring their physical servers inside the data center while the data center manager is responsible for providing power, cooling, and space. In colocations, the data center manager lacks access to tenants' server racks to deploy any existing server-level monitoring.

Our contribution. This paper takes the foundational steps towards a novel data center power monitoring approach that overcomes the above mentioned limitations. In our approach, we take the voltage measurement from a data center cluster-level PDU

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SenSys '22, November 6–9, 2022, Boston, MA, USA © 2022 Copyright held by the owner/author(s). ACM ISBN 978-1-4503-9886-2/22/11. https://doi.org/10.1145/3560905.3568079

Existing Power Metering

Cluster PDU

Utility

ATS

Server Power Distribution Unit

ATS=Automatic Transfer Switch, UPS=Uninterrupted Power Supply, PDU-Power Distribution Unit

Figure 1: Overview of our server-level power monitoring approach. We take the voltage measurement at the cluster level to disaggregate server-wise power consumption.

(Fig. 1) and, from this "single-point measurement" extract the power consumption of all the servers connected to that PDU. Hence, the proposed power meter significantly reduces the power monitoring cost by eliminating the need for a dedicated power sensor for every server. It also enables server-level power monitoring in co-locations as the voltage data is collected from the cluster PDUs operated by the data center manager.

Working principle. We identify that the power factor correction (PFC) circuit is ubiquitously present in all server power supplies due to government-mandated EMI regulations [1], and these PFC circuits cause conducted electromagnetic interference (EMI) at high-frequency range (e.g., 10kHz~150kHz) [2]. The conducted EMIs travel through the data center power distribution network and can be extracted from the voltage measurement. More importantly, these conducted EMIs change with the server power and, hence, can be used to estimate the server power. Meanwhile, we also observe that the conducted EMIs of a server create narrow frequency spikes (e.g., 1~3 Hz) in the frequency domain, and EMIs from different servers generally are found at different frequencies, even among the servers from the same batch [2]. The combination of non-identical and narrow-band server EMI makes each server's EMI distinguishable amidst many other servers' EMIs, enabling the extraction of server-wise EMIs from a single voltage sensor.

Results. Using two commercial-grade Dell PowerEdge servers, we demonstrate that our approach can estimate each server's power consumption with $\sim 3\%$ mean absolute percentage error.

2 SERVER-LEVEL MONITORING SYSTEM

Conducted EMI Generation in Server Power Supplies. The conducted EMI generated by the PFC circuit results from rapid electronic switching to "shape" the input current of the server power supply into a sinusoid that matches the input voltage [2]. Since the power supply's input current flows through the data center power network, it creates voltage ripples due to line voltage drops following Ohm's Law [2]. More importantly, however, the ripples are generated at a much higher frequency (40 ~100kHz) than the 50/60Hz nominal grid frequency. The voltage ripples can be easily identified in the voltage readings within the power network. Fig. 2(a)




Figure 2: EMI spike at PFC frequency and power variation shows the frequency components of the voltage measurement at a nearby power outlet of a Dell PowerEdge R640 server with PFC. We can see the EMI spikes caused by the PFC circuit.

Relationship Between Conducted EMI and Server Power To extract the relationship between the server power and the conducted EMI, we run our Dell server for 20 minutes at different power levels and collect the voltage data at 250K samples/second. We apply Fast Fourier Transform (FFT) on the voltage data using one-second windows. We observe that when the server power changes, the frequency of its conducted EMI also changes. More specifically, we find that the EMI frequency decreases as the power consumption increases. Fig. 2(b) shows the relationship between the server power and the EMI frequency. We use a linear fitting, $p(t) = A \cdot f(t) + B$, where p(t) and f(t) are the server power and EMI frequency, respectively, at time t. A and B are model coefficients. We run the same experiment with another Dell server and find that a similar relationship of change in the EMI frequency with server power change exists, albeit the location of the EMI frequency is different for the two servers, i.e., A and B are different.


EMI-Based Server Power Monitoring As shown in Fig. 1, we take the voltage measurement from a common point of interconnection (e.g., cluster PDU) in the data center power network. We then convert our time domain voltage data into frequency domain components using FFT. We separate the EMIs of different servers utilizing their frequency domain orthogonality. However, each server has a different EMI frequency-to-power model, and it is impractical to individually run experiments on every server to extract their EMI vs. power model. Hence, we propose to use the cluster-level power measurements (from existing power monitoring systems in data centers) that give us the aggregate power to extract the model parameters of the servers without running separate experiments. We can write the cluster-level power measurement $p^{cluster}(t)$ for N servers connected to it as

$$p^{cluster}(t) = \sum_{i=1}^{N} (A_i \cdot f_i(t) + B_i), \tag{1}$$

where A_i , B_i , and $f_i(t)$ are the model coefficients and EMI frequency of the i-th server, respectively. By collecting more than N samples of the server EMI frequencies and cluster level power, we can extract the model coefficients of all servers from Eqn. (2).

3 EVALUATION

Experimental Setup We use two Dell PowerEdge R640 servers, each equipped with two Intel Xeon CPUs, 128GB of memory, and a 750W power supply. We use a National Instrument DAQ (PXIE-6366) for voltage data collection at a sampling rate of 250K samples/second. To reduce quantization error and lower the line voltage to fit our DAQ's maximum input range, we use an RC high-pass



(a) Experiment setup

(b) EMI of two Dell servers

Figure 3: (a) Experiment setup - ① two Dell servers, ② PDU/power strip, ③ high-pass filter, ④ NI DAQ for voltage data collection, and ⑤ laptop for data processing. (b) The EMI spikes of the two dell servers.




Figure 4: The actual and estimated power of the two Dell servers with their mean absolute percentage error (MAPE).

filter with 3kHz cutoff frequency. Our experiment setup is shown in Fig. 3(a). We create power load in our server by running CPU-intensive instructions. We run the two servers simultaneously for 18 hours. We use Matlab to collect the voltage data using the DAQ continuously. We use our intelligent metered PDU to measure each server's power to serve as the ground truth.

Results Fig. 3(b) shows the EMI frequencies of the two Dell servers. We see that they have different EMI frequencies. Next, in Fig. 4 we show the two servers' actual power and estimated power using our voltage measurement-based approach. We see that our power monitoring system can estimate the power with good accuracy having a MAPE of 2.38% and 3.11% for Server#1 and Server#2, respectively. Our results demonstrate the viability of using a conducted EMI-based server-level power monitoring.

4 CONCLUSION AND FUTURE WORK

This paper demonstrated the feasibility of using a single-point voltage measurement-based power monitoring system for data centers. Using a pair of commercial-grade servers, we showed that our approach could monitor server power with a MAPE of ~3%. Here, we utilize frequency-domain orthogonality of conducted EMI to separate different servers. However, such orthogonality cannot be guaranteed in practice, especially for large data centers with many servers. This poses scalability challenges in our proposed approach. Hence, our plan is to develop a robust frequency separation algorithm that works with many servers' EMI spikes utilizing contextual information such as the similarity of power consumption among multiple servers hosting the same cloud application. In addition, we also plan to develop a low-cost voltage sensor that will be able to do real-time frequency analysis and power extraction.

REFERENCES

- T. Hegarty, "An overview of conducted emi specifications for power supplies," Texas Instruments white paper SLYY136, 2018.
- M. A. Islam and S. Ren, "Ohm's law in data centers: A voltage side channel for timing power attacks," in CCS, 2018.

2