WattScope: Non-intrusive Application-level Power Disaggregation in Datacenters

Xiaoding Guan, Noman Bashir, David Irwin, Prashant Shenoy
University of Massachusetts Amherst

ABSTRACT

WattScope is a system for non-intrusively estimating the power consumption of individual applications using external measurements of a server's aggregate power usage and without requiring direct access to the server's operating system or applications. Our key insight is that, based on an analysis of production traces, the power characteristics of datacenter workloads, e.g., low variability, low magnitude, and high periodicity, are highly amenable to disaggregation of a server's total power consumption into application-specific values. WattScope adapts and extends a machine learning-based technique for disaggregating building power and applies it to server- and rack-level power measurements that are already available in datacenters. We evaluate WattScope's accuracy on a production workload and show that it yields high accuracy, e.g., often $<\sim 10\%$ normalized mean absolute error, and is thus a potentially useful tool for datacenters in externally monitoring application-level power usage.

Keywords

Carbon-efficiency, energy-efficiency, disaggregation

1. OUR CONTRIBUTIONS

Datacenters' rapid increase in capacity has led to increasing concern and criticism over their energy consumption and carbon footprint. As a result, many cloud providers and datacenter operators have begun to increase their emphasis on energy-efficient and sustainable operations. Indeed, prominent technology companies have set ambitious goals to eliminate their carbon footprint within the next 10-20 years [1, 5, 4]. Importantly, *the simplest and most basic prerequisite for optimizing a datacenter's energy- and carbon-efficiency is providing applications visibility into their power consumption, as they cannot optimize a metric they cannot measure.* Datacenters are well-instrumented with external power meters typically attached to rack-level power distribution units (PDUs) and individual servers. However, rack- and server-level power monitoring does not provide insight into the power consumed by individual applications, since servers are multi-tenant and host multiple applications.

A common approach to application-level power monitoring is to train a model that takes per-process hardware performance counters as input and infers a corresponding power usage. For example, PowerAPI is an open-source toolkit that uses such techniques to monitor application-level power [3]. While such approaches require privileged access to the hardware performance counters, there are many scenarios where access to hardware counters is either not available or too intrusive. In particular, cloud users often lack privileged access to hardware counters; process-level power monitoring may incur high overhead when monitoring large numbers of processes, especially at high resolution, and hardware interfaces are not standardized, and thus must be tailored to specific hardware platforms. These limitations are the primary reason that fine-grained application-level power monitoring is not offered by cloud providers. This lack of support in-turn prevents cloud applications from optimizing their energy consumption and carbon emissions. Ultimately, the lack of application-level visibility into energy consumption is a key impediment to achieving the ambitious sustainability goals above, as it is impossible to optimize a metric that cannot be effectively measured.

To address the problem, we design WattScope, a system for non-intrusively monitoring application-level power consumption using aggregate server-level power measurements. WattScope uses disaggregation techniques to apportion power data from external server- and rack-level power meters, which are typically available in power distribution units (PDUs), into individual application-level power usage without requiring intrusive access to system and application software. WattScope recognizes that datacenters already collect server- and rack-level power data for thermal management and billing purposes, which can be leveraged to also provide application-level power monitoring. Thus, WattScope analyzes power data collected from these external meters to infer each application's power usage.

More formally, WattScope disaggregates a time-series of power readings $P(t)$, over some sampling interval Δt, into a separate time-series $p_i(t)$ for each individual application i, such that $\forall t, P(t) = \sum_i p_i(t)$. Since WattScope does not require any server-level access or specific hardware/software support, it can run externally as part of the facility management system. As a result, WattScope can be deployed in nearly any datacenter facility with PDUs that measure server- and rack-level power. Our key insight is that, based on a large-scale analysis of production traces, the power characteristics of datacenter workloads, e.g., low variability, low magnitude, and high periodicity, are highly amenable to disaggregation. WattScope adapts and extends a deep learning-based technique, originally designed for disaggregating building power, and applies it to servers and racks. We implement WattScope and experimentally evaluate its

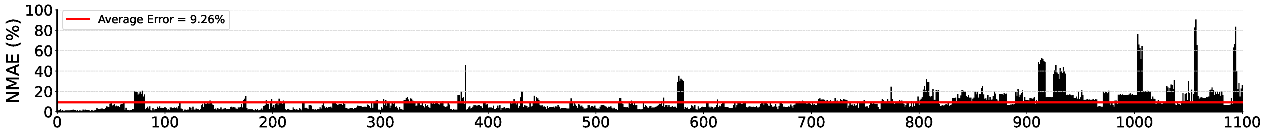

Figure 1: ***Normalized Mean Absolute Error (NMAE) in disaggregating a job's power consumption on the y-axis for 1100 servers from Google trace on the x-axis. The average error across all the servers is 9.26%. Each server runs 40 jobs on average. Servers are sorted in the order of increasing Coefficient of Variation (CoV) for the disaggregated job from 0.01 (left most) to 3.75 (right most).***

accuracy on a production workload.

Our hypothesis is that disaggregating server- and rack-level power using `WattScope` can enable highly accurate and non-intrusive application-level power monitoring without requiring any server-level hardware/software support. In evaluating our hypothesis, we make the following contributions.

Production Workload Analysis. We first analyze the job characteristics of a large-scale production workload from a major cloud provider that includes 5-minute resource usage readings for 2.7 million jobs over a 30 day period encompassing more than 100 million job-hours. Our analysis reveals that job usage patterns exhibit multiple characteristics, including low variability, low magnitude, and high periodicity, that `WattScope` can potentially exploit for disaggregation. Our analysis also shows that, while server applications can operate arbitrarily and irregularly in general, they have a high degree of regularity in practice.

WattScope Design. We present `WattScope`'s design, which adapts and extends a deep learning-based disaggregation algorithm originally applied to building power data. `WattScope`'s design includes a library of models trained for different classes of applications based on their variability, magnitude, and periodicity. `WattScope` then integrates with a cluster scheduler to learn the number and type of applications running on each server, i.e., based on their attributes, to select an appropriate model for disaggregation.

Implementation and Evaluation. Finally, we implement and evaluate a `WattScope` prototype. We implement `WattScope`'s disaggregation technique by modifying nilmtk-contrib [2], an open-source reference implementation of multiple algorithms for building energy disaggregation, to instead disaggregate server- and rack-level power, and evaluate accuracy across multiple dimensions using our production workload trace. We evaluate `WattScope`'s accuracy on a production workload and show that it yields high accuracy, e.g., often $<\sim$10% normalized mean absolute error, and is thus a potentially useful tool for broadly enabling application-level power monitoring in datacenters.

2. KEY RESULT

Figure 1 shows our key result by capturing the error in disaggregating power consumption of a given job on 1100 different servers from the Google cloud trace [6] that differ in the number and characteristics of the jobs they are running. We order the servers by the Coefficient of Variation (CoV) for the disaggregated job from low (left) to high (right). We make two key observations from this experiment. First, most of the jobs (760 out of 1100 or $\sim$69%) have a very low error of 10% or less, and a very small number of jobs (86 out of 1100 or $\sim$7.81%) have a higher than 20% error. The worst-performing job has an NMAE of 90%, but less than 3W of mean absolute error (MAE). This shows that `WattScope` is highly accurate in disaggregating the power consumption of real production jobs even in the presence of a large number of jobs on the server in practical settings. The average error is 9.26%, which is small considering the variations across servers and jobs. Second, overall, the value of NMAE increases as the CoV increases indicating that disaggregation accuracy decreases as jobs' variability in their power consumption increases. However, the trend is not smooth as other factors, such as the regularity and the intensity of the power consumption for a job also affect the power disaggregation accuracy.

Key Point. *Disaggregation accuracy is high for the vast majority of jobs in production due to their low variability and high regularity.*

Acknowledgements. This research is supported by NSF grants 2213636, 2136199, 2106299, 2102963, 2105494, 2021693, 2020888, 2045641, as well as VMware

3. REFERENCES

[1] Reuters, Amazon Vows to be Carbon Neutral by 2040, buying 100,000 Electric Vans. https://www.reuters.com/article/us-amazon-environment/amazon-vows-to-be-carbon-neutral-by-2040-buying-100000-electric-vans-idUSKBN1W41ZV, September 19th 2019.

[2] N. Batra, R. Kukunuri, A. Pandey, R. Malakar, R. Kumar, O. Krystalakos, M. Zhong, P. Meira, and O. Parson. Towards Reproducible State-of-the-Art Energy Disaggregation. In *ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys)*, November 2019.

[3] M. Colmant, M. Kurpicz, L. Huertas, R. Rouvoy, and P. Felber. Process-level Power Estimation in VM-based Systems. In *ACM European Conference on Computer Systems (EuroSys)*, April 2015.

[4] D. Etherington. TechCrunch, Google Claims Net Zero Carbon Footprint over its Entire Lifetime, Aims to only use Carbon-Free Energy by 2030. https://techcrunch.com/2020/09/14/google-claims-net-zero-carbon-footprint-over-its-entire-lifetime-aims-to-only-use-carbon-free-energy-by-2030/, 2020.

[5] K. O'Sullivan. The Irish Times, Facebook Commits to Net-Zero Carbon Emissions by 2030. https://www.irishtimes.com/news/environment/facebook-commits-to-net-zero-carbon-emissions-by-2030-1.4354701, 2020.

[6] M. Tirmazi, A. Barker, N. Deng, Md E. Haque, Z. Qin, S. Hand, M. Harchol-Balter, and J. Wilkes. Borg: The Next Generation. In *ACM European Conference on Computer Systems (EuroSys)*, 2020.