O-026

Power Consumption Estimation of Applications to Enable Decarbonization Goals of Data Center Customers

Chaudhari Pritam Jaywant‡ 金子 聡‡ 岡村 拓‡
Chaudhari Pritam Jaywant Satoshi Kaneko Taku Okamura

1. Introduction

Global data center energy demand is expected to increase exponentially in the coming three decades from a projected 3000 TWh in 2030 to 500,000 TWh in 2050 [1]. To meet this energy demand sustainably, a target has been set to match data center energy demand by 75% renewable energy or hourly carbon-free energy by 2025 and 100% by 2030 [2]. To achieve these targets, we need to provide the customer with the ability to manage its application workload in tune with the availability of renewable energy and visualize the CO_2 emissions of the workload. To realize these solutions, tracking individual application workload-level power consumption accurately is crucial. Most of the existing solutions for application-level power estimation are theoretical as they do not consider the availability of metrics in real-world. Our research focuses on developing a practical solution for estimating application-level power consumption using few practically available workload metrics while achieving reasonable accuracy.

2. Conventional application-level power monitoring

An application workload is any interactive application or batch job running on servers which results in the consumption of server resources like CPU, memory, disk, network and consequently power. However, only server power can be measured externally, which consists of a constant component called idle server power and the varying workload power. The idle server power can be determined empirically by running server with no workloads. Conventional methods of individual workload power estimation used a lot of low-level metrics accessed through OS that provide hardware-level information like clock cycles, instruction counts, cache misses, etc., sensor data like CPU temperature, fan speed, etc., and server power [3],[4]. However, all these metrics are not accessible in practice. Low-level metrics accessed through OS which are especially crucial for high accuracy require administrator-level access. Also, not all servers are provided with intelligent platform management interface (IPMI) through which sensor data can be accessed. Thus, it is difficult to have accurate power models in practice depending on the availability of metrics. Our research goal is to develop a practical solution that can provide accurate application-level power monitoring for diverse applications with few practically accessible metrics.

3. Proposed method

3.1 General method overview

Fig. 1 illustrates an overview of proposed solution. The proposed solution for application power estimation uses few high-level metrics, namely, workload-related metrics like CPU usage, memory usage, disk usage together with server power data. To achieve high accuracy using few high-level metrics, workload-specific modeling approach is used. Using a workload classifier, the workload metrics are classified. For each workload class, a dedicated workload power model is developed, and a library of workload-specific power models is maintained. During the training phase, several benchmark applications are run on the server and workload-related metrics and server power data are recorded. The recorded workload metrics are used to train the workload classifier to classify the metrics into distinct workload classes. For each workload class, the classified workload metrics and corresponding server power data (minus the idle server power) are used to train regression models with workload metrics as predictor variables and server power (minus the idle server power) as predicted variable. During prediction phase, the applications’ workload metrics are monitored and fed to workload classifier which identifies the workload class. From the library of workload-specific models, pre-trained power model specific to the workload class is used to estimate each individual application’s power consumption using the monitored workload metrics.

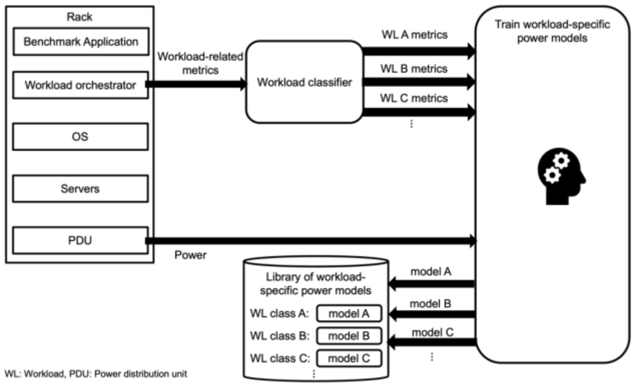

Fig. 1 Proposed method overview of application power modeling

3.2 Rationale behind workload-specific modeling

Conventional methods in literature mostly relied on many low-level, hardware performance counter (HPC) data accessed through OS kernel interface such as clock cycles, instruction counts, cache misses, etc. for developing workload power models. Thus, the conventional models are closer to the machine hardware and tend to work for any application (detached from high-level software). However, from the real-world perspective, our workload management service has access only to the workload orchestrator and so we can use only high-level workload-related data such as CPU usage, memory usage, disk usage by the application accessed through the orchestrator interface. If we use only one model for the server that was built

‡ Hitachi, Ltd.

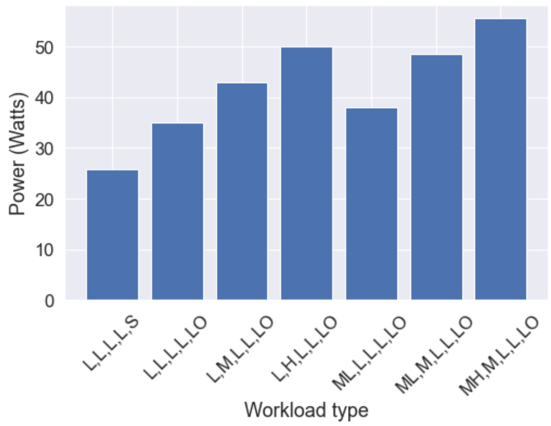

Fig. 2 Average power consumption per workload type in Alibaba data center

using such high-level data of a particular benchmark application, the model will perform poorly if the real application is different from benchmark application because our model is closer to the high-level software (detached from hardware). Thus, one tradeoff with using high-level metrics would be the need for several benchmark applications during training phase.

4. Evaluation

4.1 Workload classification

The objective of the evaluation was to estimate application-level power using few high-level metrics with mean absolute percentage error (MAPE) of less than 10% as claimed by most conventional methods. Towards this objective, we carried out two studies, workload classification and workload-specific model-based application power estimation. For workload classification, we used Alibaba cluster trace dataset which contains information about workload's CPU usage, memory usage, disk usage, network usage and duration [5]. The workloads from Alibaba cluster trace are classified into distinct types based on levels of CPU usage, memory usage, disk usage, network usage and workload duration (denoted in this order in Fig. 2) [6]. The levels are low (L), medium low (ML), medium (M), medium high (MH), high (H) and short (S), long (L) for duration. Fig. 2 shows the average power consumption of each workload type. It can be inferred that long-running resource-intensive jobs have large average power consumption. However, jobs with low resource-usage and low power consumption like L,M,L,L,LO account for about 47% of total number of workloads, while resource-intensive jobs like MH,M,L,L,LO account for only 4%. Thus, most of the workloads hosted by a typical cloud data center like Alibaba Cloud, are low resource usage, low power consuming workloads, which need to be targeted from energy-aware workload management perspective. This study bolsters the argument that identifying workload type is crucial in workload power consumption monitoring.

4.2 Application power estimation

We carried out an experiment with Cloudsuite benchmark applications in the form of docker containers running on a server [7]. The workload metrics like CPU usage, memory usage, disk

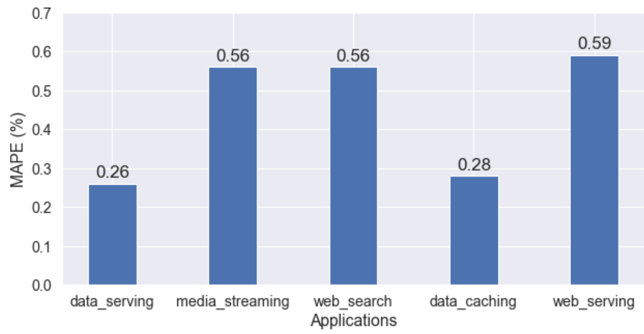

Fig. 3 Server power consumption estimation error

usage, network usage together with server power data was collected. A single benchmark application was run at a time on the server and the workload metrics and server power was monitored. A regression model was trained specific to each application with the corresponding workload metrics as predictor variables and the server power as predicted variable. We could estimate server power with the MAPE of less than 1% for most applications using application-specific models as shown in Fig. 3. By subtracting the empirically determined idle server power, we could get application power estimations.

5. Conclusion

To achieve DC customer's decarbonization goals through renewable-aware workload management and feedback workload's CO_2 emission performance, accurate monitoring of power consumption by workload is crucial. Existing solutions for workload-level power consumption monitoring used mainly low hardware-level information which is not accessible in practice. We proposed the workload-specific power modeling approach that uses few high-level workload metrics like CPU usage, memory usage, disk usage and can still achieve accuracy comparable to conventional methods. We plan to validate our method using workload data from a real data center.

References

[1] "Impact of Progress of Information Society on Energy Consumption (Vol. 3): Current Status and Future Forecast of Network-Related Energy Consumption and Technical Issues," Center for Low Carbon Society Strategy, Japan Science and Technology Agency, 2021.

[2] "Climate Neutral Data Centre Pact," [Online]. Available: https://www.climateneutraldatacentre.net/. [Accessed 17 June 2022].

[3] M. Colmant, P. Felber, R. Rouvoy and L. Seinturier, "Wattskit: Software-defined power monitoring of distributed systems," in 2017 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID), Madrid, 2017.

[4] M. Canuto, R. Bosch, M. Macias and J. Guitart, "A methodology for full-system power modeling in heterogeneous data centers," in Proceedings of the 9th International Conference on Utility and Cloud Computing, 2016.

[5] "Alibaba Cluster Trace Program," Alibaba Group, [Online]. Available: https://github.com/alibaba/clusterdata. [Accessed 17 June 2022].

[6] V. Shekhawat, A. Gautam and A. Thakrar, "Datacenter workload classification and characterization: An empirical approach," in 2018 IEEE 13th International Conference on Industrial and Information Systems (ICIIS), Rupnagar, 2018.

[7] "Cloudsuite," EPFL PARSA - 1015 Lausanne, Switzerland, [Online]. Available: https://www.cloudsuite.ch/. [Accessed 17 June 2022].