ACE: Abstracting, Characterizing and Exploiting Peaks and Valleys in Datacenter Power Consumption

Di Wang[*], Chuangang Ren[*], Sriram Govindan[†], Anand Sivasubramaniam[*], Bhuvan Urgaonkar[*], Aman Kansal[‡], Kushagra Vaid[†]
[*]The Pennsylvania State University, [†]Microsoft, [‡]Microsoft Research
{diw5108, cyr5126, anand, bhuvan}@cse.psu.edu, {srgovin, kansal, kvaid}@microsoft.com

ABSTRACT

Peak power management of datacenters has tremendous cost implications. While numerous mechanisms have been proposed to cap power consumption, real datacenter power consumption data is scarce. To address this gap, we collect power demands at multiple spatial and fine-grained temporal resolutions from the load of geo-distributed datacenters of Microsoft over 6 months. We conduct aggregate analysis of this data, to study its statistical properties. With workload characterization a key ingredient for systems design and evaluation, we note the importance of better abstractions for capturing power demands, in the form of peaks and valleys. We identify and characterize attributes for peaks and valleys, and important correlations across these attributes that can influence the choice and effectiveness of different power capping techniques. With the wide scope of exploitability of such characteristics for power provisioning and optimizations, we illustrate its benefits with two specific case studies.

Categories and Subject Descriptors

C.0 [**Computer Systems Organization**]: General

Keywords

Datacenter, Power Characterization, Peaks and Valleys

1. INTRODUCTION

The cost, scalability and environmental concerns arising from the power consumption of datacenters has come under extensive scrutiny. While much of the prior work in the area has looked to reduce energy of computing and cooling systems, the importance of how this energy is dissipated over time (i.e. the power) has gained a lot of recent attention. Power dissipation, particularly the peak or high power draws, impact both operational (op-ex) and capital (cap-ex) expenditures. Electric utilities can charge differentially (op-ex) for peaks, especially if such high power draws coincide with high demand across the grid because of supply-demand mismatches that can lead to potential black or brown-outs. Peak power draws also determine the capacity of the power distribution and cooling infrastructure that is provisioned within the datacenter. Prior studies [1, 2] have pointed out that provisioning costs can range between $10-20 per watt, which is incurred even if that watt is not actually consumed.

SIGMETRICS'13, June 17-21, 2013, Pittsburgh, PA, USA.
ACM 978-1-4503-1900-3/13/06.

To address this problem, numerous prior optimizations have been proposed. However, there is a lack of real world datacenter power consumption data to guide the design and enable thorough evaluation of these optimizations. Detailed power consumption data at different temporal scales (from seconds to months) and spatial granularities (from servers, chassis, racks, to datacenters) for datacenters serving important workloads is not easily available. This paper intends to fill this critical void by providing an in-depth analysis of measured power characteristics from the datacenter infrastructure of Microsoft corporation. Below, we present a summary of the results, with more details in our technical report [4].

2. AGGREGATE CHARACTERISTICS

In this section, we present a spatio-temporal analysis of power demand, focusing on its aggregate characteristics.

2.1 Tracing and Data Collection

We collected power measurement data from multiple geo-distributed datacenters run by Microsoft over a six month period, between July-December 2011. We specially give results here for data pertaining to 8 representative server clusters (see Table 1) in the interest of clarity. Each such cluster comprises several hundreds of servers that span multiple chassis and racks. These clusters run a variety of workloads including web-search, email, Map-Reduce jobs, and several other online cloud applications, catering to millions of users around the world. We name the 8 clusters as C_1, C_2, ..., C_8 and present the trace collection resolution and the type (soft-realtime, batch and interactive) of application that each hosts in Table 1. Apart from the IT power of the clusters, we have also collected cooling power.

Cluster Names	Data Resolution	Application Type
C_1, C_2	20 **seconds**	**Soft-realtime, Batch**
C_3, C_4, ..., C_8	120 **seconds**	**Interactive**

Table 1: Data collection is done for a period of six months.

2.2 Statistical Properties

Spatial Characteristics: Figure 1 shows the CDF of power consumption at multiple (spatial) levels for one of the clusters, C_1, starting from cluster-level to rack-level to chassis-level to server-level. The x-axis is normalized with respect to the sum of the maximum demand (over time) of all the servers within that cluster. At the lower levels (e.g. server), there is higher variance in power demands. However, statistical multiplexing effects of these demands, tend to smoothen the fluctuations as we go higher up the hierarchy. Thus, there is fairly good statistical multiplexing, and the likelihood of simultaneous peaks across all equipment at the same time reduces as we move higher up in the power hierarchy. These results corroborate observations made in [1], showing the rarity of peak needs at the aggregate level, further motivating the attractiveness for under-provisioning the power hierarchy (especially as we go higher up the hierarchy).

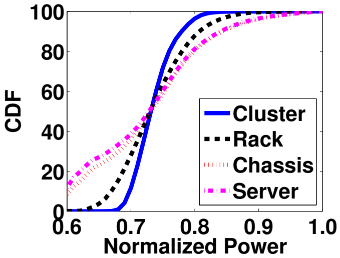

Figure 1: CDF of Power Demand at 4 different levels of C_1

Cluster Name	Hurst Parameter
C_1	**0.93**
C_2	**0.91**
C_3	**0.89**
C_4	**0.90**
C_5	**0.90**
C_6	**0.82**
C_7	**0.87**
C_8	**0.86**

Figure 2: Hurst parameter values of 8 clusters.

Temporal Characteristics: One way to understand the temporal time series of power demands is through an Autocorrelation Function (ACF) plot with different time-lags. We note that there is a fairly good time-of-day behavior that is exhibited by the power demands - lags of 24 hours (and multiples) have high correlations, and lags of 12 hours (and its odd number of multiples) are the least correlated. Furthermore, the slower than exponential decay of the ACF indicates that the demands do not follow a Poisson process, with possibility of self-similarity over time. To investigate the presence of self-similarity, we calculate the Hurst parameter, using several techniques [3], including variance, R/S method, and periodogram plots. The results are consistent across these techniques. We find a high value for the Hurst parameter, over 0.8, for all clusters (Figure 2). These quantitatively show the existence of self-similarity in the power demands.

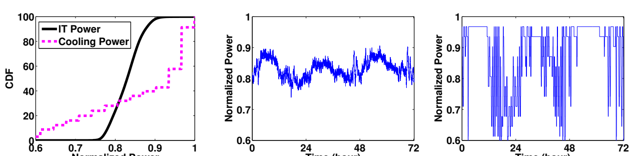

(a) CDF of IT and Cooling Power **(b) IT power** **(c) Cooling power**

Figure 3: Normalized IT and cooling power of C_1.

IT and Cooling Power: Figure 3 (a) compares the CDF of IT equipment (servers and networking devices) and cooling power consumptions both normalized with respect to their individual maximum demands for C_1. There are several interesting observations from these results: (i) Cooling power and IT power are correlated as can be seen from Figure 3(b) and (c). The *Pearson* correlation coefficient between the two is 0.841. However, the cooling power, as is to be expected due to thermal time constants, lags 2 minutes behind the IT power to reach the maximum correlation coefficient of 0.844. As expected based on the high correlation with IT power, the cooling power also exhibits time-of-day behavior and self-similarity. The Hurst parameter value is 0.90. (ii) The variation in cooling power is much more pronounced than that in IT power (also seen visually in Figures 3 (b) and (c)). Beyond its dependence on the IT power draw, cooling power also depends on other parameters including external temperatures, air-flow, etc. High user demand and consequently a high IT power consumption often occurs at times of the day when external temperatures are also among the highest for the day, leading to this wider fluctuation. (iii) The CDF shows that the cooling system is operating closer to its maximum actual draw, much more often than the IT systems. This is expected because cooling systems have fewer power states, resulting in more discrete modes of operation.

3. PEAKS AND VALLEYS OF POWER DEMANDS

Abstractions: We formally define peaks and valleys, their important attributes (height, width, area), and the correlations between peaks and valleys that need to be studied towards designing and understanding the potential of any peak suppression mechanism.
Characterizing Peaks and Valleys: We extensively characterize peak and valley attributes individually, and their cross-correlations. Results show that while there are an overwhelming number of small duration and small amplitude peaks, we cannot afford to ignore the few large ones that have very stringent demands. Further, while on the average, valleys do offer enough slack for load deferment or peak preparation, there are bursts of peaks which do not have sufficient valleys immediately following or preceding them. Further, there is significant potential of migrating load to exploit peaks and valleys across clusters, as long as we can restrict the number of such migrations to avoid the consequent performance penalties. These suggest aggregated optimizations of peaks and valleys.
Exploiting Characteristics: There are numerous use-cases for our characteristics, and we illustrate two specific case-studies. The first uses the characteristics to fine-tune load deferring and migration based on the kinds of peaks, in an aggregated manner. The second shows a simple approach to energy storage provisioning that only uses aggregate characteristics, rather than an extensive approach considering every possible eventuality in the entire power demand time series.

4. ACKNOWLEDGMENTS

This work was supported by NSF grants 0811670, 1152479, 1205618, 1213052, 1147388 and CAREER award 0953541.

5. REFERENCES

[1] X. Fan, W.-D. Weber, and L. A. Barroso. Power Provisioning for a Warehouse-sized Computer. In *Proceedings of ISCA*, 2007.
[2] J. Hamilton. Internet-scale Service Infrastructure Efficiency, ISCA Keynote 2009.
[3] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Trans. Netw.*, 2(1), 1994.
[4] D. Wang, C. Ren, S. Govindan, A. Sivasubramaniam, B. Urgaonkar, A. Kansal, and K. Vaid. ACE: Abstracting, Characterizing and Exploiting Peaks and Valleys in Datacenter Power Consumption. Technical Report CSE13-003, The Pennsylvania State University, 2013.