Proceedings of the ASME International Mechanical Engineering Congress & Exposition
IMECE 2010
November 12 - 18, 2010, Vancouver, British Columbia, Canada

IMECE2010-39503

EXPLORATORY ANALYSIS OF AGGREGATE POWER METRICS IN DATA CENTERS

Manish Marwah
Hewlett-Packard Laboratories
Palo Alto, CA

Martin Arlitt
Hewlett-Packard Laboratories
Calgary, AB

Christopher Hoover
Hewlett-Packard Laboratories
Palo Alto, CA

Cullen Bash
Hewlett-Packard Laboratories
Palo Alto, CA

Ratnesh Sharma
Hewlett-Packard Laboratories
Palo Alto, CA

ABSTRACT

In recent years, climate change, depletion of conventional energy sources and rising energy costs have led to an increased focus on sustainability. Within the Information Technology (IT) sector, data centers are significant energy consumers. The first steps towards reducing power consumption in data centers are to monitor it and to determine the heavy hitters.

Unfortunately, fine-grain power information is often not readily available within data center environments. In this paper, we conduct an exploratory analysis of aggregate power data in a data center. We collect data from the power infrastructure of a data center in Palo Alto, CA, as well as from a data center in Bangalore, India. We examine the data in increasing detail, and reveal the opportunities and challenges for disaggregating data center power consumption data.

INTRODUCTION

In recent years, the demand for data centers has seen tremendous growth. Many of the largest data centers in the US are experiencing a growth at 20% per year and over 40% of enterprises are refurbishing or building new data centers to support ongoing business operations and anticipated future demand (Greiner2008). Energy consumption of data centers is a growing concern. The Environmental Protection Agency (EPA) calculated that in 2006, 61 billion kilowatt-hour (kWh) of electricity costing $4.5 billion was consumed by data centers in the US. This amount accounts for 1.5% of the total US electricity consumption (USEPA2007). Of this, the cooling infrastructure could be responsible for up to 50% (Belady2007). It is estimated that data center power consumption will increase 4% to 8% annually and is expected to reach 100 billion kWh by 2011 (ClimateGroup2008).

This paper makes two primary contributions. First, it characterizes power use in two actual data centers. Second, it explores the opportunities and challenges in disaggregating power usage data in a data center environment. These are important initial steps towards improving the energy efficiency of data centers.

The remainder of the paper is organized as follows. The next section provides background information on general trends in electricity consumption, and a brief introduction to data centers. The Related Work section focuses on previous efforts to disaggregate power usage in residential and commercial settings. The paper then describes the data sets and provides characterization results. Next, the “Nonintrusive Application Load Monitoring” technique is described, followed by the results of its application to power use in two data centers. A discussion of the opportunities and challenges for disaggregating data center power use and a summary of our work concludes the paper.

BACKGROUND

This section provides background information on two topics. First, general trends (i.e., data center-agnostic) in electricity demand are examined. Second, a brief introduction to data centers is provided.

Electrical Usage Patterns

Figure 1 shows the estimated electricity demand in California for the week February 7-13, 2010. The demand can conceptually be broken into two parts: the *base load* and the *variable load*. The base load represents the minimum demand that must be supplied at all times; for this week, the base load was about 16.8 GW. The remainder of the demand is considered the variable load. Figure 1 shows several common patterns. First, there is a distinct time of day pattern, with demand lowest during the early morning hours, a sustained surge in demand during typical work hours (e.g., 9am-5pm), and the peak demand (up to 30.5 GW during the observed week) in the evening. This daily behavior is clearly influenced by human activities, such as sleeping, working, and food preparation. A second pattern occurs across the days of the week. In particular, weekdays have higher average and peak demands than do weekend days. Weather can also affect the demand for electricity (Dryar1944).

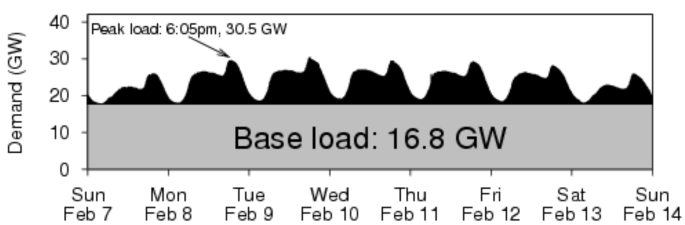

Figure 1. California Electrical Demand, Feb. 7-13, 2010. (Data from http://oasis.caiso.com/mrtu-oasis/)

Data Centers

Data centers are controlled environments for operating IT equipment such as servers, storage, and networking components. Data centers have three primary types of infrastructure: the IT infrastructure, the cooling infrastructure, and the power infrastructure. A conceptual view of these infrastructures is shown in Figure 2. The IT equipment is typically installed in industry standard racks, which are then configured in rows. The rows are aligned to form either "hot aisles" or "cold aisles". The cold aisles include perforated floor tiles which allow cool air from the floor plenum to flow to the inlets on the IT equipment. The figure shows a photo of an exemplary cold aisle. The cool air then flows through the IT equipment, and removes the heat from the operating electronic components. This air is then exhausted into a hot aisle. The warm air circulates (e.g., via a ceiling plenum) to Computer Room Air Conditioning (CRAC) units located in the data center. The cooling infrastructure typically consists of one or more CRACs, the chillers which generate chilled water for use by the CRACs, and a cooling tower which exhausts heat from the water returned from the CRACs. The power infrastructure connects to either the utility grid or to an onsite source of electricity generation. From here the power feeds into switching gear, which directs the power to either the IT infrastructure or the cooling infrastructure. The power infrastructure typically includes an Uninterruptible Power Supply (UPS) to enable the IT equipment to continue operating during any short term disruptions to the grid or on-site power generation or transmission. Lastly, Power Distribution Units (PDUs) are used to distribute the power to the individual IT components. Note that power to the cooling infrastructure does not need to go through a UPS and PDU.

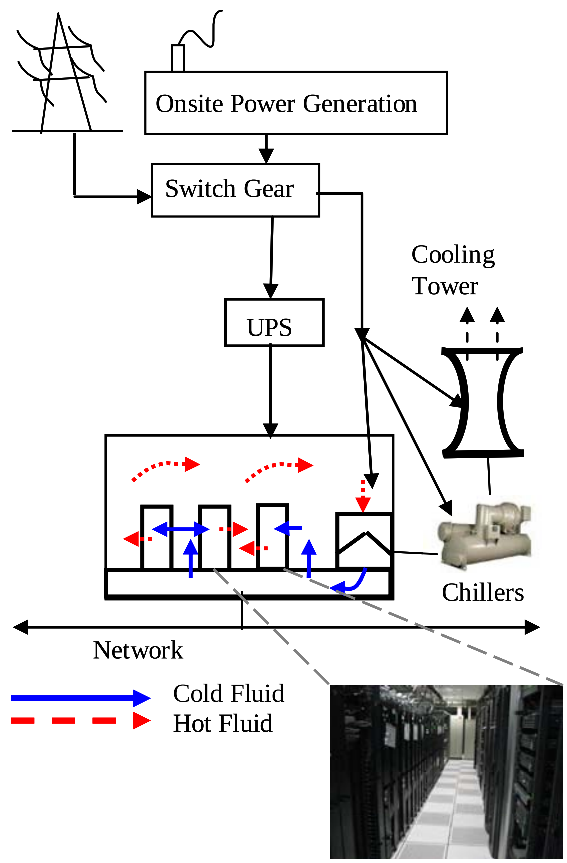

Figure 2. Data center power, cooling and IT infrastructures.

RELATED WORK

Prior work on disaggregating power consumption data has focused primarily on residential use. The seminal work on disaggregating electricity consumption data was conducted by George Hart two decades ago (Hart1992). Hart introduced the "Nonintrusive Application Load Monitoring" (NALM) technique to simplify the collection of device-specific energy consumption data in residential settings. NALM analyzes current and voltage waveforms of the aggregate residential electrical load for signatures of different devices. The information can be used for numerous purposes, such as assisting with energy conservation efforts or identifying failing devices.

While NALM has numerous beneficial qualities, it also has limitations. Three main limitations are: it does not work well for identifying "small" appliances (i.e., those that consume < 100W), for devices that use a continuously varying amount of electricity, or for distinguishing between similar devices. A number of subsequent researchers have investigated solutions to these limitations. For example, Laughman *et al.* looked at using higher harmonics to distinguish between loads with overlapping signatures (Laughman2003). Farinaccio and Zmeureanu use a rule-based pattern recognition technique to disaggregate the total electrical load in a home into the major end uses (Farinaccio1999). Prudenzi recommends a neural net approach to identify the domestic applications in a residential electrical load (Prudenzi2002). Matthews *et al.* provide a more detailed explanation of these and other recent works on disaggregating residential electrical loads.

In a more recent study, Patel *et al.* examined how to detect a variety of electrical events throughout a home using a single plug-in device (Patel2007). They monitor the broadband electrical noise (both transient and continuous) that is generated by abruptly switched mechanical or solid-state devices. This is then used to construct "features" or inputs that can be used to train a classifier such as a support vector machine. The authors claim that the technique may be able to distinguish events among a dense collection of devices that have similar switching and load characteristics. However, they have not yet demonstrated this.

Norford and Leeb were the first to apply NALM (which they renamed "NILM", or "Non Intrusive Load Monitoring") to a commercial setting (Norford1996). They indicate that disaggregating a commercial electrical load is more difficult than for a residential load, as many large electrical devices in commercial settings have more complex consumption patterns. Their work considers how to address this challenge for space-conditioning equipment of an office building.

DATA COLLECTION

We examine power consumption information from two data centers. Our initial analyses use data from the HP Labs data center in Palo Alto, CA. For this data center, we have two types of data: (1) power consumption of the IT equipment in the data center, recorded at four PDUs; and (2) power consumption of the chiller that provides chilled water to the CRAC units in the data center as well as to three buildings on the HP Labs campus. The Palo Alto data we use is for the period October 1st, 2009 until May 31st, 2010. Three of the PDUs recorded consumption every 20 seconds, while the fourth PDU and the chiller recorded consumption every 30 seconds.

We also have power consumption data of chillers units and pumps from an HP data center in Bangalore, India. The cooling infrastructure there contains three air-cooled and two water-cooled chillers. However, the chillers are not individually instrumented with power meters. Instead, power meters are installed at subpanels in the data center serving one or more chillers and pumps. Here, we look at data from one such metering point that measures aggregate power consumption of an air-cooled chiller and three pumps. Data is available at about one minute intervals for about 14 days from October 2009.

EXPLORATORY ANALYSIS

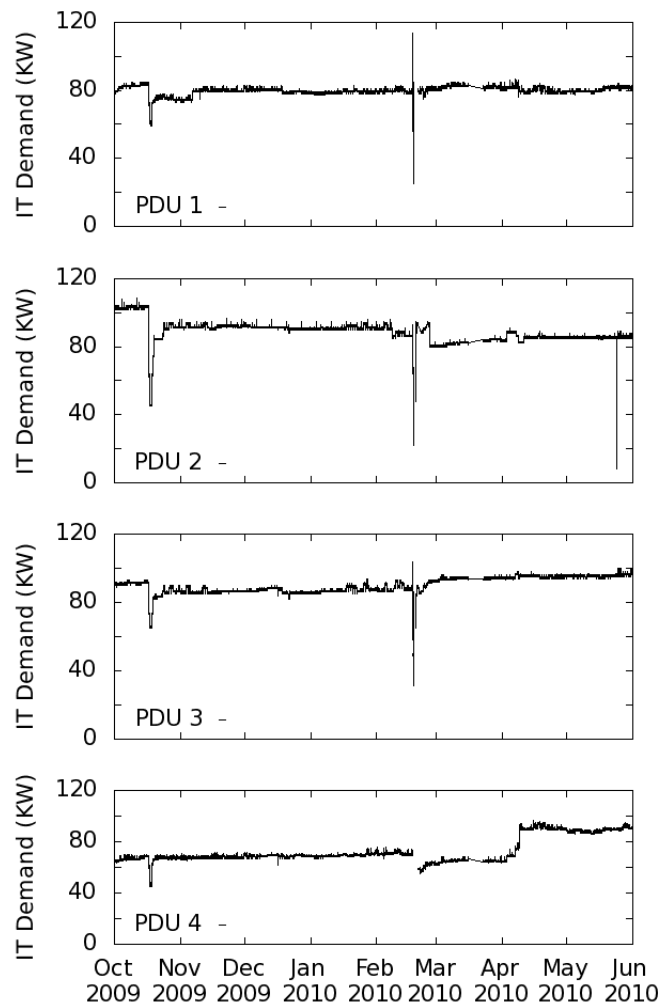

Figure 3. Power consumption of IT equipment in Palo Alto data center, Oct. 2009 - May 2010.

Each piece of IT equipment in the Palo Alto data center is connected to one of four Power Distribution Units (PDUs). Each PDU records the aggregate power consumption of the devices connected to it. The initial question we consider is what macro-level information we can determine from these four aggregated measures.

Figure 3 shows the power consumption of the IT equipment in Palo Alto. An important observation is that the consumption is quite constant, with the exception of a few key events. For example, in mid-October 2009, a scheduled shutdown of some

IT equipment occurred, to facilitate maintenance work on the cooling infrastructure. In mid-February 2010, an unscheduled shutdown occurred when an hours-long, city-wide power outage occurred. Lastly, in mid-April 2010 the demand on PDU 4 increased noticeably, as new IT equipment was added to an area of the data center supported by this particular PDU.

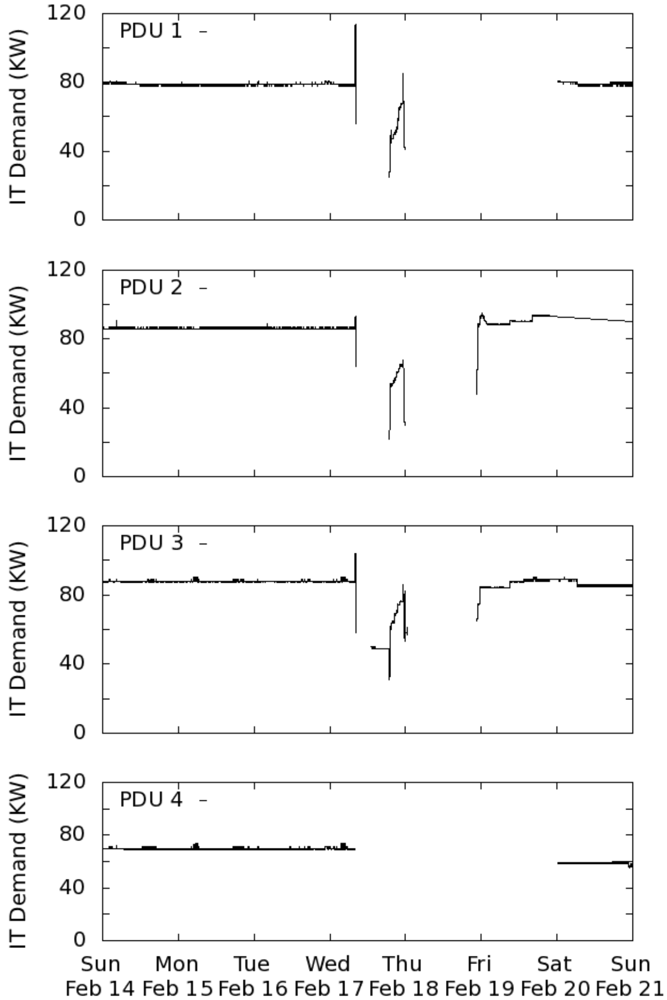

Figure 4. Power consumption for week of Feb. 14-21, 2010.

Figure 4 provides a closer look at the power consumption of the data center during the week of February 14, 2010. For the first three days of the week, the demand was relatively constant on each PDU. The demand is roughly equal across PDUs 1, 2 and 3, and slightly lower on PDU 4. This occurs because IT equipment is statically assigned to a specific PDU. On Wednesday, February 17th, 2010, a plane crash in Palo Alto created a city-wide power outage (Perry2010). The outage started around 8am; once the backup power in the data center's UPSs was exhausted, the IT equipment abruptly shut off (in some cases preceded by a momentary surge in demand). A few hours later, power was restored in Palo Alto, and the data center equipment started to come back online. However, an equipment failure in the data center's power infrastructure took the data center offline again. The data center was partially operational within a day of this failure, and fully operational within two days.

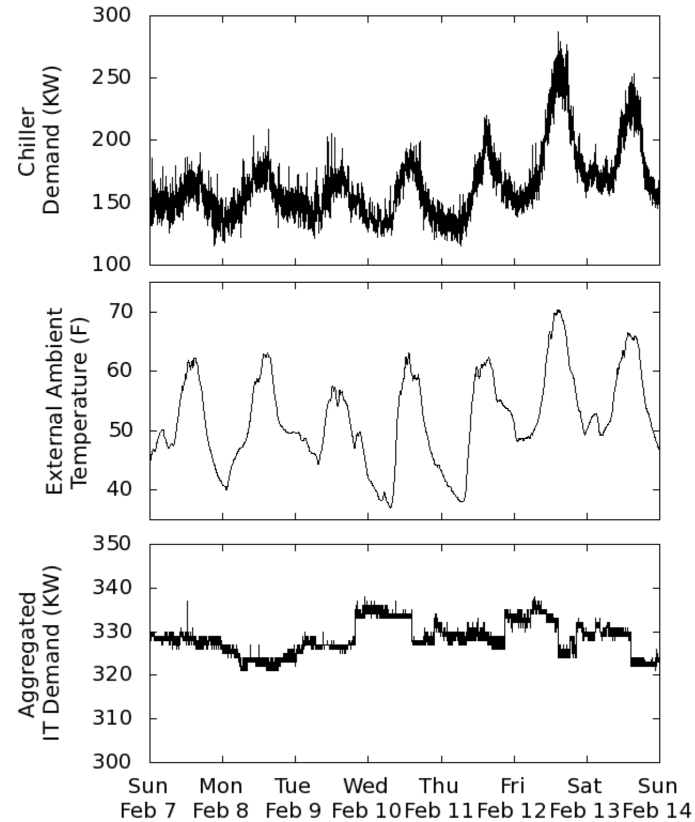

Figure 5. Examining the chiller power consumption.

For economic efficiency, the HP Labs campus uses a chiller to provide chilled water to both the data center and several buildings. Figure 5 shows the power consumption of the chiller for the week of Feb. 7th-13th. The top graph in Figure 5 reveals the chiller's power consumption ("demand") is quite variable, with the peak demand more than double the minimum demand (base load). This is quite different from the IT equipment's consumption (bottom graph), which only varied about 20 KW from a base load of 320 KW. As the IT demand is quite constant, we expect that it primarily contributes to the base load of the chiller. The variable load of the chiller is largely influenced by the external ambient temperature, as can be seen by comparing the top and middle graphs of Figure 5.[1]

Although the aggregated IT demand is relatively stable, the bottom graph of Figure 5 does reveal three noticeable (and similar) events occurred that week. For example, during the night of Tuesday, February 9th, there was a sudden increase of approximately 10 KW in the demand. This increase was sustained until the following evening. Similar events can be

[1] This graph uses weather data from http://www.wunderground.com/, for weather station KCAPALOA9 – Barron Park, Palo Alto, CA.

seen from Thursday night through Friday evening, and from Friday night until Saturday evening. In the next section we investigate such events systematically rather than visually.

PROBLEM DESCRIPTION AND TECHNIQUES USED

While our visual exploration of the data can provide information on macro-level events, to understand the micro-level events a more in-depth examination of the data is necessary. In particular, we consider the "Nonintrustive Application Load Monitoring" or NALM technique (Hart1992). This technique has been shown to work reasonably well in both residential and commercial environments. A desirable feature of NALM is that it can extract component-level behaviors from aggregated power data. Like any technique though, it has its limitations. We discuss these more in the Challenges section.

NALM has five basic steps (Hart1992):

- Power measurement
- Detection of on-off events
- Clustering of similar events
- Matching of on/off events over time
- Equipment identification

For this paper, we are using power measurements from the sources described earlier. Simple "ON/OFF" events can be identified by calculating the change in consumption from one measurement interval to the next. Increases in consumption indicate that one or more components were turned on during the past measurement interval, while decreases in consumption indicate that one or more devices were turned off. However, in many instances an increase and decrease occurs over several data points. To address these cases, we look for a "run length" of increases (or decreases) and then sum them up to extract rising (or falling) edges.

All the edges thus detected are then partitioned into groups using clustering [13], which is a generic technique used to partition a set of elements into groups (clusters) such that elements within a cluster are more similar than those across clusters. We used a simple algorithm for clustering called *k-means* which divides a data set into *k* clusters. The algorithm requires the number of clusters, *k*, and the data as input, and aims to minimize the variance within each cluster. It assumes an initial set of "centers" for the *k* clusters and then alternatively performs the following two steps.

1. For each center *i*, it identifies the set of points that are closest to it as compared to other centers. All of these points belong to cluster *i*. Typically, Euclidean distance is used to determine the closeness of two data points.
2. It computes the new center of each cluster by averaging all the points belonging to that cluster.

These two steps are repeated until convergence is reached; that is, when the centers of the clusters no longer change. Typically, the initial centers are randomly chosen from among the data points. Multiple restarts of the algorithm are performed and the best result, that is, the one with minimum total intra-cluster variance, is chosen.

Since the edges detected are assumed to correspond to ON/OFF events and noise, an odd number of clusters (k) is chosen with one cluster corresponding to noise and an equal number corresponding to ON and OFF events. The result of clustering is visually inspected to determine the noise cluster (which typically has the most number of points and a center very close to zero), and matching ON/OFF events. With some domain knowledge, it may be possible to associate these events with distinct pieces of equipment and then their power consumption can be monitored over time, without the need to install a meter at that component. However, there are several situations where it is challenging to uniquely disaggregate the data as discussed in a later section.

EXPERIMENTAL RESULTS

To demonstrate the effectiveness of the disaggregation techniques described in the previous section, we apply it to two sets of data: 1) Cooling infrastructure data collected from a subpanel at the Bangalore data center; and, 2) IT power data collected from a PDU of the Palo Alto data center.

Bangalore Data

The top graph in Figure 6 shows the raw power consumption data for 14 days. As described earlier, this data corresponds to the aggregate power consumption of an air-cooled chiller and three pumps. Figure 6 reveals that there is much more variation here than in the IT power consumption data. The abrupt changes of several hundred KW are due to the air-cooled chiller being turned on or off. The more gradual changes in demand over the course of a day are likely due to the external ambient temperature in Bangalore, shown in the bottom graph of Figure 6. A more detailed view of the time series is shown in Figure 7, where each plot shows one entire day beginning at midnight. Both large and small variations in power consumption can be seen. To disaggregate the data, we apply the aforementioned techniques.

In all, there are 19,450 data points in the 14 day time series. The rising and falling edges are detected and k-means is applied to cluster the edges. A plot of the edges is shown in Figure 8 with the positive values indicating rising edges, and negative falling edges. In addition to the noise cluster (described later), two types of clusters are discovered: ON clusters that contain positive points and correspond to rising edges; and OFF clusters that contain negative points and correspond to falling edges. A number of values for k (number of clusters) were tried during clustering: 3, 5, 7, 9 and 11. In each of these, the cluster corresponding to noise stood out and was easily identified. Its mean was always approximately zero and had an order of magnitude more points (between 16,000 to 18,000 points for the above k values) than other clusters. To determine the best k, we matched the ON and OFF clusters in each case. A close match (both in terms of cluster center magnitudes and their sizes) indicates that those clusters likely correspond to ON/OFF events of the same devices (or set of devices). In our case, k=5 provided a good match as can be seen in Table 1, which describes the properties of all the five clusters. The table lists the number of points in each of the clusters, the cluster centers, standard deviation, and the minimum and maximum elements in the cluster. These clusters are also marked by different symbols in Figure 8. More information regarding the distribution of points within each cluster is provided in Figure 9 through *box and whisker plots* of each cluster. A box and whisker plot is a convenient way to summarize sample data. The bottom and top of the box are the 25^{th} and 75^{th} percentiles, respectively, while the band near the middle of the box is the median. The minimum and maximum values are shown at the ends of the "whiskers". Points considered outliers are plotted separately.

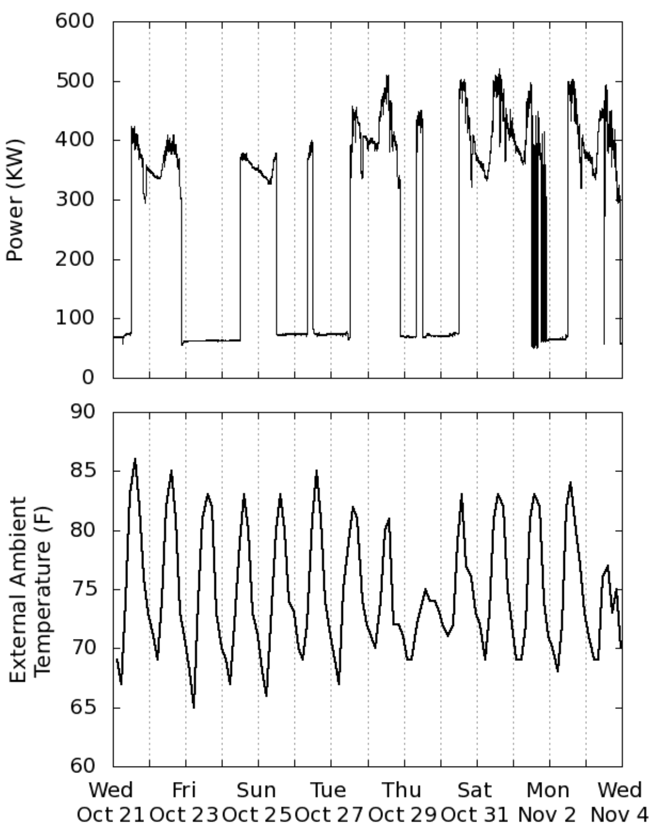

Figure 6. Power time series for cooling infrastructure subpanel of Bangalore data center.

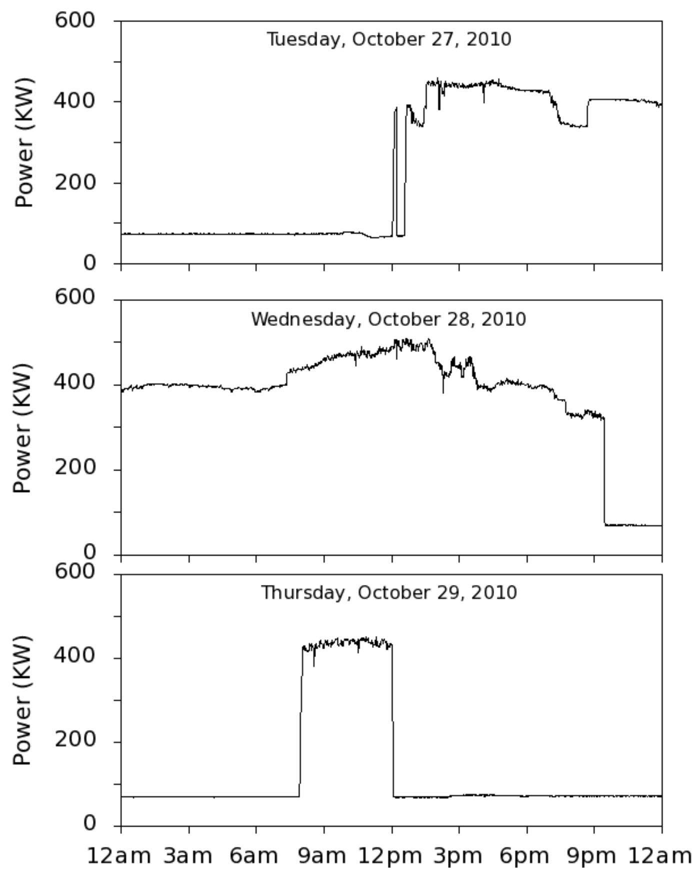

Figure 7. Detailed view of the power time series in Figure 6 over a three day period.

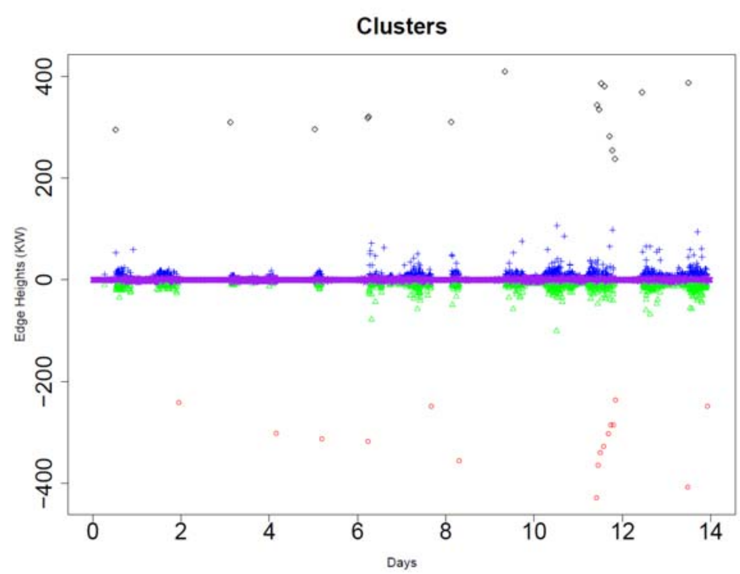

Figure 8. Edges detected from the power time series. The symbol (or color) of the point indicates the cluster it belongs to.

Table 1. Details of the five clusters.

Cluster ID	Number of points	Mean (KW)	SD (KW)	Min (KW)	Max (KW)
1	16	-312.9	56.8	-428.0	-236.8
2	744	-14.7	9.1	-100.4	-8.2
3	717	15.0	11.7	8.2	106.5
4	17,957	-0.004	1.6	-7.2	6.1
5	16	327.2	49.6	237.6	409.6

Looking at Figure 9, cluster 4 is the noise cluster, as it is tightly clustered around 0. Clusters 1 and 2 are OFF clusters; while 3 and 5 are ON clusters. OFF cluster 1 and ON cluster 5 match particularly well – the number of points are identical and the magnitude of the means differ by less than 5%. Given the large magnitude of the edges in these clusters, these likely correspond to the turning ON/OFF of the air cooled chiller. The variation within these clusters exists since the chiller is not a single state device. That is, it can consume different amounts of power depending on its operational mode. In the case of a chiller, different utilization levels correspond to different levels of power consumption. Furthermore, in some cases one or more pumps can also be turning on/off about the same time as the chiller. In this respect, we are limited by the one minute granularity of the samples.

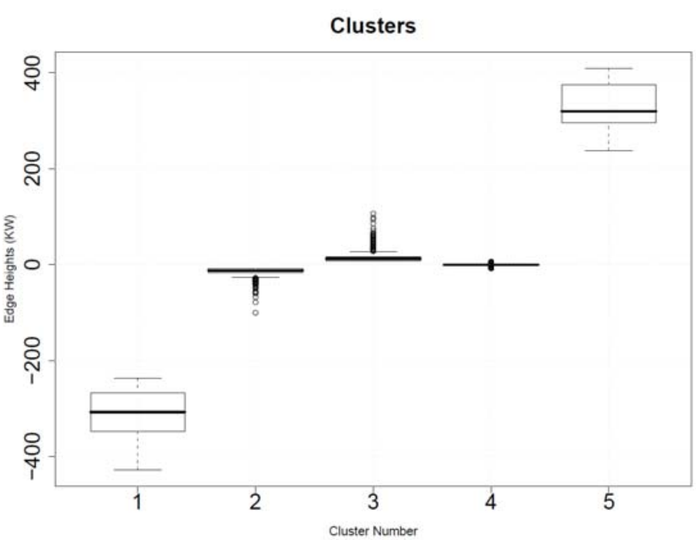

Figure 9. Box and whisker plots of the five clusters.

The remaining two clusters, ON cluster 3 and OFF cluster 2, also match closely. While the number of points is slightly different, the magnitude of the means is almost identical. However, relating these clusters to devices is more challenging in this case. These clusters correspond to one or more pumps turning ON/OFF. Note that the variance within these two clusters is higher (as a percentage of the mean) compared to the chiller clusters. A higher sampling rate would simplify the disaggregation problem by separating the pump ON/OFF events. Further, additional domain information – such as the type and power rating of these pumps, their operating schedule, how correlated is their operation with each other or the chiller, etc. – would help in disaggregating the pump power usage.

Palo Alto Data

When the above analysis is repeated on PDU 1 data from the Palo Alto data center (see Figure 3), it does not yield any useful results. Due to the static nature of the electrical load, there are no edges detected that correspond to ON/OFF events (other than a few related to the failure event). Figure 10 shows the data for PDU 1 for the month of February 2010. The large variation corresponds to the failure event described earlier, but other than that the power usage is relatively constant. The edges obtained are mainly noise events and subsequent clustering only results in forming noise clusters. Tables 2 and 3 provide details of the application of k-means with k set to 3 and 5. In Table 2, cluster 1 consists of small negative values and zeros, while cluster 2 contains small positive values. Cluster 3 can be ignored since it contains only one data point, which corresponds to the failure event. With five clusters, two (clusters 1 and 2) are spurious containing very few points, while clusters 3, 4 and 5 contain negative noise, positive noise and zeros, respectively.

Table 2 Clustering of PDU data with k = 3

Cluster ID	Number of points	Mean (KW)	SD (KW)	Min (KW)	Max (KW)
1	104,853	-0.11	0.36	-41	0
2	11,584	1.04	0.89	1	81
3	1	-105	NA	-105	-105

Table 3 Clustering of PDU data with k = 5

Cluster ID	Number of points	Mean (KW)	SD (KW)	Min (KW)	Max (KW)
1	2	58	32.5	35	81
2	1	-105	NA	-105	-105
3	11,574	-1.03	0.47	-41	-1
4	11,582	1.03	0.37	1	26
5	93,279	0	0	0	0

CHALLENGES AND OPPORTUNITIES

Our preliminary exploration indicates that power consumption data in data center environments contains a wealth of interesting and potentially useful information. However, there are challenges in extracting such information. While the general issues are known limitations of NALM (Hart1992), we re-state them here as they apply in the data center content.

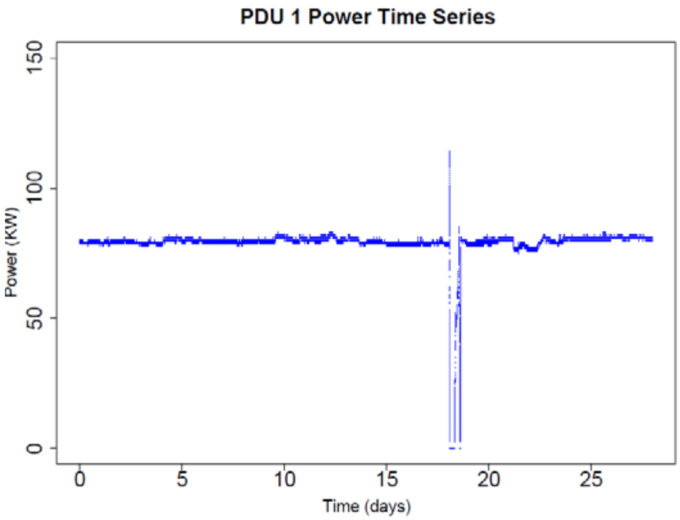

Figure 10. PDU 1 data for the month of February 2010 from the Palo Alto data center.

Number of devices. Data centers have many devices that consume power. For example, there could be thousands or tens of thousands of pieces of IT equipment in a data center. Even if sub-metering is used (e.g., at the PDU level rather than the data center level), there could still be dozens or hundreds of unique devices to identify and monitor. One practical issue is to collect the power consumption data at short durations, so that more distinct on-off events can be detected.

Number of ON/OFF events. In many data center environments, the IT, cooling and power infrastructures all operate on a 24/7 basis. As a result, there are relatively few ON/OFF events. This means it can be difficult to bootstrap the analysis process, if the data center is already in operation. Scheduled or unscheduled downtimes offer an opportunity to observe ON/OFF events, although short measurement durations will be needed to distinguish between these events.

There are several reasons why ON/OFF events could become more or less scarce in the future. For example, more ON/OFF events could occur if dynamic provisioning of VMs becomes more prevalent (Kusic2009, Chen2010). Changes in such an environment could occur for durations of tens of seconds (the time scale of controller decisions). Similarly, ON/OFF events could also increase if individual servers turned ON/OFF more quickly to eliminate electricity consumption during idle periods (Meisner2009). These could occur over very short time scales (< 1 second), which would affect granularity of the power consumption measurements needed. However, if "energy-proportional computing" becomes a reality (Barroso2007), energy savings could occur without turning servers off. In this scenario, ON/OFF event would remain scarce, and identifying the sources of changed electricity consumption would still be problematic. Similarly, cooling infrastructure components such as chillers – although not completely energy proportional – can run at varying utilization levels consuming varying amount of power.

Complexity of device signatures. NALM works best for devices that have simple power consumption signatures (e.g., "ON" state consumes *N* watts, "OFF" state consumes 0 watts). However, this may not be the case for many devices in data centers. For example, many modern servers offer CPU frequency or voltage scaling to reduce energy consumption. There are many factors that could influence the choice of frequency/voltage over time.

Data quality. NALM is at the mercy of the power consumption data. If data is missing or recorded values are incorrect, the quality of NALM's results will be affected.

Competing technologies. An advantage of NALM is that it does not require pervasive instrumentation to be put in place. However, in data center environments some of this can happen in an automated manner. For example, some modern servers (e.g., HP ProLiant G6 servers) are now manufactured with embedded power meters. However, this does not negate the potential usefulness of disaggregation techniques. For example, NALM could be used to help identify the devices that contribute the "unknown" power consumption in the data center, once the consumption of the instrumented devices is subtracted from the aggregate load.

CONCLUSIONS AND FUTURE WORK

In this paper we provided an exploratory analysis of aggregate power metrics in data centers. We showed that power consumption data contains a wealth of information. We revealed macro-level information by visualizing the data. We then explored micro-level details by applying NALM on empirical data sets. Lastly, we described the challenges and opportunities that exist for using NALM in a data center.

There are numerous opportunities for future work on this topic. One is to collect or otherwise obtain finer-grained power consumption measurements from a data center, and explore how well it can identify components like servers, for example during scheduled or unscheduled power cycle events. Other disaggregation techniques could also be explored. For example, Patel *et al.* indicate noise signatures may disambiguate devices that turn on at the same time (Patel2007). Such techniques may require additional metrics to be collected, to distinguish between numerous identically configured devices (e.g., servers).

ACKNOWLEDGMENTS

We thank the IT staff at HP Labs and HP Bangalore who made the power consumption data available to us, and who provided information on the macro-events found in the data.

REFERENCES

[1]. L. Barroso and U. Holzle, “The case for energy-proportional computing”, *IEEE Computer,* Vol. 40, No. 12, pp. 33—37, Dec. 2007.

[2]. C. Belady. In the data center, power and cooling costs more than the IT equipment it supports. Electronics Cooling, 13(1), Feb 2007.

[3]. Y. Chen, D. Gmach, C. Hyser, Z. Wang, C. Bash, C. Hoover and S. Singhal, “Integrated management of application performance, power and cooling in data centers”, *IEEE/IFIP NOMS,* Osaka, Japan, Apr. 2010.

[4]. The Climate Group. Smart 2020: Enabling the low carbon economy in the information age, 2008. http://www.theclimategroup.org.

[5]. H. Dryar, “The Effect of Weather on the System Load”, *Transactions of the American Institute of Electrical Engineers,* Vol. 63, No. 12, pp. 1006—1013, 1944.

[6]. L. Farinaccio and R. Zmeureanu, “Using a pattern recognition approach to disaggregate the total electricity comsumption in a house into the major end uses”, *Energy and Buildings,* Vol. 30, pp. 245—259, 1999.

[7]. C. Greiner. Considerations for a 'green' energy-efficient data center. Ovum, 2008.

[8]. J. Han and M. Kamber, Data Mining: Concepts and Techniques, 2nd edition, Morgan Kaufmann, 2006.

[9]. G. Hart, “Nonintrusive appliance load monitoring”, Proceedings of IEEE, Vol. 80, No. 12, pp. 1870—1891, Dec. 1992.

[10]. D. Kusic, J. Kephart, J. Hanson, N. Kandasamy and G. Jiang, “Power and performance management of virtualized computing environments via lookahead control”, *Cluster Computing,* Vol. 12, No. 1, pp. 1—15, 2009.

[11]. C. Laughman, K. Lee, R. Cox, S. Shaw, S. Leeb, L. Norford and P. Armstrong, “Power Signature Analysis”, *IEEE Power and Energy,* Vol. 1, No. 2, pp. 56—63, Mar-Apr 2003.

[12]. H. Matthews, L. Soibelman, M. Berges, E. Goldman, “Automatically disaggregating the total electrical load in residential buildings: a profile of the required solution”, *Intelligent Computing in Engineering*, Plymouth, UK, Jul. 2008.

[13]. D. Meisner, B. Gold, and T. Wenisch, “PowerNap: Eliminating Server Idle Power”, *ASPLOS*, Washington, DC, Mar. 2009.

[14]. L. Norford and S. Leeb, “Non-intrusive electrical load monitoring in commercial buildings based on steady-state and transient load-detection algorithms”, *Energy and Buildings*, Vol. 24, pp. 51—64, 1996.

[15]. S. Patel, T. Robinson, J. Kientz, M. Reynolds, and G. Abowd, “At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line”, UbiComp, pp. 271—288, 2007.

[16]. T. Perry, “Tesla Tragedy: How Much Engineers Matter”, *IEEE Spectrum*, Vol. 47, No. 4, pp. 8, April 2010.

[17]. A. Prudenzi, “A neuron net-based procedure for identifying domestic appliances pattern-of-use from energy recordings at meter panel”, *IEEE Power Engineering Society Winter Meeting*, Vol. 2, pp—941-496, 2002.

[18]. U.S. Environmental Protection Agency. Report to congress on server and data center energy efficiency public law, pages 109-431, 2007.