

# An Adaptive Model of Energy Consumption Predictor for Big Data Centers

Xuexia Xu<sup>1</sup>, Gang Lin<sup>2</sup>, Jianzong Wang<sup>3</sup>

<sup>1</sup>Sports Engineering and Information Technology Department, Wuhan Sports University, Wuhan, China

<sup>2</sup>Experimental Center of Teaching and Research, Wuhan Sports University, Wuhan, China

<sup>3</sup>NetEase, Inc., Guangzhou, China

<sup>3</sup>jzwang@188.com

**Abstract** - It is well-known that with the explosive growth of data, the age of big data has arrived. However, the power consumption of the big data center is huge and will be a major obstacle to its wider extension. How to save the power towards green computing is a potential tendency by utilizing the elastic computing capability to analysis the power consumption over cloud computing platforms in big data centers. Anyway, the high energy cost of data centers have highlighted the energy models in the configuration of cloud providers' big data centers, which can estimate the energy consumption and adopt specific strategies to save the power within a given power budget. Moreover, the design of an adaptive energy model which can self-model the energy consumption on different conditions is quite challenging research. In this paper, we propose pBigData (Power Saving for Big Data Centers), an adaptive self-modeling paradigm, which can predict data centers' energy consumption, warn to exceed the threshold value and dynamically construct new model when existing ones become inadequate due to changes in hardware or different workloads. pBigData is to monitor the energy consumption according to the collecting data of devices including disks, processors, networks and temperatures, and then provide in-depth statistical analysis of the energy consumption. Our experiments reveal conclusively how accurate pBigData can enhance system energy efficiency while maintaining performance.

**Index Terms** - Big Data, Data Centers, Power Saving, Energy Consumption, Predictor, Early-warning

## 1. Introduction

Recently, power consumption has become a growing concern. Recent studies show that the energy used inside all U.S data centers is 1-2% of total U.S energy consumption. A typical data center with 1000 racks of 10MW total power consumption, costs 7M dollars for power and 4M-8M dollars for cooling per year. From the viewpoint of thermal effects, the increasing energy consumption leads to higher temperatures increasing the possibility of CPU and disk failures, which degrades the reliability of the whole system and constrains the further growth of system capacity. The increased power consumption also raise several environmental issues. The increased amount of carbon dioxide generated by hot servers can worsen global warming, while large water-cooled data centers located near lakes, produce heat pollution for the aquatic ecosystem. All of these issues result in strong incentives and motivations for designing energy efficient systems for current data centers and proposing early-warning to reduce the possibilities of devices' failures, thereby motivated the importance of and challenges to an adaptive modeling approach in data centers. There has been significant

amount of work on power and energy management strategies at both component and system level. At the component level, Dynamic Voltage Scaling (DVS) based control for CPU power, adjusting the power states of memory devices, and some researchers have investigated the use of multi-speed disks whose power consumption can be controlled by setting the speed. The principle behind these methods is to exploit the multiple power states available in the underlying devices, to trade off power consumption for speed in non-critical portions of the application.

Today, existing models for estimating power consumptions for data centers, mainly for single disk drive, cannot serve as simulator for large scale data centers. Energy models are highlighted the dependencies on configuring the data centers. Therefore, there are strong incentives to estimate energy consumption and adopt specific strategies to save the power within a given power budget. In this paper, we present pBigData, an adaptive self-modeling paradigm, which can predict data center's energy consumption, warn to exceed the threshold value and dynamically construct new model when existing ones become inadequate due to changes in hardware or different workloads. The accurate energy consumption of data centers in our model is determined by following factors: workload, CPU utilization, disk scheduling, caching effects, and networked storage configurations. The basic idea behind pBigData is to predict and monitor the energy consumption according to the statistical data produced by our adaptive energy model. It provides accurate energy consumption value and lower error rate by comparing with the value of actual measurement.

Specifically, this paper makes the following three main contributions: 1, Model big data centers' power consumption. We study the design and realization of pBigData, a novel paradigm for data centers to construct system energy models without assistance and with little knowledge regarding the system configuration or usage. 2, Apply the model to novel system energy management and optimization. We use and refine pBigData to diagnose system energy characteristics in order to account for energy consumed by hardware components and workloads. 3, Deploy the model for cloud computing. We design our frame-work as a module in Hadoop[1], a software framework that supports data-intensive distributed applications to prove that pBigData is extendability to the cloud computing environment. The rest of

this paper is structured as follows. Section II presents motivations, pBigData architecture and predictor models to estimate the power consumption. The prototype implement and experiment results are shown in Section III. Finally, discussion and future work are presented in Section IV.

## 2. The pBigData Frame Work

In this section, we first present our motivation for our energy modeling, then we introduce the architecture of pBigData.

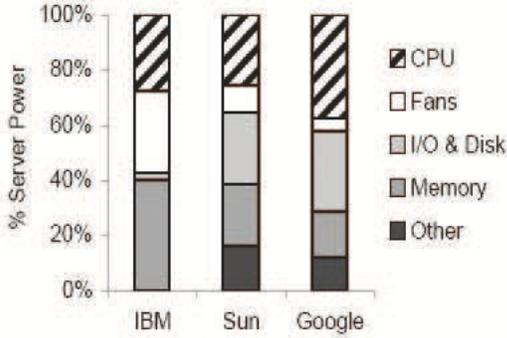


Fig. 1. Various servers with different configure

### A. Motivation

We demonstrate energy models can change drastically due to configuration changes of the system [5], as show in Figure 1. Existing system modeling methods provide energy estimation at a fixed, low resolution. In contrast, we seek to generate energy models with different temporal resolutions. It is worth nothing that most processor energy models can achieve much higher resolutions based on a detailed knowledge of the processor architecture and often require circuit-level component-wise energy models. pBigData, in contrast, needs little hardware information regarding the system and automatically selects the appropriate system parameters to build the model. Now for simplicity and without losing generality, we choose to model the power consumption of a server, whose hardware information given in Table I as done in [7]. We study three hardware configurations with only minor differences: a fixed CPU frequency of 2GHz, 800MHz; and flexible frequency as determined by DVFS. Figure 2 presents a linear function based its CUP utilization and extend to the linear energy model for each of the hardware configuration. The graph clearly shows that the differences between the three models are not negligible (as high as 25%) and we can predict the accurate energy consumption based on the linear relationship.

The system power can be expressed by a linear regression model. Therefore, it can indeed harvest the dependency on configuration and usage to construct and select models according the observed configuration and construct a model that is most accurate for the most observed usage. In next section, we detail design the system architecture of pBigData.

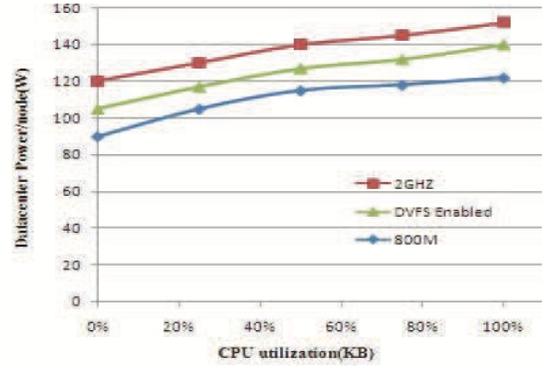


Fig. 2. Linear system power model based on CPU utilization with different configurations

Table I Hardware details of Data centers for data nodes

Name	Desription
OS	Fedora Core 8.0
Disks	Seagate ST3160023AS, 160GB7200RPM
Mainboard	GA-945GCMX-S2
CPU	Intel(R) Celeron(R) CPU 2.00GHz
NIC	Reltek 8169
Memory	512MB DDR2

Table II Hardware details of data centers for namenodes

Name	Desription
OS	Fedora Core 8.0
Disks	1 Seagate ST3160023AS, 160GB, 7200RPM. 8 Seagate ST3500320AS, 500GB, 7200RPM.
Mainboard	SUPER X7DVL-I
CPU	Intel(R) Xeon(R) CPU 5110 @ 1.60GHz
NIC	Intel(R) PRO/1000)
Memory	1G DDR2

### B. Energy model architecture

Figure 3 presents the architecture of our model, an extension of Hadoop cluster which is a cloud computing framework that supports data-intensive distributed applications. There are three major components: energy consumption data collector, model constructor, and energy predictor.

**Data Collector:** The data collector gathers Data Nodes' information from various hardware/software interfaces and standardizes it into equivalent quantizes, or predictor/response values through Name Nodes, and passes these values to the model constructor level of pBigData. The data collector also detects hardware and workload changes of the data centers and alerts the upper level, model construction to update the energy model accordingly. The data collector monitors the disk activity, CPU utilization, network performance and creates a level of abstraction which allows the rest of pBigData to focus on the statistical analysis.

**Model Constructor:** The model constructor is the core of pBigData. This adaptive energy model needs hardware and workload information regarding the data centers, and

automatically selects the appropriate parameters to build the model. Our energy model combines the linear model and non-parametric regressors for the numerical calculation of model generation. The constructor generates an initial set of predictors, and then it starts collecting data periodically. We measure total system powers by the initial model and try to find a curve that approximates the aggregate behavior to construct the accurate model. The model constructor will run if the accuracy of the model in use degrades beyond a preset threshold.

**Energy Predictor:** This component is the control center for dynamic energy management. It aims to propose the statistical predictor selection, and address the early-warning by recasting the data layout and task distribution of data centers. It is possible to enable significant portions of Data Nodes to be power down to avoid exceeding the power budget.

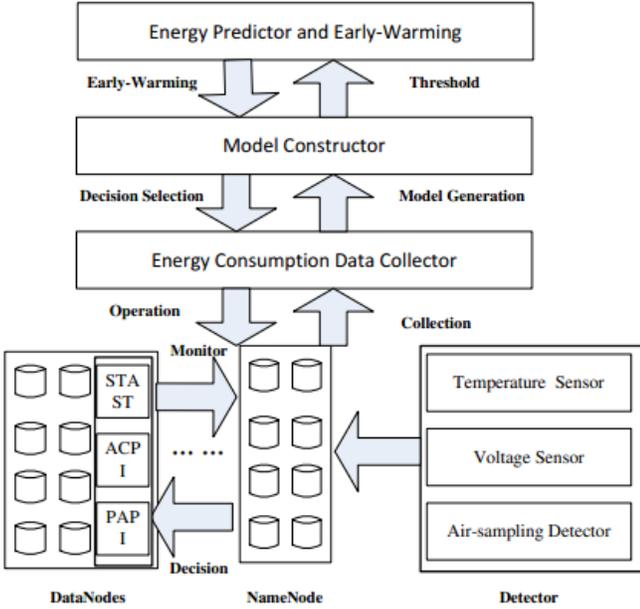


Fig.3. pBigData architecture

### C. Energy prediction modeling

We study the energy prediction model that is calculation unit of the pBigData. The model constructor essentially generates both the linear and regressive models. The combined model estimates the energy with the regressive model if the value of the predictor vector has been observed before. That is, if a predictor vector falls into a non-empty bin of the regressors, the regressive model is used for estimation. Otherwise, the linear model is used. This combination method is particularly useful when the number of predictors is large and there can be a huge number of empty bins in the regressive model.

Table III lists the basic set of parameters encountered in our model. After collection a large set of data samples of predictors and responses through the data collector, we generate the sample matrix in Equation 1. The model

constructor performs a correlation analysis by multiplying the correlation factor of various component then calculates the total energy consumption through adding all data nodes' by the energy predictor model. The accuracy of the model depends on the correlation analysis, so we will study the trade-off between accuracy and predictor selection in next section

$$\varepsilon = \frac{P_{pred} - P_{meas}}{P_{meas}} \quad (1)$$

Table III prediction model parameters

Parameter	Description
$P_{pred}$	Total energy consumption of data centers
$P_{pred,i}$	Energy consumption of $i$ th data node
$n$	Number of data nodes
$i$	The $i$ th data nodes
$P_{meas}$	The total power consumption measured by P3 International's Kill A Watt
$\varepsilon$	The accuracy of power prediction

### 3. Experiments

In this section, we first describe a way of setting up a cluster as a testbed then presents our experimental methodology that we use to study quantitatively the prediction accuracy of the pBigData prototype in Hadoop cluster while maintaining performance.

#### A. Testbed

The testbed is illustrated in Figure 4. Our experimental setting consists of a data center that forms a name node server, 2 data node servers and a monitor that is monitoring the temperature and voltages of all nodes. All components are interconnected using the Cisco 3750 Gb switcher. The hardware details of name nodes and data nodes are described in Table II and Table I respectively.

#### B. Overhead of time and space in pBigData

Our prototype code is implemented by modifying the Hadoop program. We use the Ganglia [2], a scalable distributed monitoring system for high-performance computing systems such as clusters, to collect information of devices. Then the overall structure of pBigData is shown in Figure 3 as a software module embedded in the standard Hadoop written by metrics API in order to generate the model and predict the energy consumption.

From data storage point of view, the data collector will generate the sample data set about 14M/day, so the data should clean up per week as described in Figure 5.

The below picture in Figure 5 depicts the time overhead (10ms-50ms) of various predictor ways are acceptable if using the predictor selected by pBigData. To reduce the overhead, we design a new system call that is able to read multiple data structure in the Ganglia at the same time over data collector module. It shows that this schema reduces the time by five times compared with all predictors.

### C. Prediction accuracy comparison

The adaptive model is valuable or not depends on its prediction accuracy so we formulize the accuracy  $\varepsilon$  in Equation 2, where  $P_{meas}$  represents the total power consumption measured by P3 International's Kill A Watt in data centers' electricity supply.

$$\varepsilon = \frac{P_{pred} - P_{meas}}{P_{meas}} \quad (2)$$

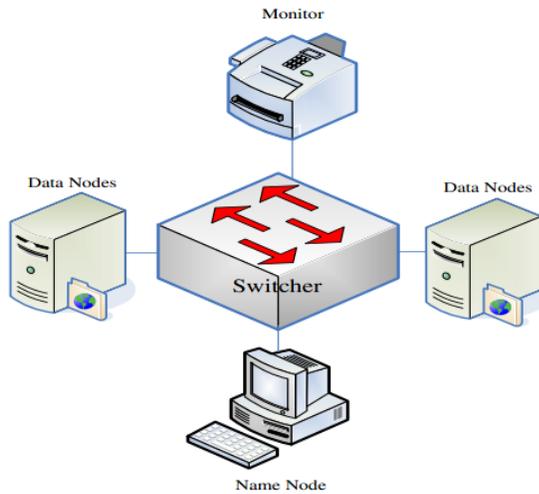


Fig.4. pBigData testbed in Hadoop cluster

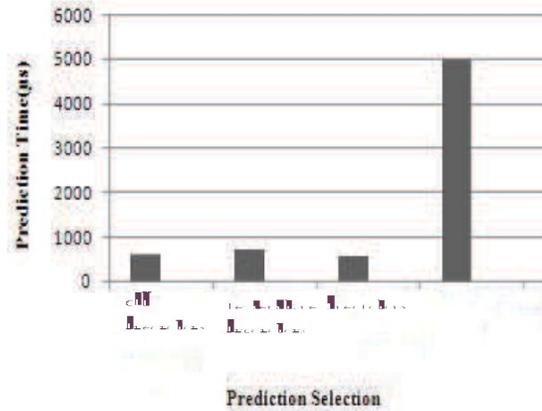
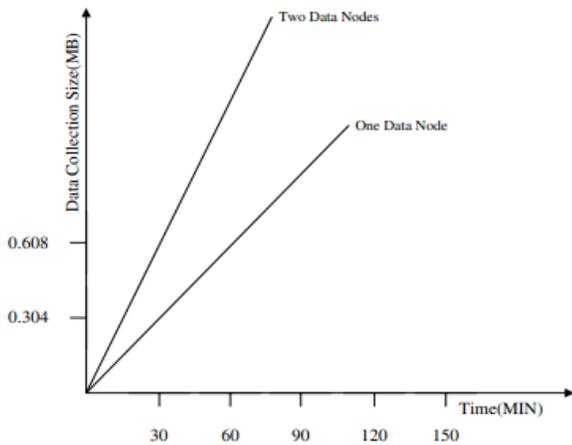


Fig.5. Overhead of space time in pBigData

Figure 6 summarizes the results. The results show that the prototype was able to build energy models; it converged in several hours of usage and provided energy models with error below 15%. Such accuracy is encouraging, and comparable with state-of-the-art system energy modeling methods that heavily leverage knowledge regarding system power characteristics and perform data collection and model construction within name node. More importantly, our prototype achieved this without the loss of performance.

### 4. Conclusion

Optimizing energy efficiency of data centers remains an open issue. We change the point of view in this research that we address the perdition problem of energy consumption then process the early-warning to achieve the aim of power saving. The key idea behind pBigData is to select the appropriate predictors using the corresponding regressors function to obtain the system energy consumption through multiplying the correlation factor of various components. We derive and validate performance and energy models. The results measured on the prototype are indicating that pBigData are valuable for the energy consumption prediction in data centers. We plan to expand the underway research to address the alerts generated by the model. The model could select the appropriate schema to save the power by recasting the data layout and task distribution among the cluster. Our long-term goal is to develop energy consumption prediction benchmark embedded in Hadoop as a part of standard over the cloud computing platform.

### References

- [1] Hadoop. <http://hadoop.apache.org/>.
- [2] Ganglia. <http://ganglia.sourceforge.net/>.
- [3] X.J. Ruan, S. Yin, A. Manzanaraes, J. Xie, Z.Y. Ding, J. Majors, and X.Qin. ECOS: An Energy-Efficient Cluster Storage System. In IPCC, Phoenix, Arizona, USA, Dec. 2009.
- [4] X. Fan, W. Webber, and L. Barroso. Power provisioning for a warehouse- sized computer. In ISCA, San Diego, USA, 2007.
- [5] David Meisner, Brian T. Gold, Thomas F. Wenisch. PowerNap: Eliminating Server Idle Power. In Proceedings of ASPLOS09, Washington DC, USA, 2009.

- [6] J. S. Chase, D. C. Anderson, P. N. Thakar, A. M. Vahdat, and R. P. Doyle. Managing Energy and Server Resources in Hosting Centers. SOSP2001.
- [7] D. Economou, S. Rivoire, C. Kozyrakis, and P.Ranganathan. Full-system power analysis and modeling for server environments. In MOBS, Boston, MA, USA, 2006.
- [8] M. Allalouf, Y. Arbitman, M. Factor, R. Kat, K. Meth, and D. Naor. Storage modeling for power estimation. In Proceedings of SYSTOR. ACM New York, NY, USA, 2009.
- [9] J. Zedlewski, S. Sobit, N. Garg, F. Zheng, A. Krishnamurthy, and R.Y. Wang. Modeling hard-disk power consumption. FAST2003.
- [10] J.Z.Wang. Research on the framework of high performance computing towards Big Data. HPC China2013.
- [11] C.Yin, J.Z.Wang, C.S.Xie, J.G.Wan. Robot: An Efficient Model for Big Data Storage Systems Based On Erasure Coding. BigData2013.
- [12] J.Z.Wang, W.J.Gong, P. Varman, C.S.Xie. Reducing Storage Overhead with Small Write Bottleneck Avoiding in Cloud RAID System. GRID2012.