Online Power Measurement and Prediction of PCs

JIANG Congfeng^{1, a}, ZHAO Yinghui^{2, b}

¹Grid and Service Computing Technology Lab, Hangzhou Dianzi University, Hangzhou, 310037, China

²Department of Hydraulic Engineering and Agriculture, Zhejiang Tongji Vocational College of Science and Technology, Hangzhou, 311231, China

^acjiang@hdu.edu.cn, ^bzhaoyinghuihust@gmail.com

Keywords: online measurement, power consumption, workload characterization, power aware computing system

Abstract. Since more power consumption results in more failures and degradations in system performance, reliability, and power bills, it has been a critical problem for not only large scale server system but also personal computers (PCs). Though much literature has focused on energy management and power budgeting for server systems, power consumption of PCs does not attain sufficient attentions fairly. In this paper an online power measurement and prediction framework is proposed and used to save more energy considering the PC as a whole controlled system. The framework includes parts such as power measurement unit, power prediction unit and a simple execution unit of power reduction decisions. A hardware-software joint prototype is implemented based on an intelligent digital multimeter. Experiments on a desktop PC and a laptop show that PC with the framework can save more power consumptions than that of the PCs without this framework.

Introduction

In the past decades, the performance of the PC has increased approximately 500 times (e.g., CPU frequency from 6MHz to 3GHz) due to the intensity increase of VLSIs. However, with the tremendous growth of personal computers (PCs) utilization in homes and offices, their power consumptions are often a topic of discussion concerning energy utilization, carbon emissions, and sustainable development in recent years. Upon examination, few of the users of PCs in homes and offices sector have a consistent track record of increased product performance with increased power consumptions. Currently, PCs have much potential to explore to reduce the overall power consumptions because power consumption of PCs and the peripherals will remain a potentially and increasingly significant part of the overall power budgets in the long run. Since a typical PC is configured for peak power consumption in the worst-case, PCs are often underutilized and the actual power consumption while running most software applications is far less than the peak power. Most of the time, actual power consumption for a PC is far less than the nameplate rating. Therefore, this leads to over-demand for power and cooling infrastructure and provide a significant opportunity in simply enabling power management technologies.

Moreover, with the scale of computing system increases, power consumption has become the major challenge to system performance and reliability [1, 2, and 3]. For example, server farms today consume more than 1.5% of the total electricity in the U.S. at a cost of nearly \$4.5 billion, which is estimated to about 61 billion kilowatt-hours (kWh) in 2006 and it is more than the electricity consumed by the nation's color televisions and similar to the amount of electricity consumed by approximately 5.8 million average U.S. households [4]. According to the energy consumption, 40% of energy use is contributed to desktop PCs, most of which can be turned off at night or when not in use. The U.S. EPA estimates that if all PCs in the U.S. used power saving technology, the US could save more than \$300 million in energy costs each year. This kind of power reduction would reduce air pollution equivalent to removing more than a half-million cars from America's roads each year.

Various companies have developed some kinds of solutions for large scale enterprises or organizations with hundreds of computers installed, such as NightWatchman [5], LANDesk [6],

Faronics Power Save [7], Verdiem SURVEYOR [8]. However, these package solutions are more passive than active in operation modes. Although modern CPUs and graphics cards can change their power consumption code according to real workload and instructions execution, if more features, such as self-learning, workload characterization and prediction are integrated into the existing software, they can save more energy for computers running periodical tasks.

In this paper, we investigate the feasibility of hardware-software joint measurement and prediction of PC power consumption based on simple hardware implementation. In the proposed framework, we use a simple feedback algorithm to accumulate and predict the behavior and power consumption. In the framework, the system performance and different workload characteristics will be collected dynamically through a background module and it is implemented in a software package. A testbed for investigating this framework is implemented on a PC using Intel quad core processor. We use a multimeter with USB connection cable to measure the real time power consumption of the system. In the multimeter there is a special microcontroller for integrating the current and voltage over time to calculate the amount of active power consumed in the load. The performance and overhead of the algorithms are assessed under different usage workload.

The remainder of this paper is organized as follows: In section 2 we review some related work on power aware computing system and energy management. In section 3 we propose the metrics and the measuring framework. In section 4 we present the experiment implementation of the proposed framework and we present experiment results and analysis. Finally, we summarize the work in Section 5.

Related Work

Power consumption has become a main concern for enterprise server system in order to achieve a high quality of service level in terms of, for example, availability and reliability. The adaptive power management has to ensure peak power safety but should also intelligently schedule individual jobs to computing resources to guarantee negotiated Service Level Agreements or Quality of Service. The Service Level Agreement is an electronic contract between a service user and a provider, and defines service quality like online time, response time, failure percentage, etc. From a general-purpose viewpoint, performance can be defined by SLA constraints for corresponding underlying workload heterogeneity [9].

Since the overall system power consumption has strong relationship to processor resource usage, various excellent Dynamic Voltage Scaling/Dynamic Frequency Scaling (DVS/DFS) algorithms have been proposed to reduce the power consumption of processors [10,11]. Memory is another source contribution for power consumptions in server systems. Existing techniques usually manage power for the main memory by passively monitoring the memory traffic and regulation. Some algorithms are proposed to predict when to power down which memory units and into which low-power state to transition [12,13]. In large scale data centers, server systems, or in the enterprise storage system, power consumption of hard disks is a critical issue where data-intensive applications exhaust disk storage extensively. Since energy consumptions directly affects hard disk drive reliability and system performance, reducing energy consumptions of disks can dramatically save overall enterprise IT costs [14,15].

Since virtual machine provides time-sharing and resource-sharing on the highly expensive hardware and makes it possible to execute, develop, and test applications concurrently, it was widely used to reduce the hardware acquisition cost and improving the productivity by working on it simultaneously. Currently, server virtualization is a mainstream technology and various worldwide enterprises and organizations have adopted it as basic IT infrastructure to provide many benefits such as server consolidation, application and service consolidation, etc[16]. Moreover, virtualization is ever-growingly deployed in data centers for power and cost savings.

Mor Harchol-Balter et al [17] investigates the performance of task assignment policies for server farms, as the variability of job sizes (service demands) approaches infinity and they found that the

Size-Interval-Task-Assignment policy (SITA), which assigns each server a unique size range, was inferior to the much simpler greedy policy, Least-Work-Left (LWL), for certain common job-size distributions, including many modal, hyper-exponential, and Pareto distributions.

As the PC becomes more and more cheap and multi-functional and high performance, it consumes more and more electricity at the same time, especially for the desktop computers. The power management technology, advanced power management (APM), was introduced by Intel, Microsoft and other manufacturers to enable hardware and software interaction to reduce power consumptions in PCs. For example, the Intel SpeedStep® technology can adjust the speed of the processor dependent on power source and processor behavior to provide demand-based performance. When the system is underutilized or has less demanding, the processor then steps-down to a lower power state. APM is replaced by Advanced Configuration and Power Interface (ACPI). ACPI provides a standard and flexible interface between the underlying hardware and applications to implement power management capabilities via the operating system as the operating system has the most knowledge regarding a system's power demands and real time workloads and is the best way to perform power management mechanisms. For example, in ACPI specification, CPU C-states can be used to reduce power consumption during processor idle time since real-world workloads rarely saturate multi-core processors. However, the key problem is when and how to use which C-state.

In this paper, we propose an online power measurement and prediction framework for personal computers that can provide relatively accurate power consumption information of the whole hardware components. We introduce a hardware software jointly assisted power measurement system that integrates with the host platform and provides detailed power information of multiple hardware components. This framework is configured to take advantage of the power-management capabilities of the hardware. We also propose a self-learning based behavioral model to predict the workload and power consumption information on the system. In the model the prediction is configured as a background process and acts as a component of the operating system that has the capability to predict the power and application performance of the computer it manages. The experiment results show that the prediction model can provide accurate future workload information and add on little overhead.

Metrics and the Measuring Framework

Metrics. In order to measure the power reduction, the power evaluation of implemented framework must base on the practical and real-time power consumption metrics. Accurate power consumption characteristics and measurement is vital to help activate power reduction mechanism to make right decisions according with preset selection criteria. Although some computing system vendors report some power-performance figures, but they are often not directly comparable due to differences in workload, configuration, test environment, etc.

Power Usage Effectiveness (PUE) metric is often used to evaluate the efficiency of data centers [18]. PUE is the ratio of the total power consumed by a data center to the power consumed by the IT equipment that operates the facility. PUE puts emphasis on maximizing the power supplied to the equipment running customer applications and minimizing the power consumed by cooling and power distribution. For example, a PUE of 2.0 indicates that for every watt of IT power, an additional watt is consumed to cool and distribute power to the IT equipment.

However, PUE is not suitable for PC power consumption measurement because PCs usually do not use additional cooling equipments and infrastructures. And it is usually hard to get comprehensive data on the power consumptions of specific components. In order to reflect the actual power consumption of the PCs, we just use the total power consumed by the monitor and the main components of it from the wall outlet. This also makes it easy to implement and compare different scenarios directly from the wall plug.

In this paper we do not consider the situation that the computer will receive operating system patches or antivirus definition during night when it is idle. The PSU's efficiency is not taken into account either. For example, at a load of 400W, an 80% efficient PSU will consume $400/0.8=500W$

from the mains. But this does not affect the resulting measurement accuracy as we just want to show the power consumption trend under different workloads and this constant behavior can be ignored in analysis. Although the wall outlet measurement won't show how much power is consumed by the components such as graphics card and CPU, the trend of the overall consumption of the computer is already enough.

Measurement Framework. It is commonly believed that system performance can be maximized by operating servers at their highest power levels under a given power budget. The main idea of this paper is the intuition that in lower loaded periods, there is a potential to save power consumption by dynamically powering off part or whole of the computer to address the actual computing demands. Under such lower-load conditions, an appropriate power management scheme can considerably reduce power consumption. In this paper we propose a framework to help figure out the total system power consumption with simple hardware selections. The proposed framework in this paper contains three parts: workload characterization and prediction, power consumption measuring and estimation, feedback control of power consumption. Fig.2 illustrates the framework.

Fig. 1 System Framework of Feedback Based Controlling.

Fig.1 shows the logical configuration of the whole computing including measuring part, workload prediction and hardware resources such as CPU, memories, network, etc. Fig.2 gives a simplified diagram for power measurement. The power consumption is measured by a multimeter and the results are transported to the dedicated computer through a USB connection.

Fig. 2 Framework of Power Measurement

Workload Characterization. Workload characterization is an important technique to help understand the performance of system applications and demands. A good understanding of the workload of multi processor computing system should provide insights about user requests, program activities and help in improving the quality of the service provided to users. In a real computing

system, workloads are varying dynamically due to various reasons such as user behaviors, usage modes of diverse applications, unforeseen load fluctuations, etc. Therefore, workload characterization acts as a crucial component of any performance analysis process and researchers often use processors runtime statistics, trace data of real workload and algorithms such as online learning, clustering, and fractal, to predict the workload of different applications in the future. However, since there are no specific workload characterization algorithms that can characterize every kind of workloads in different computing systems and applications, workload characterization and prediction schemes must be developed for specific individual application.

Depending on the type of applications, the workload characterization model could vary in complexity ranging from a simple specification of the application demands of the server resources, especially the processor cycles, to a detailed mathematical model capturing the flow of control during application execution. We assume that the PC users execute some logic requiring a given amount of CPU time. Whenever a task is processed at a host, the actual execution time is compared with the current estimate of the task demand and if it is lower the estimate is updated. Thus, as the task progresses, the estimated task execution time is iteratively set to the more accurate observed execution time on the respective host.

In this paper, we proposed a lightweight resource-oriented workload characterization scheme which describes the consumption of system resources by the workload, such as length of processor queue, processor time, waiting time, disk operations, memory, etc. In resource-oriented workload characterization scheme, resource data can be collected easily by automated agents or third party packages which use the OS calls or APIs. We monitor the CPU and memory utilization during task execution and use it to split the measured observed time into time spent using the CPU and time spent waiting. This algorithm is conservative in that it starts with conservative estimates of the CPU cycles and refines them iteratively in the course of the execution. The workload characteristics can be represented by observations in macro and micro scale. The former includes resource-oriented performance, such as response time, processor time, memory utilization, disk operations, etc. And the latter includes processor instructions execution time, status of registers, status of threads, hardware counters, etc.

Experiment Implementation and Results

Experiment Setup. In the experiment, the computer is tested to know its actual power consumption by changing various parameters. To measure the actual power used by the computer, the power flowing into it is metered by a multimeter. To show specific power reduction potentials, the measurements were mainly done at three different configurations: idle or standby status, a mid-utilized status, and then highly-utilized status. For this to be done, the video card and processor is loaded to their maximum capacity. A typical mid-utilized usage stands for the workload editing a document, reading an audio file and surfing the Internet at the same time.

Table 1 gives a comparison on hardware parameters of two experiment platforms, including CPU frequency, memory capacity, hard drive capacity, monitor type, video card type and power adapter type. The workload is classified as normal mode, highly loaded mode, and underutilized (least usage) mode. They are N1, H1, U1 and N2, H2, U2 for case 1 and case 2 respectively.

The experiment results are shown in Figure 3 according to tests conducted. the normal mode stands for the workload including Internet surfing, audio playing, document editing and photo editing, the Highly loaded mode stands for the workload including Database service instances and connection(Oracle 11G R1 Enterprise edition and Microsoft SQL server 2008), Programming with Microsoft Visual studio 2008, 3D Gaming(World of Warcraft), and the underutilized (least usage) mode stands for the workload including only running basic services of operating system with least power management, shutting down all applications including antivirus and firewall, no hard disk spin, and the computer is with least lightness of the monitor.

Configuration Items	Case 1	Case 2
Type of testbed	Dell 1420 INSPIRON Notebook	Commodity Desktop Computer
CPU	Intel Core Duo T7250 at 2.0GHz frequency ,800MHz front bus and 2MB L2 Cache	Intel Core i7 Quad core with 8MB L2Cache
Memory	2GB	6GB
Hard drive	320GB notebook hard disk	500GB desktop hard disk
Monitor	14.1inch UltraSharp WXGA+ (1440x900)	22 inch Hanns.G DVI+ (1680x1050)
Video Card	NVIDIA 128MB DDR3GeForceTM 8400M GS	ATI 512MB 4800 platinum edition
Power Adapter	90W (110V/220V)	500W (220V)

Table 1. Hardware Parameters of Two Experiment Platforms

Fig. 3 Experiment Results (watts)

Figure 3 lists the result including the consumption of the configuration when varying the operation model and workload. Since it is hard to isolate the processor, memory and hard disk, we only give the whole power consumption.

Since most current monitor has been energy star certificated, they consume little power when power saver mode is activated. Therefore, the proposed framework affects the whole power consumption more on the box than on the monitor. However, the number will vary depending on the processor and what other peripherals connected.

Based on these results, we can see that the average power consumption of a high-end computer system is about 1540kwh a year, if the computer is left on 24/7. This cost, however, does not include the electricity used by printers, scanners, and other devices. Note that these prices don't include the cost to run the monitor, see the second table.

	N1	H1	U1	N2	H2	U2	SQL Server 2008	Visual Studio 2008	WOW
MAX	43.1	53.9	28.5	239.9	281	156.6	236.9	213.6	252.8
MIN	21.3	34	14.5	65.3	73.5	54.9	217.2	204.1	233.4
AVG	31.2	47.5	15.5	155.4	162.4	105.8	228.3	207.5	244.9

Table 2. Key Values of Various Workloads (Watts)

	N1	H1	U1	N2	H2	U2	SQL Server 2008	Visual Studio 2008	WOW
MAX	48.3	57.4	32.3	245.7	287.3	162.2	243.8	217.5	258.3
MIN	23.5	37.2	18.1	69.9	78.2	58.4	222.4	208.5	239.8
AVG	35.9	51.1	22.6	155.4	162.4	105.8	231.5	211.7	248.6

Table 3. Original Key Values of Various Workloads without Workload Characterization (Watts)

Conclusions

Although the IT industry has an excellent track record in developing energy-efficient products and PCs are currently extremely well power managed, there is still room for improvement. In this paper an online power measurement and prediction framework is proposed to save more energy considering the PC as a whole controlled system. The framework includes parts such as power measurement unit, power prediction unit and a simple execution unit of power reduction decisions. A hardware-software prototype is implemented based on an intelligent digital multimeter. Experiments on a desktop PC and a laptop show that PC with the framework can save more power consumptions than that of the PCs without this framework. The results also show that there could be a lot of room for improvement when online measurement and prediction is integrated while the system does not reach sleep mode levels. Since there is not always a heavy load for even high end PC, power reduction seems to be not only nontrivial but also a priority from an economic and environmental point of view.

Compared to the experiment without the proposed framework, we notice a decrease in the power consumption to 5W, which is related to the fact that the self-learning and prediction scheme of the proposed framework. The framework has different energy efficiency rate at 20%, 50%, and 100% of loading with a constant power factor of 0.9 or greater.

Please also note that the performance depends on many factors such as the accuracy of workload characterization and the number of power consumption is slightly different with and without reading media. It is also true with the sound card and it is negligible.

References

- [1] C. Lefurgy, K. Rajamani, F. Rawson, W. Felter, M. Kistler, and T.W. Keller: *Computer*, Vol.36 (2003), p.39
- [2] X. Wang, M. Chen, C. Lefurgy, and T.W. Keller, in: *Proceedings of 2009 18th International Conference on Parallel Architectures and Compilation Techniques (PACT'09)*, ACM Press, New York (2009)
- [3] X. Wang, and M. Chen, in: *Proceedings of IEEE 14th International Symposium on High Performance Computer Architecture (HPCA 2008)*, IEEE CS Press, Washington (2008)
- [4]http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf
- [5] <http://www.1e.com/>
- [6] <http://www.landesk.com/>
- [7] <http://www.faronics.com/html/powersave.asp>
- [8] <http://www.verdiem.com/>
- [9] Y. Guo, Solihin, L. Zhao, and R. Iyer, in: *Proceedings of the 40th Annual IEEE/ACM Symposium on Micro architecture (MICRO-2007)*, Los Alamitos: IEEE Computer Society Press (2007)
- [10] A. Miyoshi, C. Lefurgy, E. V. Hensbergen, R. Rajamony, and R. Rajkumar, in: *Proceedings of the 16th Annual ACM International Conference on Supercomputing (ICS'02)*, ACM Press, New York (2002)
- [11] W. Kim, M. Gupta, G. Wei, D. Brooks, in: *Proceedings of 14th International Symposium on High-Performance Computer Architecture (HPCA2008)*, IEEE CS Press, Washington (2008)
- [12] H. Huang, C. Lefurgy, T. Keller, and K. G. Shin, in: *Proceedings of the International Symposium on Low-Power Electronics and Design (ISLPED'05)*, IEEE CS Press, Washington (2005)
- [13] I. Hur, C. Lin, in: *Proceedings of 14th International Symposium on High-Performance Computer Architecture (HPCA2008)*, IEEE CS Press, Washington (2008)
- [14] Q. Zhu, Z. Chen, L. Tan, Y. Zhou, K. Keeton and J. Wilkes: *ACM SIGOPS Operating Systems Review*, vol. 39(2005), p.177
- [15] S. Gurumurthi, A. Sivasubramaniam, M. Kandemir, and H. Franke, in: *Proceedings of 30th Annual International Symposium on Computer Architecture (ISCA-30)*, Los Alamitos: IEEE Computer Society Press (2003)
- [16] R. Figueiredo, A.P. Dinda, J. Fortes: *Computer*, Vol.38 (2005), p. 28
- [17] Mor Harchol-Balter, A. Scheller-Wolf, A. Young, in: *Proceedings of ACM Sigmetrics 2009 Conference on Measurement and Modeling of Computer Systems (Sigmetrics '09)*, ACM Press, New York (2009)
- [18]http://www.thegreengrid.org/gg_content/TGG_Data_Center_Power_Efficiency_Metrics_PUE_and_DCiE.pdf