*International Journal of Advances in Computer Science & Its Applications – IJCSIA*
*Volume 4: Issue 4* [ISSN 2250-3765]
*Publication Date : 27 December,2014*

# A Technique to Minimize the Power Consumption of Computer Fleets

[Daniel F. García, Joaquín Entrialgo, Ramón Medrano, Javier García]

***Abstract*— Energy consumption of current computer fleets of most organizations can be substantially reduced by using power management systems. Energy saving is an important matter as, with the sustained increment of the cost of electricity, the monetary savings can be significant. This paper presents a technique to minimize the power consumption of the computer fleet or an organization. The technique is based on a distributed system where a central manager coordinates the local power manager of each computer. A method to determine the profitability of this type of management system is also proposed. The key element to save energy is the idle timer used by local power controllers. When the controller detects that the computer is inactive, it starts the timer, and when the timer expires, the controller puts the computer in suspension or hibernation, or turns off the computer. This paper proposes a systematic technique to estimate the appropriate values of the idle timers for the computers of the fleet. Any administrator of the fleet can take advantage of the technique proposed in this paper to achieve an appropriate trade-off between energy saving and user productivity. The paper also shows the application of the technique to the computer fleet of a company.**

***Keywords*—computer power management systems; enervy saving applications; consumption minimization algorithm**

## I. Introduction

The cost of electricity is increasing continuously and its proportion of the total operation costs of computer systems is growing [1]. Therefore, most organizations are implementing measures to control the cost of energy consumed by their computers [2] [3]. This is really important in medium and large sized companies, in which the total number of computers in their fleets is very significant.

All types of computers and devices can be considered, from servers to desktops and laptops. But in particular, this article is focused on desktop computers. The reason is that this type of computer is the most frequently used in organizations and offers the greatest opportunities for power saving. The number of servers used by most organizations is much lower than desktop computers and they are more closely supervised by the IT department of the organization. Laptops are usually directly managed by their own users and consume less power than desktops.

Daniel F. Garcia, Joaquín Entrialgo, Ramon Medrano, Javier Garcia
University of Oviedo
Spain

Many studies have been carried out concluding that the vast majority of desktop computers are left powered on after being used for a certain period of time [4], that the power usage depends directly on the number of computers [5], and that there is much space for power saving in optimal conditions [6].

Dynamic Power Management is a general concept that was introduced some years ago to save energy in computing systems. Benini et al. [7] defined how a power manager can control and optimize the use of the different power states featured by a computer in order to save energy. Their methods have been used in some techniques for power management of the OS [8]. This work defined the road map for power management strategies for single computers, but did not provide an approach to the management of large computer fleets from the power optimization perspective.

Large companies and institutes might have computer fleets with several thousand computers. The general management of these computers is a problem which has already been tackled by system administrators. Tools like Microsoft Active Directory, NIS (Network Information Service), Puppet, AFS (Andrew File System) or CVMFS (CernVM File System) make the management and the data and software distribution [9] to thousands of workstations, laptops and servers across many countries relatively straight forward with a few system administrators. However, the power management functionality of the computers is not always integrated in this kind of system management tools.

Nevertheless, there are several centralized Personal Computer Power Management (PCPM) commercial tools [10]. 1E NightWatchman [11], IBM Tivoli [12], and Verdiem Surveyor [13] are the most prominent. These tools offer a centralized dashboard to system administrators to define different power policies for each group of computers, block users from changing the power policy of their computers and perform some basic monitoring on the usage of the computers.

There are also totally free tools, such as SpiceWorks PowerManager [14], or others like Granola Enterprise [15] that provide the analysis of the energy consumption for free but require the user to pay for the control functions.

Before using one of these tools permanently in an organization, the IT administrator must perform tests to determine whether they will generate substantial savings or not. The vendors of these tools typically provide a trial period (30 days) for their evaluation and give guidance on how to use them, although not always in a clear and systematic manner. This paper provides a comprehensive and systematic method, which includes several phases, to determine the potential savings that one of these tools can provide.

Once the decision to deploy a PCPM system is taken, the administrator must assign a proper value to the idle timer of the computers. For that, several new algorithms have been proposed. They extend the seminal works, including adaptation [16] and learning capabilities [17] in the optimization process used to find the proper value of the idle timer. The use of these algorithms is complex, as they require the modeling of the computer system and its workload, the definition of cost functions and their minimization with complex optimization techniques.

In this paper a more practical method is proposed to estimate the idle timer values appropriated to minimize the energy consumed by the computer fleet of an organization. Utilizing the available idle periods (previously measured to determine the potential savings), two metrics (energy waste and user satisfaction) are represented as a function of the value of the idle timer. Then, the administrator can easily select an adequate value for the idle timer by weighing the two metrics.

## II. System Architecture

A simple approach to managing the power consumption of the computer fleet is using the mechanisms integrated in the operating systems of computers. As an example, the Windows operating system (OS) allows the user to select a power plan. The plan contains several settings for the power management module of the OS. Basically, the plan contains the required period of inactivity of several components in order for the OS to turn them off. This is typically applied to hard disks and the display. The plan also contains the idle period required to put the whole computer in suspension or hibernation.

The OS monitors the activity of many components of the computer continuously. When a component has been idle during the predefined period of time, the OS puts the component in a low-power state or turns the component off. In particular, the monitoring of the user input is very important to detect the inactivity due to the absence of keystrokes or mouse movements.

Using only the mechanisms integrated in the OS, the IT administrator must prepare a power plan for each computer of the organization and activate it in the target computer, eliminating the permissions of any user to change the active power plan. Our approach uses a distributed system to create and distribute these power plans automatically.

Fig. 1 shows the system architecture. In each computer, there is an agent with two main objectives:

1) Monitor the activity of the user of the computer and send the information about the utilization/idle periods to a central management server.

2) Receive power plans from the central management server and apply them depending on the time. Typically, there will be different plans for working and non-working hours.

The central management server stores the information gathered from agents in a database. It also provides an administration console that allows a centralized control of the computers of the fleet. It is typically implemented as a web application and permits the visualization of past or current power states of the computers of the fleet and the interactive elaboration and enforcement of power control strategies.

Another important element of the centralized approach, depicted in Fig. 1, is the wake-up console. It is also implemented as a web application. It provides users with a mechanism to wake up their computers remotely, if they have enough privileges, to be able to work at times in which the computers should be switched off.

This distributed system architecture also provides fault tolerance. When there is no communication available, any agent continues controlling the computer with the latest settings received and gathering and storing information about the operation of the computer locally. When the communication is reestablished, the agent can send the information and receive new control settings.

## III. Determining the Profitability of the Power Management System

The start-up and utilization of a power management system in an organization requires a methodology based on several sequential phases.

### A. *Measurement of the on/off and active/idle periods*

In the first phase, the agent must be installed in all the computers of the fleet to be controlled. In this phase, the agents never modify the power plan used by the operating system, and the computers must be used normally, as prior to the installation of the agent. The agents gather information about the power state of the computers and, within the turned on periods, the inactivity intervals.

It is necessary to have data from at least two full weeks of operation, although it would be desirable to extend the period to three or four weeks. The weeks must not contain holidays, as they would distort the information. With this information, a graph such as that shown in Fig. 2 can be obtained.

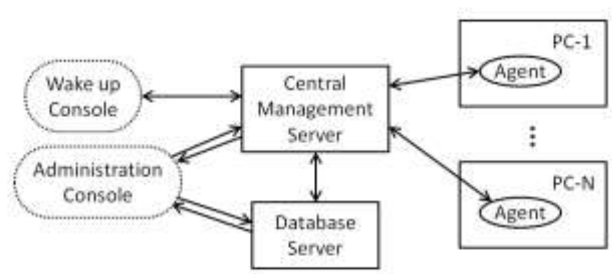


Figure 1. Power management approach integrating a central manager for the agents.

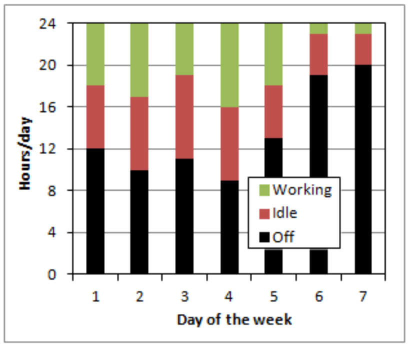


Figure 2. Weekly utilization of a computer or a group of computers

For each day of the week, the black portion of the bars represents the percentage of time that the computer is off. The colored portions represent the percentage of time that the computer is on. Two parts can be distinguished: when the computer is on but idle (red part) and when the computer is on and is being used (green part). If the size of the red segment is significant, the utilization of a power management system can generate substantial economic savings.

The main objective of the power management system is converting a great part of the red segment in a black segment, that is, when a computer is idle it must be turned off (or put in suspension or hibernation). Of course, the shutdown should be done without affecting user productivity, that is, a computer that is doing useful work should not be turned off.

### B. *Estimation of the total energy consumption*

To estimate the energy consumption of the computer fleet over a defined period of time, the hours that each computer is turned on have to be converted in the watt hours consumed. Generally this is done with power consumption models.

Each computer of the fleet must be identified and its components (type, brand and model of CPU, disk, etc.) must be perfectly known. We must also have a database of the computer components that contains the power consumed by each component. Generally, this information can be obtained from manufacturers. Then, the power consumption of any computer when switched on can be estimated by simply adding the consumption of all its components.

Another option is to use data of the power consumed by complete computers with a configuration perfectly defined by the brand and model of the computer (e.g. HP, DELL). This power consumption data is provided by the manufacturer of the computer.

Knowing the power consumption of each computer c, $W_c$, and the periods of time that it has been on, $T_{cp}$, the calculation of the energy consumption of the computers of the fleet during a given interval of time, $E_C$, is easy. Simply, for all computers, multiply the power consumption of each computer by the duration of periods in which it has been on inside the interval. Equation (1) represents this calculation.

$$E_C = \sum_{c=1}^{n} \sum_{p=1}^{m_c} \left[ W_c \cdot T_{cp} \right] \qquad (1)$$

Where $c$ is a computer, $p$ is an idle period, $n$ is the number of computers of the fleet, and $m_c$ is the number of turned on periods of the computer $c$.

### C. *Estimation of the total energy cost*

In the next phase, the cost of energy consumed by the computers of the fleet must be calculated. To carry out this calculation we need to know the cost of electricity. If the cost is always the same, the calculation is very easy, simply multiply the KWh consumed by the price of the KWh.

However, if the cost of electricity varies according to the time of day and/or with the day of the week, the cost of each period that each computer is on must be calculated independently. Furthermore, if during one on period, the cost of the electricity changes, that period must be divided into two consecutive periods, each with its own cost of electricity.

At the end of this phase, the administrator has an estimation of the total cost of the electricity consumed by the computers of the fleet in the interval of time under analysis.

### D. *Decision to purchase and deploy a PCPM system*

With all the estimations obtained in the previous phases, the administrator of the fleet must take the decision to purchase and deploy a PCPM system or abandon the project.

The administrator of the computer fleet will compare the cost of the power management system with the potential savings, estimating an ROI index (Return of Investment) or the amortization period of the system. After this period the power management system will generate real savings. It is important to remark that the decision taken will be optimistic because a power management system cannot reach the potential savings. Any system will only eliminate a high fraction of the total inactivity measured, but it will never eliminate all the inactivity because the productivity of the users would be severely affected.

## IV. Estimation of the Idle Timer Values to Control the Energy Consumption of the Computers

To achieve the potential energy savings, we must progressively turn off computer components and, eventually, the whole computer when a controller detects that the computer components and/or the whole computer have been inactive for a certain period of time.

In the analysis of the idle periods of a computer or group of computers, two features are important: 1) the instants when the periods happen and 2) the duration of the periods.

The instants (location in time) are important because if the periods are concentrated near a specific instant (for example, in the middle of the day, at lunch hour), it is an indication that most workers leave their computers around that instant. This fact can be useful to generate centralized orders for turning off and on the computers.

The duration of the idle periods is essential to estimate the waiting time that the power manager must wait from the detection of inactivity until taking the decision to turn off a component or the complete computer. The idle timer of the OS must be initialized with this waiting time, which is commonly called the value of the idle timer.

Fig. 3 shows the idle periods of a computer or a group of computers ordered by their length. It also includes the selection of three values for the idle timer. Selection 1 provides great energy savings. For each idle period, the computer remains on during a short time at the beginning of the period. However, this selection also causes a great inconvenience to users, because the user would have to resume the computer in most of the idle periods, reducing his productivity. Selection 3 provides little energy saving. The value of the idle timer is so high that the computer is suspended (hibernated) only in a few idle periods. Furthermore, in most of the periods in which the computer is suspended, it has been waiting turned on before the suspension for a long time. This selection causes little disturbance to users, because the user only has to resume the computer after a few idle periods of long duration. Selection 2 for the idle timer shows an intermediate behavior between the two previous selections.

To select the most appropriate value for the idle timer, $V_{IT}$, the manager has to assign a weight to two aspects:

The value should minimize the total inactivity of the computers, reducing the periods that the computers remain idle. This requires a low value of $V_{IT}$.

The value should also minimize the number of times that the user has to resume (turn on) the computer. This requires a high value of $V_{IT}$.

In order to establish a trade-off between these two aspects, they can be represented as a function of the value of the idle timer in the same graph. To obtain a general representation of the two aspects, two metrics are defined as percentages:

- Percentage of removed inactivity, *RI*: defined by eq. (2) represents the relation between the part of all idle periods which is removed by the expiration of the idle timer and the addition of all idle periods (the total inactivity time observed). In (2), $T_i$ represents the idle period *i* and *NT* the total number of idle periods of all computers used in the analysis.

$$RI = \frac{\sum_{i=1}^{NT} (T_i - V_{IT}) \; if \; T_i > V_{IT}}{\sum_{i=1}^{NT} T_i} 100 \qquad (2)$$

- Percentage of user satisfaction, *US*: defined by eq. (3) represents the relation between the number of idle periods that are less than or equal to the idle timer and the total number of periods observed.

$$US = \frac{\sum_{i=1}^{NT} 1 \; if \; (T_i \le V_{IT})}{NT} 100 \qquad (3)$$

Calculating these two metrics as a function of the value of the idle timer, for the idle periods represented in Fig. 3, the curves shown in Fig. 4 are obtained.

The value of the idle timer defined by the intersection of the curves can be considered a good starting value to find an appropriate value for $V_{IT}$, but it should not be automatically considered as the best optimal value for $V_{IT}$. Other factors could be considered to select an appropriate value for $V_{IT}$. In the example shown in Fig. 4, three ranges can be considered for $V_{IT}$.

When $V_{IT}$ is less than 4, the removed inactivity is greater than 60%, and therefore, the energy saving is high. However the user must resume the computer after all idle periods, making satisfaction very low, under 5%.

When $V_{IT}$ is greater than 12, the removed inactivity is always very low (less than 5%) and the energy saving is very low. The percentage of resumptions is low and does not decrease significantly when $V_{IT}$ increases, so the satisfaction of the user is high (greater than 80%). This range of $V_{IT}$ is useless for saving energy.

When $V_{IT}$ is between 4 and 12, the decrement of the removed inactivity and the increment of the user satisfaction are approximately linear. This is the useful range to select $V_{IT}$. The particular value to select depends of the importance given to the energy saving (low $V_{IT}$) against the satisfaction, and also the productivity, of the user (high $V_{IT}$).

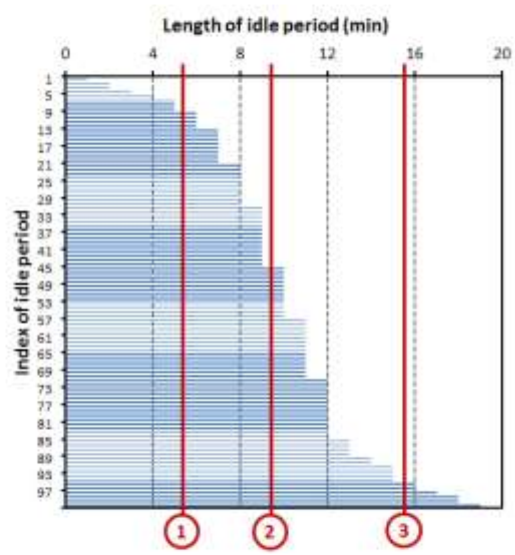


Figure 3. Representation of idle periods

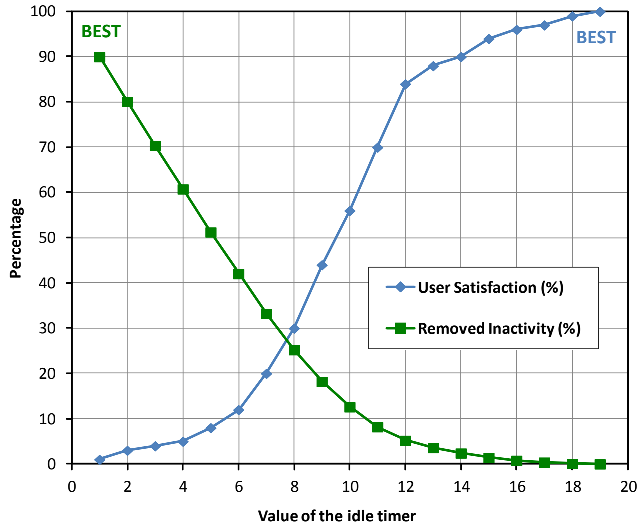


Figure 4. User satisfaction and removed inactivity as a function of the value of the idle timer

## V. Experimental Results in a Computer Fleet

A reduced scale experiment has been carried out in a Spanish company in order to study real lengths of inactivity intervals and the relation of removed inactivity and user satisfaction. The technique developed has been applied to 53 computers of different departments with very different configurations and usage patterns.

The agents measured the inactivity every 60 seconds, and therefore, all idle periods shorter than this value are not taken into account in the computations, due to the high uncertainty about their real length. In addition, it does not make sense to consider very short periods as real inactivity, because the user may be reading information on the screen.

In an interval of four weeks, 27833 idle periods were measured. Fig. 5 shows the idle periods when ordered by their length. The horizontal axis represents the length of the idle periods in seconds. In the vertical axis, the first idle periods (near index 1) are the shortest and the last idle periods (near index 27833) are the longest.

The maximum value of the horizontal axis that is shown in Fig. 5 is 86400 seconds, which cuts the longest idle periods at the bottom of the figure. The number of idle periods cut is 110 and the longest idle period is of 842723 seconds.

In Fig. 5, there is a large set of very short periods in the upper part and a small set of very long periods at the bottom. This shape of the ordered idle periods is characteristic in many computer systems. Therefore, a relatively small value of the idle timer will be required to reach a high percentage of the potential savings.

Fig. 6 shows the metrics RI and US. The removed inactivity, and therefore the power savings, decays exponentially when the value of the idle timer increases, while the user satisfaction grows very fast and then tends asymptotically to 100%. Fig. 7 shows a zoom in the curves for low values of the idle timer, in the range of 60 to 3600 seconds.

The curves RI and US intersect around 860 seconds. This value saves 85% of the energy with 85% of user satisfaction.

Considering an average consumption of 80Wh per computer and that the total inactivity time is 16014.98 hours, if 100% of energy consumed in the idle periods was saved, 1281.20 KWh would be saved. However, Fig. 7 shows that the maximum savings reachable, using the value of the idle timer of 860 seconds obtained from the intersection of curves RI and US, is 1089.02 KWh. These savings are for a period of four weeks and 53 computers. Extrapolating these data to one year and the 400 computers of the company, the total energy saved with this value of the idle timer would be 106.84 MWh.

In the surroundings of the intersection points, the slope of curve US is greater than the slope of curve RI. The slope is noticeably greater at the left of the intersection point than at the right. Therefore, the administrator can reduce the value of $V_{IT}$ to save a little more energy, drastically decreasing the user satisfaction. Similarly, the administrator can increase $V_{IT}$ to augment the satisfaction, saving just a little less energy. These conclusions are valid for the analyzed computer system, but in other organizations, the utilization patterns and the possibilities of saving energy may be different.

## VI. Conclusions

In this paper a centralized approach for power management of computers fleets has been proposed and implemented. The approach benefits from the existing power control mechanisms already embedded into the OS and uses three layers to develop the power management functions.

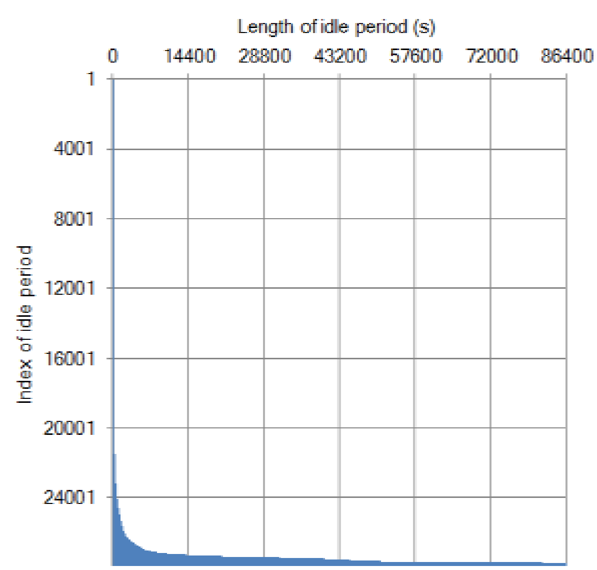


Figure 5. Idle periods measured in the computers of a company

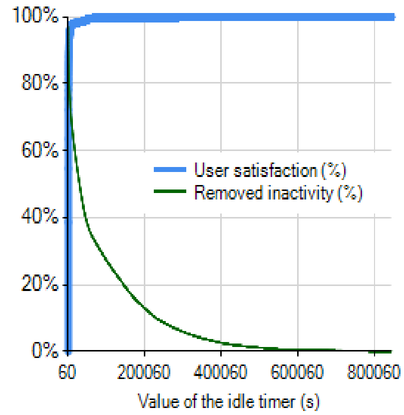


Figure 6. User satisfaction and removed inactivity as a function of the value of the idle timer

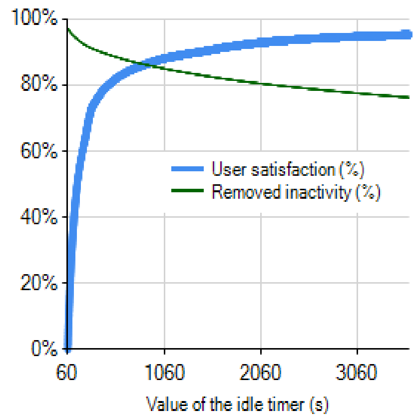


Figure 7. Zoom of curves of user satisfaction and removed inactivity for low values of the idle timer

A systematic technique has been developed to analyze the idle periods and the energy waste generated by these periods, before deploying a power management system to minimize the energy consumption. The correct estimation of the potential savings is very important to take the decision to use, or not, a power management system.

Finally, as the key control element of the proposed approach is the idle timer used by the OS to detect a period of inactivity and, when the timer expires, put the computer in suspension, a systematic technique to estimate the appropriate value for the idle timer has also been proposed.

To show the applicability of the technique, a distributed monitoring tool has been developed and used in several computers of the computer fleet of a company. With the idle periods measured over a period of four weeks, the methodology has been used to estimate an optimal value for the idle timer, reaching high percentages of energy savings and user satisfaction simultaneously.

To sum up, the technique proposed in this work can be used to improve the deployment of any power management system, implemented with commercial software packages or self-constructed software tools. Future work will be focused on developing techniques to adjust the value of the idle timer, daily or weekly, using the latest data measured.

## Acknowledgment

This work has been developed in the context of the project TIN2001-24903 of the Spanish National Plan for Research, Development and Innovation.

## References

[1] S. Murugesan, “Harnessing green IT: principles and practices,” IEEE IT Professional, vol. 10, iss. 1, January/February 2010, pp. 24-33.

[2] C.A. Webber, J.A. Roberson, M.C. McWhinney, R.E. Brown, M.J. Pinckard, and J.F. Busch, “After-hours power status of office equipment in the USA,” Energy, vol. 31, iss. 14, Nov. 2006, pp. 2823-2838, doi:10.1016/j.energy.2005.11.007.

[3] G. Procaccianti, L. Ardito, A. Vetro, and M. Morisio, “Energy Efficiency in the ICT - Profiling Power Consumption in Desktop Computer Systems,” in Energy Efficiency - The Innovative Ways for Smart Energy, The Future Towards Modern Utilities, M. Eissa, Ed. INTECH, 2012, pp. 353-372. http://dx.doi.org/10.5772/48237

[4] M. Chetty, A.J. Brush, B.R. Meyers and P. Johns, “It's not easy being green: understanding home computer power management,” Proc. SIGCHI Conference on Human Factors in Computing Systems, Apr. 2009, pp. 1033-1042.

[5] S. Buzzi, M. Mellia, “PoliSave: Efficient power management of Campus PCs,” Software, Proc. Telecommunications and Computer Networks (SoftCOM), Sep. 2010, pp.82-87.

[6] C. Piseth, A. Oxley, Low Tan Jung, “Dynamic PC power state transitioning based on a monitoring service,” Proc. Computer & Information Science (ICCIS), vol.1, Jun. 2012, pp.237-241.

[7] L. Benini, A. Bogliolo and G. De Micheli, “A survey of design techniques for system-level dynamic power management,” IEEE Trans. on Very Large Scale Integrated Systems, vol. 8, iss. 3, Jun. 2000, pp. 299-316.

[8] Y. Lu, L. Benini and G. De Micheli, “Operating-system directed power reduction,” Proc. International Symposium on Low Power Electronics and Design (ISLPED), Jul. 2000, pp. 37-42.

[9] A. De Salvo, A. De Silva, D. Benjamin, J. Blomer, P. Buncic, A. Harutyunyan, A. Undrus and Y. Yao, “Software installation and condition data distribution via CernVM File System in ATLAS”, Proc. International Conference on Computing in High Energy Physics (CHEP), May 2012.

[10] R. Ascierto, Selecting a PC Power Management Solution Vendor, Research Report of OVUM, December 16, 2011. http://ovum.com/research/selecting-a-pc-power-management-solution-vendor/

[11] 1E, “NightWatchman,” http://www.1e.com/it-efficiency/software/ nightwatchman-enterprise-pc-power-management/

[12] IBM, “Tivoli - EndPoint Manager for Power Management,” http://www-142.ibm.com/software/products/us/en/ ibmendpmanaforpowemana/

[13] Verdiem, “Surveyor,” http://www.verdiem.com/products/pc-power-manager

[14] SpiceWorks, “Power Manager,” http://www.spiceworks.com/free-pc-power-management-tool/

[15] MiserWare, “Granola Enterprise,” http://www.miserware.com

[16] Q. Jiang, H.-S. Xi, and B.-Q. Yin, “Adaptive optimisation of timeout policy for dynamic power management based on semi-Markov control processes,” IET Control Therory and Applications, vol. 4,iss. 10, 2010, pp. 1945–1958, doi: 10.1049/iet-cta.2009.0467.

[17] Y. Tan, W. Liu, Q. Qiu, “Adaptive power management using reinforcement learning,” Proc. IEEE/ACM International Conference on Computer-Aided Design (ICCAD 09), Nov. 2009, pp.461-467.