An Analysis on Energy Efficient System Design in Grid Computing

D. Ramesh[1] and A. Krishnan[2]

[1] Anna University of Technology, Tiruchirappalli
drameshphd@gmail.com
[2] K.S. Rangasamy College of Technology, Tiruchengode

Abstract. The primary focus of research in computing systems has been on the improvement of the effective designon system performance. In order to fulfil this objective, the performance has been steadily growing driven by more efficient system design and algorithms. The performance of grid is improved in many aspect based on various research direction of last few year. In grid computing, load sharing is the major research issue. In addition to the load sharing, at present, the power management is attracting current researchers. This paper further explains basic power management scheme in the general computing as well as grid computing. And this paper strongly performed an analysis on various categories of real time grid systems. The power consumption on various grid levels based on multiple volumes in the organization level is analysed. The conclusion is focused the future requirement of research direction in the energy efficient system design of grid computing.

Keywords: Grid Computing, Power Management, Energy Efficient System Design.

1 Introduction

The modern computing industry focus for grid computing as the requirement is over whelming. Although the performance per watt ratio has been constantly rising, the total power drawn by computing systems is hardly decreasing. Oppositely, it has been increasing every year that can be illustrated by the estimated average power use across three classes of servers presented in Table 1. The table describes the power consumption of every year from 2000 to 2006 on various grid levels and on various volumes. If this trend continues, the cost of the energy consumed by a server during its lifetime will exceed the hardware cost.

The problem is even worse for large-scale compute infrastructures, such as clusters and data centres. It was estimated that in 2006, IT infrastructures in the United States consumed about 61 billion kWh for the total electricity cost about 4.5 billiondollars. The estimated energy consumption is more than double from whatwas consumed by IT in 2000. Moreover, under current efficiency trends, theenergy consumption tends to double again by 2011, resulting in 7.4 billion dollarsannually.Energy consumption is not only determined by hardware efficiency, but it is alsodependent on the resource management system deployed on the infrastructure and theefficiency of applications running in the system.

N. Meghanathan et al. (Eds.): CCSIT 2012, Part I, LNICST 84, pp. 421–428, 2012.

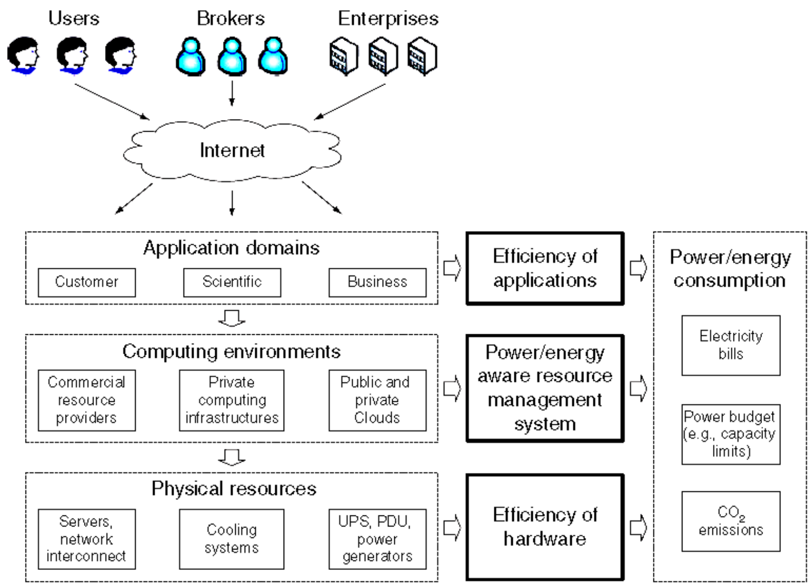

Fig. 1. General Overview of Energy Consumption in Computing Industry

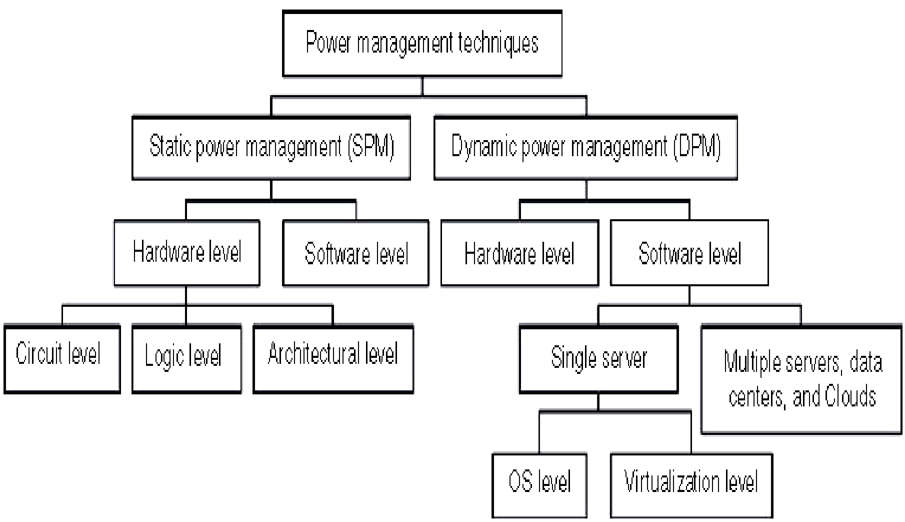

Fig. 2. Power Management

Table 1. Estimated Average Power Consumption / Grid Class

Grid Class	2000	2001	2002	2003	2004	2005	2006
Volume	186	193	200	207	213	219	225
Mid-Range	424	457	491	524	574	625	675
High-end	5534	5832	6130	6428	6973	7651	8163

The general overview of energy consumption in the various organization level of computing industry is explained in the figure 1. The power management with various factors of energy parameters is shown in the figure 2.

2 Related Work

There are a number of industry initiatives aiming at the development of standardized methods and techniques for the reduction of the energy consumption in computer environments. They include Climate Savers Computing Initiative (CSCI), Green Computing Impact Organization (GCIO), Green Electronics Council, The Green Grid, International Professional Practice Partnership (IP3), with membership of companies such as AMD, Dell, HP, IBM, Intel, Microsoft, Sun Microsystems, and VMware.

Energy-efficient resource management has been first introduced in the context of battery-powered mobile devices, where energy consumption has to be reduced in order to improve the battery lifetime. Although techniques developed for mobile devices can be applied or adapted for servers and data centres, this kind of systems requires specific methods.

To reduce the power and energy consumption in modern computing systems, as well as recent research works that deal with power and energy efficiency at the hardware and firmware, operating system (OS), virtualization, and data centre levels. The main objective of this work is to give an overview of the recent research advancements in energy-efficient computing, identify common characteristics, and classify the approaches. On the other hand, the aim is to show the level of development in the area and discuss open research challenges and direction for future work.

According to data provided by Intel Labs [1], the main part of power consumed bya server is accounted for the CPU, followed by the memory and losses due to thepower supply inefficiency. The data show that the CPU no longer dominatespower consumption by a server. This resulted from the continuous improvement ofthe CPU power efficiency and application of power-saving techniques (e.g., DVFS)that enable active low-power modes. In these modes, a CPU consumes a fraction ofthe total power, while preserving the ability to execute programs. As a result, currentdesktop and server CPUs can consume less than 30% of their peak power in lowactivitymodes, leading to dynamic power range of more than 70% of the peak power. In contrast, dynamic power ranges of all other server's components are muchnarrower: less than 50% for dynamic random access memory (DRAM), 25% for diskdrives, 15% for network switches, and negligible for other componentsPower supplies transform alternating current (AC) into direct current (DC) to feedserver's components. This transformation leads to significant power losses due to theinefficiency of the current technology.

Dhiman et al [2] have found that although regression models based on just CPU utilization are able to provide reasonable prediction accuracy for CPU-intensive workloads, they tend to be considerably inaccurate for prediction of power consumption caused by I/O- and memory-intensive applications. The authors have proposed a power modelling methodology based on Gaussian mixture models that predicts power consumption by a physical machine running multiple virtual machine

(VM) instances. The main reason of the power inefficiency in data centres is low average utilization of the resources.

3 Energy Management – Concepts and Theory

To understand power and energy management mechanisms, it is essential to clarifythe terminology. Electric current is the flow of electric charge measured in amperes.Amperes define the amount of electric charge transferred by a circuit per second.Power and energy can be defined in terms of work that a system performs. Power isthe rate at which the system performs the work, while energy is the total amount ofwork performed over a period of time. Power and energy are measured in watts (W)and watt-hour (Wh), respectively. Work is done at the rate of 1 W when 1 A istransferred through a potential difference of 1 V. A kilowatt-hour (kWh) is theamount of energy equivalent to a power of 1 kW (1000 W) being applied for onehour.

The difference between power and energy is very importantbecause a reduction of the power consumption does not always reduce the consumedenergy. For example, the power consumption can be decreased by lowering the CPUperformance. However, in this case, a programmay require longer time to complete itsexecution consuming the same amount of energy. On one hand, a reduction of the peakpower consumption results in decreased costs of the infrastructure provisioning, suchas costs associated with capacities of UPS, PDU, power generators, cooling system,and power distribution equipment. On the other hand, decreased energy consumptionleads to a reduction of the electricity bills. The energy consumption can be reducedtemporarily using dynamic power management (DPM) techniques or permanentlyapplying static power management (SPM). DPM utilizes the knowledge of the realtimeresource usage and application workloads to optimize the energy consumption.

However, it does not necessarily decrease the peak power consumption. In contrast,SPM includes the usage of highly efficient hardware equipment, such as CPUs, diskstorage, network devices, UPS, and power supplies. These structural changes usuallyreduce both the energy and peak power consumption.The main power consumption in complementary metal-oxide-semiconductor(CMOS) circuits comprises static and dynamic power.

The static power consumption,or leakage power, is caused by leakage currents that are present in any active circuit,independently of clock rates and usage scenarios. This static power is mainly determinedby the type of transistors and process technology. The reduction of the staticpower requires improvements of the low-level system design. Dynamic power consumption is created by circuit activity (i.e., transistorswitches, changes of values in registers, etc.) and depends mainly on a specificusage scenario, clock rates, and I/O activity. The sources of the dynamic powerconsumption are short-circuit current and switched capacitance. Short-circuit currentcauses only 10–15% of the total power consumption and so far no way has beenfound to reduce this value without compromising the performance. Switched capacitanceis the primary source of the dynamic power consumption.

The efficiency of power supplies depends ontheir load. They achieve the highest efficiency at loads within the range of 50–75%.However, most data centres normally create a load of 10–15% wasting the majority ofthe consumed electricity and leading to the average power losses of 60–80%. As aresult, power supplies consume at least 2% of the US electricity production. Moreefficient power supply design can save more than a half of the energy consumption.

Power efficiency and energy conservation are key design considerationsfor embedded systems. Various techniques have beenproposed over the years to reduce the energy consumption ofprocessor and memory subsystems as they are the two major contributorsof overall system energy dissipation. Dynamic voltagescaling (DVS) can be effectively used to reduce the power requirementquadratically while only slowing the processor performancelinearly. Recent studies show thatmemoryhierarchy, especially thecache subsystem, has become comparable to the processor in termsof energy consumption.

Dynamic cache reconfiguration (DCR)provides the ability to change cache configuration at run time sothat it can satisfy each application's unique requirement in termsof cache size, line size and associativity. By specializing the cachesubsystem, DCR is capable of improving cache energy efficiency aswell as overall performance

4 Energy Management – Grid Environment

4.1 Grid System Power Management

There are two main strategies for power reduction in *Grid*system: Dynamic Voltage/Frequency Scaling (DV/FS) and *Grid*number controlling: Vary-On Vary-Off (VOVF). DV/FS works by reducing the voltage and frequency, consequently saving energy at the cost of slower program execution. Researchers have developed various DV/FS scheduling algorithms to save energy under timing deadlines [3, 4]. Some researchers also utilized feedback control to dynamically adjust server frequency [5]. In these existing works, control variables can be either *Grid*frequency or application-level quality of service requirements [6-8]. VOVF is a major mechanism for power reduction applied in *Grid*clusters [5, 6, 7]. VOVF dynamically turns idle *Grid*off when the system experiences a light workload, and turns the appropriate *Grid*on when the system encounters a heavy workload. VOVF dramatically improves the system energy efficiency by reducing the idle *Grid*power consumption. Virtualization as a key strategy to reduce power consumption for application services is another way of VOVF. When applying virtualization, multiple virtual servers can be hosted on a smaller number of more powerful physical servers, using less electricity [9]. In [10], researchers demonstrated a method to efficiently manage the aggregate platform resources according to the guest virtual machine (VM) relative importance (Class-of-Service), for both the black-box and the VM-specific approach.

4.2 Cooling-Aware Power Management

Increasing computation capabilities in data centres has resulted in corresponding increases in rack and room power densities. How to cool these new higher-powered racks is a question that challenges all data centre managers [11]. There is several works attempting to reduce the energy consumption in the cooling sub-system. In [12], the authors explored the physics of heat transfer, and presented methods for integrating it into batch schedulers. It reduced the amount of heat recirculation in the data centre and improved the cooling subsystem efficiency. A mathematical scheduling problem is formulated in [13] to minimize the data centre cooling cost; they also provided two heuristic methods XInt-GA and XInt-SQP to solve the problem. In [14], researchers present a unified, coordinated, thermal-computational approach to the data centre energy management problem. Another group of researchers formulated an optimization problem to reduce the power consumption in servers and cooling system by selecting frequency level and cold air supply [15]. An integer linear programming was applied to solve the problem. Future work differs from these efforts: all of the above works focus on a single data centre, none of them considered multi-mirror services and their request distribution, and how they influence on the total cost for OSPs. In addition to these related work, there are few proposals which concentrates more on energy efficiency as well as effective system design. One such example in the computer network is energy efficient and reliable communication proposed by chandramohan et al [16-19], in which the author presented swarm intelligence based methodology.

4.3 Leveraging Variability Electricity Price in Reducing Cost

In [7], the researcher first considered the variable electricity prices for data centres and proposed a scheme to shut down the data centre when the electricity price is high. In [15],the author proposed a load dispatching strategy to reduce total electricity cost. An optimization problem was formulated in [8] to minimize the electricity cost in a multi-electricity-market environment. In [9], researchers considered the problem of capping the brown energy consumption of Internet services and interacting with the carbon market. However, existing efforts focus narrowly on electricity usage of the *Grid*subsystem without considering the dynamic behaviour of the cooling system and how to leverage it to reduce cost. Also, the network cost is not well studied in relation to reducing the total operational cost. The contribution of future work is that to provide a precise modelling of electricity usage in IDCs and provide energy-efficiency strategies in both *Grid*and cooling subsystems in addition to leveraging variability of electricity price. The net-work cost is also considered to obtain the optimal load dispatching among IDCs.

The motivation of this work is to help OSPs to conserve and manage their own electricity and network cost. Future work is directed to build by mainly exploring the following two opportunities:

1) Electricity market volatility: Electricity markets have been deregulated, attempting to create more economically desirable consumer markets. In the spot electricity markets, prices exhibit both location diversities and time diversities. If the electricity price is high in one region, dynamically route more requests to the regions with lower prices in order to save the total electricity cost.
2) Applying energy-efficiency strategies in IDCs: Two distinct sub-systems account for most of an IDC power draw: the *Grid*subsystem, which accounts for 56% of total power consumption of an IDC, the cooling subsystem, which accounts for about 30% of total power consumption.

These two subsystems dominate and their power draw can vary drastically with system workload. Applying energy-efficiency strategies in those two subsystems will contribute on both energy savings and electricity cost reductions in IDCs.

References

1. Minas, L., Ellison, B.: Energy Efficiency for Information Technology: How to Reduce PowerConsumption in Servers and Data Centers. Intel Press, USA (2009)
2. Dhiman, G., Mihic, K., Rosing, T.: A system for online power prediction in virtualizedenvironmentsusing gaussian mixture models. In: Proceedings of the 47th ACM/IEEE Design Automation Conference, Anaheim, CA, USA, pp. 807–812 (2010)
3. Vasic, M., Garcia, O., Oliver, J., Alou, P., Cobos, J.: A DVS system based on the trade-off between energy savings and execution time. In: 11th Workshop on Control and Modeling for Power Electronics, COMPEL 2008, pp. 1–6 (2008)
4. Pinheiro, E., Bianchini, R., Carrera, E., Heath, T.: Dynamic cluster reconfiguration for power and performance. In: Compilers and Operating Systems for Low Power (2001)
5. Wang, X., Chen, M.: Cluster-level feedback power control for performance optimization. In: Proc. of Symposium on High-Performance Computer Architecture (2008)
6. Sharma, V., Thomas, A., Abdelzaher, T., Skadron, K., Lu, Z.: Power-aware QoS management in web servers. In: 24th Annual IEEE Real-Time Systems Symposium (2003)
7. Dovrolis, C., Stiliadis, D., Ramanathan, P.: Proportional differentiated services: delay differentiation and packet scheduling. In: Proceedings of the Conference on Applications (1999)
8. Sharma, R., Bash, C., Patel, C., Friedrich, R.: Balance of power: dynamic thermal management for Internet data centers. IEEE Internet Computing (2005)
9. Murugesan, S.: Harnessing green it: principles and practices. IT Professional (2008)
10. Kesavan, M., Ranadive, A., Gavrilovska, A., Schwan, K.: Active CoordinaTion (ACT) – toward effectively managing virtualized multicore clouds. In: 2008 IEEE International Conference on Cluster Computing (2008)
11. Patterson, M.K., Fenwick, D.: The State of Data Center Cooling. A review of current air and liquid cooling solutions, Intel White Paper (2008)
12. Moore, J., Chase, J., Ranganathan, P., Sharma, R.: Abstract making scheduling cool: temperature-aware workload placement in data centers. In: Usenix Annual Technical Conference (2005)
13. Tang, Q., Gupta, S., Varsamopoulos, G.: Energy-efficient thermal-aware task scheduling for homogeneous high-performance computing data centers: a cyber-physical approach. IEEE Transactions on Parallel and Distributed Systems 19, 1458–1472 (2008)
14. Banerjee, A., Mukherjee, T., Varsamopoulos, G., Gupta, S.K.S.: Cooling-aware and thermal-aware workload placement for green HPC data centers. In: Proceedings of the International Conference on Green Computing, GREEN-COMP 2010, pp. 245–256. IEEE Computer Society, Washington, DC, USA (2010)
15. Pakbaznia, E., Pedram, M.: Minimizing data center cooling and server power costs. In: Proceedings of the 14th ACM/IEEE International Symposium on Low Power Electronics and Design, ISLPED 2009, pp. 145–150. ACM, New York (2009)
16. Chandra Mohan, B., Baskaran, R.: Reliable Transmission for Network Centric Military Networks. European Journal of Scientific Research 50(4), 564–574 (2011)

17. Chandra Mohan, B., Baskaran, R.: Survey on Recent Research and Implementation of Ant Colony Optimization in Various Engineering Applications. International Journal in Computational Intelligent Systems 7(4) (2011)
18. Chandra Mohan, B., Baskaran, R.: Energy Aware and Energy Efficient Routing Protocol for Adhoc Network Using Restructured Artificial Bee Colony System. In: Mantri, A., Nandi, S., Kumar, G., Kumar, S. (eds.) HPAGC 2011. CCIS, vol. 169, pp. 473–484. Springer, Heidelberg (2011)
19. Chandra Mohan, B., Sandeep, R., Sridharan, D.: A Data Mining approach for Predicting Reliable Path for Congestion Free Routing using Self-Motivated Neural Network. SCI, vol. 149, pp. 237–246. Springer, Heidelberg (2008)