ISSN 2395-695X (Print)
ISSN 2395-695X (Online)
International Journal of Advanced Research in Biology Engineering Science and Technology (IJARBEST)

Vol. 2, Special Issue 10, March 2016

HYBRID POWER BUDGETING FRAMEWORK FOR CLOUD ENVIRONMENT

D.Thamilarasi , Mr.S.Sivaprakash
Department of Computer Science and Engineering
K.S.R.College of Engineering

Abstract

Cloud computing is used to access computing resources owned and operated by a third-party provider. Cloud computing is an Internet-based computing to share resources, software and information. Data centers are used to share storage data with the users. Both transactional and long-running analytic computations are comprised into workloads. Power management strategies used in enterprise servers based on Dynamic Voltage and Frequency Scaling (DVFS). DVFS allows the server to transition the processor from high-power states to low-power states. The processors are assigned to sleep states such as deep sleep to reduce energy consumption. In deep sleep the server can be configured to use Direct Memory Access (DMA) to place incoming packets into memory buffers for processing in the active state.

Cloud Data centers are constructed with energy and cost parameters. Power budgeting methods are used to manage the power consumption in data center operations. Computing power and cooling power factors are used in the power budgeting framework. Cooling power is sufficient to extract the heat of the computing power. Self-Consistent Total Power Budgeting algorithm is applied to partition the total power budget among the cooling and computing infrastructure. Throughput predictor technique is used for servers with heterogeneous workload sets. Optimal computing power budgeting technique is used to identify the suitable power caps for the servers. Energy efficiency is improved by maximizing the System Normalized Performance (SNP).

Task scheduling is upgraded with power management constraints. Power budgeting technique is integrated with task scheduling algorithms to construct power optimization framework for data centers. System Normalized Performance (SNP) estimation is improved to measure task and data center operations. Power budgeting scheme is enhanced with data transmission delay factors.

1. Introduction

A key differentiating element of a successful information technology (IT) is its ability to become a true, valuable and economical contributor to cyber infrastructure. “Cloud” computing embraces cyber infrastructure and builds upon decades of research in virtualization, distributed computing, “grid computing”, utility computing and more recently, networking, web and software services. It implies service oriented architecture, reduced information technology Over head for the end-user, greater flexibility, reduced total cost of ownership, on demand services and many other things.

A powerful underlying and enabling concept is computing through service-oriented architectures (SOA) – delivery of an integrated and orchestrated suite of functions to an end-user through composition of both loosely and tightly coupled functions, or services – often network based. Related concepts are component-based system engineering, orchestration of different services through workflows and virtualization. The key to a SOA framework that supports workflows is componentization of its services, an ability to support a range of couplings among workflow building blocks, fault-tolerance in its data- and process-aware service-based delivery and an ability to audit processes, data and results, i.e., collect and use provenance information.

Component-based approach is characterized by reusability, substitutability, extensibility and scalability, customizability and composability. There are other characteristics that also are very important. Those include reliability and availability of the components and services, the cost of the services, security, total cost of ownership, economy of scale and so on. Many categories of components are distinguished in the context from differentiated and undifferentiated hardware, to general purpose and specialized software and applications, to real and virtual “images”, to environments, to no-root differentiated resources, to workflow-based environments and collections of services and so on.

An integrated view of service-based activities is provided by the concept of a workflow. An IT assisted workflow represents a series of structured activities and computations that arise in information-assisted problem solving. Workflows have been drawing enormous attention in the database and information systems research and development communities. Similarly, the scientific community has developed a number of problem solving environments, most of them as integrated solutions. Scientific workflows merge advances in these two areas to automate support for sophisticated scientific problem solving.

A workflow can be represented by a directed graph of data flows that connect loosely and tightly coupled processing components. It illustrates a Kepler-based implementation of a part of a fusion simulation workflow. In the context of "cloud computing", the key questions should be whether the underlying infrastructure is supportive of the work flow oriented view of the world. This includes on demand and advance-reservation-based access to individual and aggregated computational and other resources, autonomics, ability to group resources from potentially different "clouds" to deliver workflow results, appropriate level of security and privacy, etc.

2. Related Work

In [6], a control system for minimizing the power consumption in blade server enclosures is proposed. The power consumption of the blade server is minimized using three techniques: blade server consolidation, adjusting the speeds of the fans and assigning P-states to processors. The P-state assignment is based on a simple utilization-based technique. A processor is assigned a P-state so that the utilization is never higher than 80%. As discussed in the introduction, this technique is not effective in a power or performance constrained data center because the utilization of each core should be close to 100%. The other two techniques can be used in future work in combination with our assignment technique to reduce the power consumption.

In [5], it is shown that using an integrated approach to managing the cooling and computational resources in a data center is more efficient than if the two resources were managed independently. Their technique is similar to ours in that it trades-off power consumption with QoS. They trade-off power by deciding the amount of compute resources will be turned on at a compute node. In this paper, we extend that work in two directions. First, we consider a power constrained data center and a reward constrained data center. Second, we show how assigning P-states at the data center level results in improved performance.

The P-state assignment problem for optimizing some objective in a computer system has been studied widely (e.g., [7], [8], [9], [11]). The primary difference between our work and these studies is that our work considers the power consumed by the CRAC units in addition to the power consumed by compute nodes. The thermal-aware scheduling problem has been previously researched (e.g., [2], [3], [4]). Unlike our study, none of these papers consider P-state assignment.

Many other techniques to increase the energy efficiency of data centers exist. For example, the Open Compute Project started by Facebook proposes the following two techniques: 1) using a 480Velectrical distribution system to reduce energy loss and 2) reusing hot aisle air in the winter to heat offices. Another example proposed by the Sustainable Ecosystems Research Group at HP is to use water evaporation for cooling instead of using compressors. Many of these techniques can be used in conjunction with our technique to obtain further performance gains.

3. Energy Management for Clouds

Energy consumption and carbon emissions have become important global issues. Companies can reduce IT energy consumption by using a public on-demand cloud service, instead of their own servers or a data center (DC). A great deal of energy is consumed at the data centers that provide the infrastructure for cloud services. Reducing the energy use and emissions from data centers is important for energy-efficient companies. There are many existing technologies for reducing energy used by non-IT equipment. Air handling is one of the biggest consumers of energy in DCs [1]. Recently, some DCs have PUE (Power Usage Effectiveness) approaching 1.0. This means that almost all of the DC power consumption is due to IT equipment such as servers, storage, or networks. Therefore, it is becoming more important to reduce the power consumed by the IT equipment.

For reducing power consumption from IT equipment, virtual machine live migration has been studied for an efficient DC. Focusing on one DC cannot optimize the energy of the entire Cloud. Unified and flexible power management is even more important when assessing the overall status of multiple DCs. There is a need to track the changing states of multiple DCs in the Cloud. Hence an advanced tracking system is absolutely necessary to assess the status in Cloud Computing and to manage the total power consumption of multiple DCs.

Cloud-based energy management system is devised for Clouds. The system is developed on a sensor management Cloud that manages physical sensors as same as other IT resources, i.e. CPU, storage and network. The system has three main logical layers: the Sensor Management, Data Management and Application layers. The Sensor Management Layer provides access to physical sensors as Cloud resources similar to CPUs, disks, or memory resources. This layer makes it possible for multiple users to remotely manage and share physical sensors even if the sensor itself has only primary functions. The Data Management Layer provides a standardized data model for sensor data and customizes the data for user requirements [10]. The Application Layer provides sensor applications that feed sensor data to the Data Management Layer. Optimized Virtual Machine allocator is developed as an application. The optimized allocator analyzes sensor data that includes CPU workload histories and power usage and then recommends VM allocations to minimize the energy consumption in data centers.

The energy management system is develped in a Cloud and experimentally evaluated how much energy was saved by the optimized VM allocation. A Cloud environment is developed as a target for energy reduction with the system. The experimental environment emulates a real reference case with 12 servers and 116 VMs. The experiment showed a power savings of over 34% for the truly required 6 servers.

4. Power Budgeting for Data Centers

Datacenters and computing clusters with hundreds or thousands of servers consume excessive amounts of power, with large facilities consuming up to 20 MW for a total cost of $12 million per year. As a result, the total cost of ownership of data centers is dominated by power consumption, which constrains total performance and scalability. The power consumption of computing infrastructures and the power of cooling units are two major components of the total power of the data center, where the power consumption of the computer room air conditioning (CRAC) units depends on the power consumption of servers and the hot spots in the layout of the center. The power consumption of a facility at any moment of time must be capped below a maximum limit that is specified by the electric grid operators and the electrical current carrying capacity of its power cables.

As the scale and the utilization are significantly increased in current years, the cooling power could take up to 42 percent of the total power consumption of the data center. Thus, management for cooling power to avoid over cooling in data centers is more and more necessary. In server power budgeting, is that different workloads trigger different power consumption patterns and thus the power management settings that work for one set of workloads do not necessarily work for another set of workloads. One needs to find settings for each server that lead to a global optimal for the entire computing facility. The goal of this paper is to devise a new power budgeting method, where the total power budget is allocated among the servers and cooling equipment to maximize the system normalized performance (SNP), or equivalently minimizes the average runtime. We summarize our contributions as follows.

1) We propose a novel method to partition the total power budget between the computing servers and cooling units in a self-consistent way, where the cooling power meets the heat removal requirements for the computing power, which is allocated using an optimal power budgeting technique.

2) We propose a novel throughput predictor for servers with heterogeneous workload sets, where the measurements from the performance counters are used to estimate the change in the throughput as a function of the server power cap, on top of which, we can estimate the application normalized performance (ANP) beyond current power cap.

3) Leveraging the throughput predictor, we propose an optimal computing power budgeting technique that is inspired by methods for solving the well known knapsack problem. The budgeting technique identifies the optimal power caps for the servers, such that the total server power meets the computing budget and the system normalized performance is maximized.

4) We setup a realistic simulation environment for a data center with thousands of servers, where the power estimates for the servers are derived from real measurements on a server executing heterogeneous workload sets. We use computational fluid dynamics (CFDs) simulations to ensure accurate modeling of air flow and heat transfer within the center and use the CFD results to compute the cooling power requirement for a given power distribution. To speedup CFD simulation, we use an approximate approach based on heat cross-interference coefficient matrix. We experimentally demonstrate the advantages of our power budgeting

method and performance improvements, in terms of SNP, slowdown norm and unfairness, over previous approaches.

5. Issues on Power Budgeting Schemes

Cloud Data centers are constructed with energy and cost parameters. Power budgeting methods are used to manage the power consumption in data center operations. Computing power and cooling power factors are used in the power budgeting framework. Cooling power is sufficient to extract the heat of the computing power. Self-Consistent Total Power Budgeting algorithm is applied to partition the total power budget among the cooling and computing infrastructure. Throughput predictor technique is used for servers with heterogeneous workload sets. Optimal computing power budgeting technique is used to identify the suitable power caps for the servers. Energy efficiency is improved by maximizing the System Normalized Performance (SNP). The following issues are identified from the current power budgeting systems.

- Thermal aware task scheduling is not supported
- Task scheduling is not integrated with data center operations
- System Normalized Performance (SNP) is not focused on computational operations
- Data distribution latency under Virtual Machines is not considered

6. Hybrid Power Budgeting Framework

Task scheduling is upgraded with power management constraints. Power budgeting technique is integrated with task scheduling algorithms to construct power optimization framework for data centers. System Normalized Performance (SNP) estimation is improved to measure task and data center operations. Power budgeting scheme is enhanced with data transmission delay factors. Cloud data and resources are shared with power management mechanisms. Power constrained scheduling schemes are adapted to allocate the data. Center and computational resources for the users. Power optimization is carried out with internal communication factors. The system is divided into six major modules. They are Cloud Resources, Workload Management, Planning For Power Budgets, Power Budgeting On Data Centers, Task Scheduling With Power Optimization and Hybrid Power Budgeting Scheme

Cloud data centers and computational resources are managed under cloud resources module. Task submission operations are carried out under the workload management. Planning for power budgets module is designed to assign the power cost limits. Power budgeting is tuned for the data center operations. Power constrained resource allocation is carried out under task scheduling with power budget module. Hybrid power budget scheme is designed to manage data center and task scheduling operations with power utilization factors.

6.1. Cloud Resources

Cloud data centers are deployed to provide shared data values. Cloud servers are connected to construct the data centers. Computational resources are provided by the cloud resource providers. User account details are managed under the cloud server environment.

6.2. Workload Management

Cloud resources are provided for the workloads submitted by the users. Data access and computational operations are handled under the cloud resources. Cloud data center sends the requested data to the users. Computational tasks are executed under the cloud resource providers

.

6.3. Planning For Power Budgets

Power budgets are assigned for the data center operations. Economical and energy level factors are considered in the power budgeting plans. Power consumption levels are estimated with cooling power and computing power values. Power budgets are also allocated for the computational resources

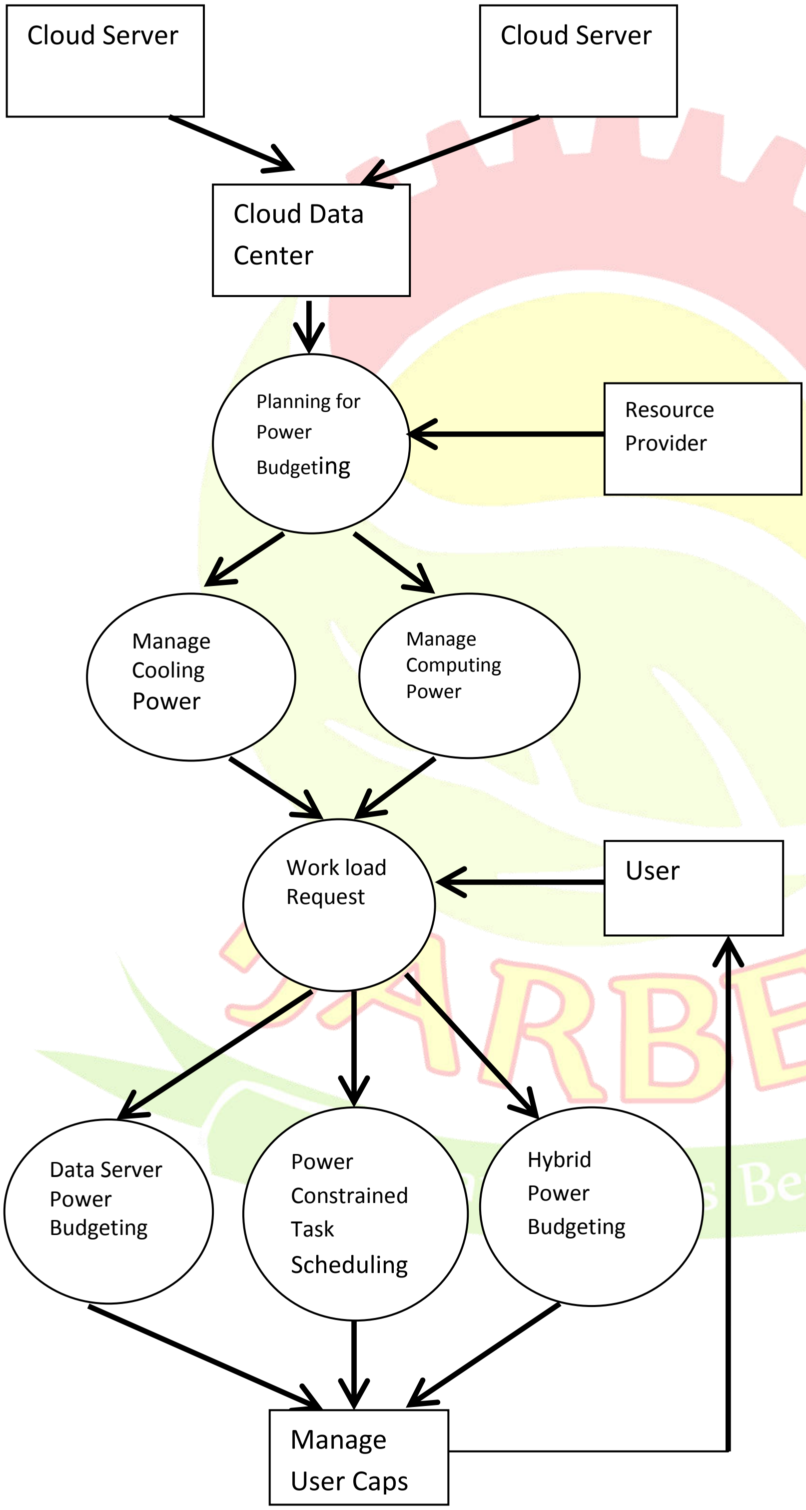
Cloud Server
Cloud Server
Cloud Data Center
Planning for Power Budgeting
Resource Provider
Manage Cooling Power
Manage Computing Power
Work load Request
User
Data Server Power Budgeting
Power Constrained Task Scheduling
Hybrid Power Budgeting
Manage User Caps

Fig. No: 6.1. Hybrid Power Budgeting Framework

6.4. Power Budgeting On Data Centers

Total power budgeting mechanism is adapted to manage cooling and computing infrastructures. Self-Consistent Total Power Budgeting algorithm is applied to divide the power budgeting estimation. Data server workloads are analyzed with throughput predictors. Server power caps are detected using the Optimal computing power budgeting technique.

6.5. Task Scheduling With Power Optimization

The power optimization process is designed to handle computational tasks. Computational resources are allocated with power budget criteria. Total and optimal power budgeting operations are tuned for the computational resource sharing environment. System Normalized Performance (SNP) estimation is integrated with the power management process.

6.6. Hybrid Power Budgeting Scheme

Computational task execution is integrated with the data sharing in cloud data centers. Hybrid power optimization algorithm is applied to manage the power consumption in data centers and resource providers. Task scheduling is combined with the data server selection process. Data transmission power consumption is also considered in the budgeting process.

7. Conclusion

Computing servers and data centers provides shared data and resources to the users. Power budgeting techniques are designed to share resources with energy controlled mechanism in data centers. Thermal aware task scheduling scheme is used to execute computational tasks with energy constraints. Hybrid power budgeting scheme integrates the task scheduling and data sharing operations in cloud environment. High data center utilization is achieved in data sharing process. Energy consumption is minimized in the task and data transmission process. The power budgeting scheme is achieves efficient application performance. The system maximizes the throughput in data distribution process.

References

[1] N. AbouGhazaleh, A. Ferreira, C. Rusu, R. Xu, F. Liberato, B. Childers, D. Mosse and R. Melhem, “Integrated CPU and l2 Cache Voltage Scaling Using Machine Learning,” SIGPLAN Notices, vol. 42, pp. 41-50, 2007.

[2] Z. Abbasi, G. Varsamopoulos and E.K.S. Gupta, “Thermal Aware Server Provisioning and Workload Distribution for Internet Data Centers,” Proc. 19th ACM Int’l Symp. High Performance Distributed Computing (HPDC’10), pp. 130-141, 2010.

[3] J. Choi, S. Govindan, J. Jeong, B. Urgaonkar and A. Sivasubramaniam, “Power Consumption Prediction and Power-Aware Packing in Consolidated Environments,” IEEE Trans. Computers, vol. 59, no. 12, pp. 1640-1654, Dec. 2010.

[4] E. Pakbaznia, M. Ghasemazar and M. Pedram, “Temperature- Aware Dynamic Resource Provisioning in a Power-Optimized Datacenter,” Proc. Conf. Design, Automation and Test in Europe 2010, pp. 124-129, 2010.

[5] L. Parolini, N. Tolia, B. Sinopoli and B. H. Krogh, “A Cyber-Physical Systems Approach to Energy Management in Data Centers,” Proc. 1st ACM/IEEE Int’l Conf. Cyber-Physical Systems, pp. 168-177, 2010.

[6] N. Tolia, Z. Wang, P. Ranganathan, C. Bash and M. Marwah, “Unified Thermal and Power Management in Server Enclosures,” Proc. ASME/Pacific Rim Technical Conf. Exhibition on Packaging and Integration of Electronic and Photonic Systems, MEMS, and NEMS (InterPACK), 10 pp., July 2009.

[7] J. Apodaca, D. Young, L. Briceno, J. Smith, S. Pasricha, A.A. Maciejewski, H.J. Siegel, S. Bahirat, B. Khemka, A. Ramirez and Y. Zou, “Stochastically Robust Static Resource Allocation for Energy Minimization with a Makespan Constraint in a Heterogeneous Computing Environment,” Proc. 9th ACS/IEEE Int’l Conf. Computer Systems and Applications (AICCSA‘11), p. 10, Dec. 2011.

[8] C. Xian, Y.-H. Lu and Z. Li, “Dynamic Voltage Scaling for Multitasking Real-Time Systems with Uncertain Execution Time,” IEEE Trans. Computer Aided Design of Integrated Circuits and Systems, vol. 27, no. 8, pp. 1467-1478, Aug. 2008.

[9] B.D. Young, J. Apodaca, L.D. Briceño, J. Smith, S. Pasricha, A.A. Maciejewski, H.J. Siegel, B. Khemka, S. Bahirat, A. Ramirez and Y. Zou, “Deadline and Energy Constrained Dynamic Resource Allocation in a Heterogeneous Computing Environment,” J. Supercomputing, vol. 63, no. 2, pp. 326-347, Feb. 2013.

[10] Abdulla M. Al-Qawasmeh, Sudeep Pasricha, Anthony A. Maciejewski and Howard Jay Siegel, “Power and Thermal-Aware Workload Allocation in Heterogeneous Data Centers”, IEEE Transactions On Computers, Vol. 64, No. 2, February 2015

[11] H. Yu, B. Veeravalli and Y. Ha, “Dynamic Scheduling of Imprecise- Computation Tasks in Maximizing QoS under Energy Constraints for Embedded Systems,” Proc. 2008 Asia and South Pacific Design Automation Conf. (ASPDAC‘08), pp. 452-455, Mar. 2008.