Energy Efficient VM Live Migration in Cloud Data Centers

[1]Jyothi Sekhar, [2]Getzi Jeba

[1, 2] Department of Information Technology, Karunya University,
Coimbatore-641114, Tamil Nadu, India

Abstract

Server consolidation is one of the techniques to reduce the overall energy consumption in the data centers besides various hardware and software features. Virtualization technology (VT) has made this possible with the help of various hypervisors which helps to create the virtual machines (VMs). During consolidation there are various parameters that needs to be thrown light into including performance, SLA (Service Level Agreement), CPU-I/O relation etc. Live migration of VM is the key to consolidating the servers without stopping them thereby with near-zero downtime for the systems. Excessive consolidation causes performance degradation which has severe impact on the QoS (Quality of Service) of the application that are running in that environment. VM allocation and placement are the main issues in consolidation. The performance- energy trade-offs have been discussed based on Greedy heuristics. Finally the simulation and live demonstration results are plotted.

Keywords: *Data centers, server consolidation, live migration, virtual machines, single threshold.*

1. Introduction

Cloud computing is the delivery of anything as a service (XaaS) rather than a product, whereby shared resources, software, and information are provided to computers and other devices as a pay-as-you-use commodity over a network (typically the Internet). Cloud computing is a marketing term for various technologies that provide computation, software and storage services that do not require end-user knowledge of the physical location and configuration of the system that delivers the services.

It is also a source of dynamically scalable and virtualized resources. The services of cloud are broadly divided into three categories: Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS) and Software-as-a-Service (SaaS). Amazon EC2, Microsoft Azure, Google App engine are the main providers of IaaS and PaaS. Google Apps, Salesforce etc are the leading providers of SaaS. A cloud service has three distinct characteristics that differentiate it from traditional hosting. It is sold on demand-- the very time of request; it is elastic -- a user can have as much or as little of a service as they want at any time; and the service is fully managed by the provider.

Therefore the hard-to-dispute benefits of cloud computing are highly reduced implementation and maintenance costs, increased mobility at a global scale of workforce, flexible and highly scalable infrastructures, increased availability of services at a very reasonable price, large database storage for entire data and many more. When Cloud becomes a source of anything and everything, there arise the question of security and privacy of data. There is a lot of challenges secure data storage, high internet speed, standardization etc. The cost of change in an ROI (Return on Investment) business case is less in Cloud Computing as the choices of Cloud services are more stable and more cost-effective than traditional ownership.

Virtualization technology renders a novel way to improve the power efficiency of data centers mainly via server consolidation. This enables several virtual machines (VMs) to reside in a single physical system making the utmost use of the system resources. Virtual machines use the resources of the physical system. In a sense, a data center can be keyed as a centralized repository both physical or virtual pertaining to a particular business; it is a restricted access area containing automated systems that constantly monitor server activity, Web traffic, and network performance. A flexible data center demands the following in order: consolidation, standardization, virtualization and utility.

The hypervisors support lifecycle management functions for the hosted VM images, and increasingly facilitate both offline and live migration for the VM [9]. Power management in data centers is crucial as it is not only limited to the energy cost but also the initial investment of cooling systems. Enormous amount of electrical energy is needed to run the data centers which results in the enormous expulsion of carbon dioxide into the atmosphere heightening the threat of global warming.

Datacenters consume 10 to 100 times more energy per square foot than typical office buildings [23]. According to Gartner, Cloud market opportunities in 2013 will be worth $150 billion [24]. Therefore the cloud service providers must adopt measures to preclude the performance degradation as well as high energy costs.

2. Related Works

Apparao et.al [12] presented a study on the impingement of consolidating several applications on a single server using xen as the hypervisor. Another study by Clark et.al [19] also dealt with virtual machine migration using Xen tabulating the various effects mainly the downtime and throws light into the practicality of live migration. Chase et.al [21] considered balancing the resource usage costs and its benefits during server consolidation. Nathuji and Schwan [22] proposed architecture of energy management system for virtualized data centers where the consolidations of VMs are done using live migration. According to Bing Wei [10] the VM with the largest memory will have the highest migration cost. He also proposed a workload adaptive algorithm for migration. Kyoon Hoon Kim et.al [23] proposed an algorithm for minimum price real-time VM provisioning. Hadi Goudarzi et.al [5] proposed an algorithm for minimizing the operational cost in the cloud computing system through SLA and effective VM placement. Norman Bobroff et.al [9] also proposes a management algorithm that provides probabilistic SLA guarantees. Reduction in the number of physical machines used was done using time series forecasting and bin packing heuristics. A hybrid approach that proactively allocates demanded resources for a predictable demand pattern and leverages a reactive controller to deal with excess demand was proposed by Anshul Gandhi et.al [8]. The system minimizes the number of SLA violations but they consider fixed CPU cycles which can cause wastage of CPU cycles.

3. Proposed Work

Server consolidation is the phenomena used to reduce the number of physical machines used to run the virtual machines. Through the process of virtualization a single physical machine can reside in it a number of virtual machines. This is done mainly to optimize the resource usage in these servers. In huge data centers, the number of powered-on servers should be reduced so that excess power consumption by poorly utilized servers can be precluded. Thus main concern is about over-utilized and under-utilized servers.

If a server is over-utilized, certain VMs from these servers are migrated to certain other servers within the same datacenter which can provide the required resources without much increase in power consumption. In the case of under-utilized servers, the VMs are moved to another efficient server and the former one is shut down thus reducing the power consumption.

The power model for servers can be given as:

$$P_{idle} + U_{CPU}\,(P_{max} - P_{idle})$$

where P_{idle} is the idle power of server, U_{CPU} is the CPU power, P_{max} is the maximum power of server.

When the server becomes over-utilized there are two main issues: which VM to migrate and to where it should be migrated.

3.1 VM Allocation

The allocation of virtual machines is done using the Greedy heuristics. The greedy heuristics sorts the VMs according to their size ie. the resource usage of the VMs and hence the power consumption contributing to the *SortingFactor.*

The VMs are sorted in the decreasing order of their sizes and the 'largest' VM is removed first. They are migrated to a host whose power consumption has the least increase $deltaP_{incr}$ after migration.

Greedy heuristics
Input : hostList, vmList
Output : BestVmAlloc in hostList

```
For all vm in vmList do
        vm.SortingFactor
end for
        vm.sortDecrOrder()
for all vm in vmList do
for all host in hostList do
/*The increase in power consumption of host*/
        deltaP_incr (host) = host.addingPower(u)
end for
/*Place the vm on the server with the least deltaP_incr*/
        host_minpower .add(vm)
end for
```

3.2 Cost of Migration

The cost of migration can be calculated for this system. Suppose $Vmdata_{size}$ is the size of data in virtual machine *Vm*, *M* bits are to be migrated; D_r is the available data rate for migration.

The entire migration duration can be calculated as

$D_{mig} = Vmdata_{size} / D_r$.

From this we can calculate the energy required for this migration. Suppose P_{mig} is the power used then

$Energy = P_{mig} * D_{mig}$

Suppose the migration duration is T seconds then

$Power = (P_{mig} * D_{mig}) / T$

Thus the cost of migration

$C_{mig}\ (Vm) = \{\ (P_{mig} * D_{mig})\ /\ T$; *if the migration has occurred in Vm,*
0; otherwise}

4. Simulation

4.1 Testbed for simulation

For the simulation of the single threshold policy, 20 virtual machines and 10 hosts were created inside a single datacenter. The host machine CPU allotted includes 1000, 2000 or 3000 MIPS with 1GB RAM and 1TB storage while the metrics for the virtual machine goes like 250, 500, 750 or 1000 MIPS, 128MB RAM and 1GB storage. The threshold value was set from 0.5 to 0.9. The workload given is about 150,000MI which equals 10mins of 250MIPS with 100% CPU utilization.

The graphs obtained on plotting the simulation result clearly show that the power consumption decreases when the utilization threshold increases.

4.2 Testbed for live migration of virtual machines using virt-manager

Migration of virtual machines was done live using the open source operating system ubuntu 12.04. Two physical systems each having core i5 processor, with 4GB RAM and 500GB hard disk were virtualized to include a single virtual machine in each of them. The hypervisor used was KVM/QEMU. The virtual machine manager called as the virt-manager was installed. The systems were connected using LAN and application protocol ssh used is. The nfs storage was created in a server and mounted on both the host machines. The utilization threshold was included in the open source software codes of the virt-manager in the source system.

5. Working Process

Workloads were introduced into the source virtual machine and are run repeatedly. The source host CPU consumption was observed using a shell script run separately in the source host. Virt-manager was run in both the machines. When the CPU utilization in the source machine exceeds the particular threshold the migration starts. The migration in the virt-manager is pre-copy migration as there is iterative copying and migration of the virtual machine CPU status. When the entire virtual machine is transferred to the destination, it continues to run there until stopped manually.

We assume that the destination system is able to provide the necessary resources for the migrating VM as it was idle. Here the performance degradation that may occur due to lack of resources is prevented and hereafter only the destination system needs to run. In the case of datacenters where there are many servers and VMs, this technique will certainly help in shutting down some of the under-utilized servers and preserve the QoS (Quality of Service) of the applications in the over-utilized servers.

6. Performance Evaluation

Simulation of the single threshold policy using greedy heuristics was done. The metric is the energy consumed by the servers during the simulation of single threshold policy.

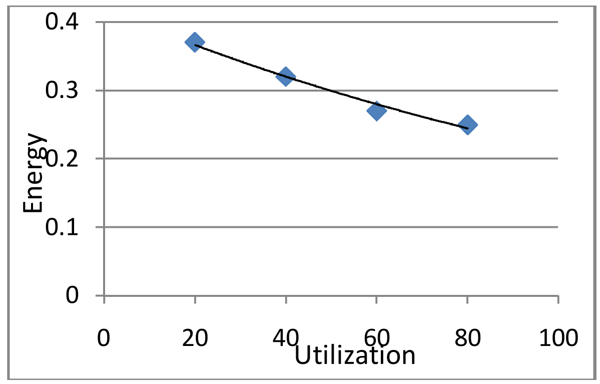

Fig. 1 Utilization Threshold-Energy graph – cloudsim

The single threshold policy implemented using cloudsim depicts that the increase in utilization threshold helps to reduce the energy consumed by the server. The live demonstration of energy consumption by the system during migration and with a single threshold was done on a small scale.

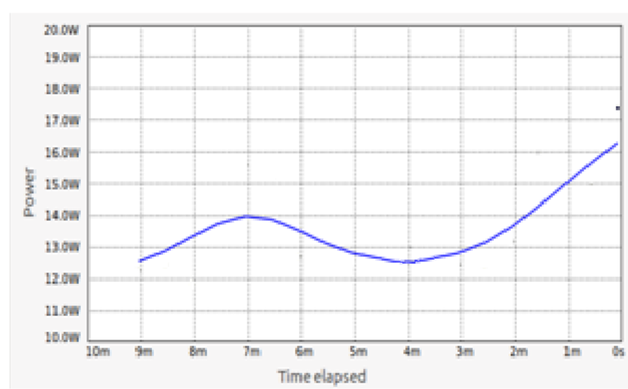

Fig. 2 Energy-Time graph – virt manager

The Energy-time graph was generated when the migration of virtual machine was taking place. The smaller curve in the beginning appears when the virtual machine along with an application within it starts running. It gradually stabilizes and when the threshold is reached it starts to migrate to the destination host which is indicated by the increase in power consumption.

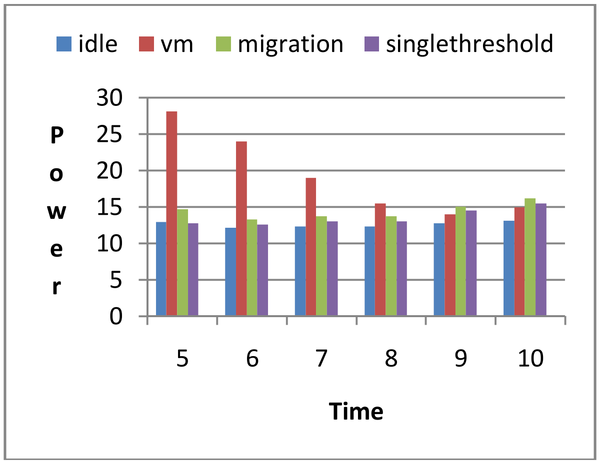

Fig. 3 Power-Time graph – virt-manager

In this graph, the power in watts across time in minutes are plotted to formulate the power consumption of the physical host at various stages. The idle system power consumption is almost constant throughout the graph. When a VM is started there is a tremendous increase in the power consumed which linearly decreases and becomes almost constant.

During migration the resource utilization always increases and therefore the power also increases but at a slower rate. The setting up of utilization threshold for a particular host system helps to reduce the power consumed by the entire system as the VM is migrated entirely to the destination and now only the destination system needs to run and henceforth shutting down the source machine will eventually reduce the total power consumption of the entire system.

7. Enhancement and conclusion

Live migration, though has its own overhead, help to consolidate the servers and thus reduce the number of physical machines that are powered-on and thereby reducing the power consumption of these systems. The quality of service is a very important factor that should not be compromised. At higher threshold values there occur slight SLA violations. But the greedy algorithm is found to be 3-5% better in energy-performance trade-offs than the previous heuristics used.

The splitting up of the workload among the VMs is to be experimented to reduce the degree of consolidation. The cost of migration can be reduced by selecting the VM with the least memory size. Our future work will concentrate on reducing the overhead of migration and also the cost of migration.

References

[1] Gergo Lovasz, Florian Miedermeier, Hermann de Meer, “Performance Tradeoffs of Energy-Aware Virtual Machine Consolidation” Cluster Computing; The Journal of Networks, Software Tools and Applications, 2012

[2] William Voorsluys, James Broberg, Rajkumar Buyya, Srikumar Venugopal, “Cost of Virtual Machine Live Migration in Clouds: A Performance Evaluation”, CloudCom '09 Proceedings of the 1st International Conference on Cloud Computing, 2009, pp 254-265

[3] Anshuman Khandual, “Performance Monitoring in Linux KVM Cloud Environment”, IEEE, 2012

[4] Young Choon Lee and Albert Y. Zomaya, “Energy Efficient Resource Allocation in Large Scale Distributed Systems” Ninth International Symposium on Distributed Computing and Applications to Businee, Engineering and Science, IEEE, 2010, pp 580-583

[5] Hadi Goudarzi, Mohammad Ghasemazar and Massoud Pedram, “SLA-based Optimization of Power and Migration Cost in Cloud Computing”, CCGRID '12 Proceedings of the 2012 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, pp 172-179

[6] Lionel Eyraud-Dubois, Hubert Larcheveque, “Optimizing Resource Allocation while handling SLA Violations in Cloud Computing Platforms”, 27th IEEE International Parallel & Distributed Processing Symposium, 2013

[7] Richa Sinha, Nidhi Purohit, Hiteishi Diwanji, “Energy Efficient Dynamic Integration of Thresholds for Migration at Cloud Data Centers”, International Journal of Computer Applications 2011, pp 44-48

[8] Anshul Gandhi, Yuan Chen, Daniel Gmach, Martin Arlitt, Manish Marwah, “Minimizing Data Center SLA Violations and Power Consumption via Hybrid Resource Provisioning”, IGCC '11 Proceedings of the 2011 International Green Computing Conference and Workshops, IEEE 2010,pp 1-8

[9] Norman Bobroff, Andrzej Kochut, Kirk Beaty, “Dynamic Placement of Virtual Machines for Managing SLA Violations” International Symposium on Integrated Network Management, 2007. IM '07. 10th IFIP/IEEE , 2007, pp 119-128

[10] Bing Wei, “A Novel Energy Optimized and Workload Adaptive Modeling for Live Migration”, International Journal of Machine Learning and Computing, 2012, pp 162-167

[11] Anton Beloglazov and Rajkumar Buyya, “Optimal Online Deterministic Algorithms and Adaptive Heuristics for Energy and Performance Efficient Dynamic Consolidation of Virtual Machines in Cloud Data Centers”, Concurrency and Comptutation: Practice and Experience, 2011, pp 1-24

[12] Apparao P Iyer R, Zhang X, Newell D, Adelmeyer T, Characterization & analysis of a server consolidation benchmark. In: VEE '08: Proceedings of the fourth ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, New York, NY, USA, ACM (2008) 21-30. 13. Qi Zhang et.al, “Cloud Computing: State-of-the-art and research challenges”, Springer 2010

[13] Anton Belaglozav, R. Buyya, “Adaptive Threshold-Based Approach for Energy-Efficient Consolidation of Virtual Machines in Cloud Data Centers”, Proceedings of the 8th

International Workshop on Middleware for Grids, Clouds and e-Science, ACM 2010
[14] Anton Beloglazov, "Energy Efficient Resource management in Virtualized Cloud Data Centers", In Proc. of the 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing, IEEE Press, 2010, pp.826-831.
[15] Anton Belaglozav et.al, "Energy-aware resource allocation heuristics for efficient management of data centers for Cloud computing", Future Generation Computer Systems, Volume 28, Issue 5, ACM 2012
[16] Aziz Murtazaev, Sangyoon Oh, "Sercon: Server Consolidation Algorithm using Live Migration of Virtual machines for Green Computing", IETE Technical Review, Vol 28, Issue 3, 2011.
[17] Pablo Graubner, Matthias Schmidt, Bernd Freisleben, "Energy-efficient Management of Virtual Machines in Eucalyptus", 4th International Conference on Cloud Computing, IEEE 2011.
[18] Ching-Chi Lin, Pangfeng Liu, Jan-Jan Wu, "Energy-Aware Virtual Machine Dynamic Provision and Scheduling for Cloud Computing", 4th International Conference on Cloud Computing, IEEE 2011
[19] Clark. C, Fraser. K, Hand, S. Hansen J.G, Jul E, Limpach C, Pratt I, War eld, A.: Live migration of virtual machines. In: NSDI'05: Proceedings of the 2nd Conference on Symposium on Networked Systems Design & Implementation, Berkeley, CA, USA, USENIX Association (2005) 273-286.
[20] Cloud Computing for Dummies, Wiley Publishing, Inc.
[21] J.S. Chase, D.C. Anderson, P.N. Thakar, A.M. Vahdat, R.P. Doyle, Managing energy and server resources in hosting centers, in: Proceedings of the 18th ACM Symposium on Operating Systems Principles, ACM, New York, NY, USA, 2001, pp. 103–116.
[22] R. Nathuji and K. Schwan. Virtualpower: Coordinated power management in virtualized enterprise systems, ACM SIGOPS Operating Systems Review, pp. 265–278, 2007.
[23] Kyong Hoon Kim, Anton Beloglazov, Rajkumar Buyya, "Power-aware Provisioning of Cloud resources for Real-time Services", Proceedings of the 7th International Workshop on Middleware for Grids, Clouds and e-Science, ACM 2009.
[24] Pablo Graubner, Matthias Schmidt, Bernd Freisleben, "Energy-efficient Management of Virtual Machines in Eucalyptus", 4th International Conference on Cloud Computing, IEEE 2011.
[25] Anton Beloglazov, Jemal Abawajy, Rajkumar Buyya, "Energy Aware Resource Allocation Heuristics for Efficient Management of Datacenters for Cloud Computing", Future Generation Computer Systems, Elsevier, 2012, pp 755-768.

Jyothi Sekhar has done Btech in Information Technology from the Institute of Engineering and Technology, Kerala in 2010 and is now doing her final year Mtech in Network and Internet Engineering from Karunya University, Coimbatore. The main areas of interest include energy efficiency in cloud data centers and virtual machine live migration. The main areas of interest include energy efficiency in cloud data centers and virtual machine live migration.

Getzi Jeba is currently working as Assistant Professor in Karunya University, Coimbatore. She had completed her Mtech in Network and Internet Engineering from the same university in 2009. She has her area of interest in routing in computer networks and cloud computing. She is pursuing her PhD in Energy Efficient live migration techniques. She is a member of the Computer Society of India.