Enabling Energy Efficient Resource Management in Cloud Systems

Anup H. Gade
PG Scholar, RGPV
Shri Ram Institute of Tech.
Madhotal, Jabalpur

Sandeep Sahu
Assistant Professor
Shri Ram Institute of Tech.
Madhotal, Jabalpur

ABSTRACT
Rapid growth of the demand for computational power by scientific, business and web-applications has led to the creation of large-scale data centers consuming enormous amounts of electrical power. Generally it has been proposed that energy efficient resource management system for virtualized Cloud data centers that reduces operational costs and provides required Quality of Service (QoS). Energy savings are achieved by continuous consolidation of VMs according to current utilization of resources, virtual network topologies established between VMs and thermal state of computing nodes. From so long improving performance has been the primary motive of computing systems industry whose growth is steadily fuelled by resource hungry applications. A typical data centre with 1000 racks needs 10 MW of power to operate, which is sufficient enough to power a small city. With increasing shift in migrating services to cloud the power demands of the backend data centers will continue to grow. Apart from consuming enormous power, 'Cloud's also increase global carbon foot print. It is mentioned in the popular reports that industry generates about 2% of the total global CO2 emissions, which is equivalent to the aviation industry. The primary objectives of this research are to create Simplified Energy models for physical systems and virtual machines by monitoring their resource usage.

Keywords
Energy efficient resource management, virtualized cloud data centers, Virtual Machine (VM), Quality of Service (QoS), computational power, simplified energy models.

1. INTRODUCTION
We Rapid growth of the demand for computational power by scientific, business and web-applications has led to the creation of large-scale data centers consuming enormous amounts of electrical power. Generally it has been proposed that energy efficient resource management system for virtualized Cloud data centers that reduces operational costs and provides required Quality of Service (QoS). Energy savings are achieved by continuous consolidation of VMs according to current utilization of resources, virtual network topologies established between VMs and thermal state of computing nodes. It has presented first results of simulation-driven evaluation of heuristics for dynamic reallocation of VMs using live migration according to current requirements for CPU performance. The results of the system show that the proposed technique brings substantial energy savings, while ensuring reliable QoS. This justifies further investigation and development of the proposed resource management system.

2. REVIEW OF LITERATURE
An authentic paper about Cloud computing [6] published from researchers in University of California, Berkeley paints a clearer picture of Cloud Systems along with its revenue models. This section summarizes some of the important points regarding Cloud Systems internals. Cloud Systems interface with data centers which can be spread over wide geographical regions and provide the ability to use computing resources on demand. This on demand resource usage is coupled with a pay as you go business model much like an electricity grid. Clouds systems, which are open for general public, are termed as public clouds and those internal to an organization are called private clouds. Commercial cloud service providers like Amazon, Google, Microsoft etc. provide APIs to interact and customize physical resources in their public clouds. This model of utility computing has been tremendous success owing to clouds elastic scalability and savings in setting up own infrastructure. The risk of under or over provisioning data centre resources for organizations are completely migrated to Cloud Computing provide which performs statistical multiplexing of compute resources over a much larger set of users. This idea is proved to be beneficial to both cloud vendors and users of clouds. Various services offered by cloud vendors can be broadly classified into the following three services as SaaS, PaaS and IaaS.

2.1 VM Migration
Since Virtual Machines run in a sandboxed environment using virtualized resources their execution state can be persisted to disk which are referred to as virtual machine snap-shots and are very useful for scenarios like disaster recovery, sand-boxing etc. Snapshots of such VMs can be made accessible to other nodes in the cloud if the VM resources are used from shared network storage. A simple VM Migration can be performed by saving snapshot of a running VM onto network storage from the source node and then requesting the destination node to resume VM state. This idea also requires that VM configuration settings are in sync on all nodes in the Cloud which is possible by using shared network storage. The alternative to using network storage is to make use of high available disk mirroring solution like DRBD [11]. The potential problem with using a snapshot based migration is if there are clients connected to applications like web, database, email servers etc running on the VM, they will be disconnected during the snapshot migration process and all of these clients have to be restarted after VM resumes its operation on the destination node. This operation causes downtime to end users and which is not acceptable in typical enterprise deployments.

Live migration [12] uses an improved technique using which virtual servers can be mi-grated between different physical machines without loss of clients/application connectivity. This concept iteratively copies memory pages of running VM to destination node, before pausing VM brie y for a short period of time (as low as 60ms) to copy remain-ing pages and signals the destination node to resume execution. Please make a note that live migration requires all nodes in the Cloud to have access to virtual machine resources (hard-disks, configuration settings etc.) on shared network storage.

2.2 Power & Energy Models

Rajkumar Buyya et al have carried out a detailed survey of energy efficient cloud systems in [15] which also points to some important work done in modeling physical and virtual machines. Summary of these models follow beginning with definitions of these terms.

Power is deified as the rate at which the system performs work. Electric Power is deified as the rate at which electrical energy is transferred by an electric circuit. Power is represented by Equation 1.2.1 and its SI units are Watts. Energy is total amount of work performed over a period of time and is given by Equation 1.2.2 Energy is measured in watt-hour (Wh) or kilowatt-hour(kWh) etc.

$$P = V\,I = I^2R \qquad (1.2.1)$$

where V is the Voltage, I is the Current owing and R is the Resistance

$$E = PT \qquad (1.2.2)$$

where P is the Power obtained from Equation 1.2.1 and T is the Time period.

The difference between power and energy is important here. Energy has a time component associated where as Power is an instantaneous quantity. In other words Energy is the accumulated power over a time period. Decreased energy consumption will lead to smaller electricity bills. Information of the peak power utilized by individual components in a data center is required during infrastructure provisioning (UPS, PDU, generators, cooling systems etc).

2.3 Power model for Physical Systems

System modeling in general can be classified into practical models or simulated models depending upon how they are derived. Values produced from practical models tend to have less deviation from the original value since they are pro led using special external hardware. Simulated models are simple to construct and can also be referred to as black-box models. These models are used to generate a range of reasonable values for experimental use and don't necessarily have to produce accurate predictions. To summaries from previous section Energy models sum up power consumed over a given time period and thus are dependent on Power models.

2.4 Power model for Virtual Machines

Pioneering work in this area is done by Microsoft research with 'Joule meter'[17] project which approximates power consumed by VM's initially using standard power models and constantly represent them as time progresses using machine learning techniques. This developed software works only on windows platforms and can also be used to estimate energy consumption of physical nodes. There is a provision for calibrating initial readings using external power meters which will produce ne grained results. The power model proposed in this project is given by

$$E_{sys} = cpu + mem + io + disk \qquad (1.4.1)$$

where values correspond to the utilization rates and all other parameters are obtained using linear regression. In case of memory $_{mem}$ corresponds to last level cache miss rate.

3. RESEARCH METHODOLOGY

This section deals with the implementation details and algorithms used for developing cloud centre software. This project makes use of several functionalities of these external components and builds a cloud control centre on top of them. The cloud system developed in this project has energy policies which work perfectly only with homogeneous VM workloads. This means that if the collective load on the VM's exceeds the host’s physical limitation, performances will dramatically sure. To mitigate risks imposed by this dynamic behavior in clouds robust performance triggered migration policies are to be designed. The reason for choosing cloudsim framework is to simulate and understand how energy saving policies work with non-homogeneous clouds. Analysis of energy saved in various CloudSim scenarios by using a single threshold energy policy in comparison with non power aware policy are explained below.

2.1 Simulation Design

The following steps are carried out in this simulation experiment

1. VM's are created randomly with either 250, 500, 750 or 1000 MIPS rating to cater the needs of the cloudlets of the same MIPS requirement.
2. The host nodes will be created randomly with either 1000, 2000 or 3000 MIPS rating.
3. The threshold value used for Single Threshold policy is 80%. Since this is a CPU driven simulation, this means that if the CPU utilization of a node is more than the set threshold, the cloud centre automatically consolidates load distribution by migrating some running VM's to another node,
4. After the initial conditions are configured, simulations for both Non Power Aware and Single Threshold policies are run with
    1. 10 Nodes and 20 VM's
    2. 100 Nodes and 200 VM's
    3. 500 Nodes and 1500 VM's
    4. 1000 Nodes and 3000 VM's

3.2 Discussion

The average SLA violation in Table 5.1 varies between 10%-14% when using single threshold policy. This SLA violation occurs because of over-subscription of VM's into as less physical machines as possible. When load on these VM's increase their host node cannot guarantee requested MIPS performance factor and thus VM performance will get aected. When cloud centre observes such a scenario it consolidates VM's in the cloud as de ned by the threshold value. The advantage with using this simulator is since it is able to model real cloud deployments which will have non-homogeneous VM workloads, it would be possible to obtain accurate energy saving estimates.

Policy	Energy Consumed(kWh)	No of Migrations	Avg SLA violation
NPA	0.91	-	-
ST(80%)	0.35	37	10.00%
10 Nodes-20 VM's, Energy Saving: 61.53%			
NPA	9.1	-	-
ST(80%)	1.84	446	13.44%

100 Nodes-200 VM's, Energy Saving: 79.78%			
NPA	46.35	-	-
ST(80%)	17.93	3304	11.74%
500 Nodes-1500 VM's, Energy Saving: 61.31%			
NPA	92.36	-	-
ST(80%)	36.44	6620	11.42%
1000 Nodes-3000 VM's, Energy Saving: 60.54%			

Table 2.2.1: Analysis of energy consumption in simulated private clouds

3.3 Energy Savings in Cloud Control Centre

This section explains the experiments conducted on real test-bed (private cloud). Please make a note that due to lack of external power meters the power consumption values are estimated using power models. Initially the experiment is started with 10 VM's and 5 nodes using load balanced start approach. These policies effectively load balances all VM's by starting 2 VM's on each node. This is considered as a base case and will be compared with 'Greedy Energy Saving Migration' and 'Energy Saving with Load Balancing' policies. Repeat the above procedure for one VM on each node and mainly with only 4 VM's running on 4 nodes. A power saving greater than 45% is reported in each of these cases. Around 45% less power consumption is reported when started with 2 VM's on each node. There was even lesser power consumption of 61% when the experiment is started with one VM on each node. The last experiment estimates the maximum achievable power saving with this private cloud by initially starting 4 VM's across 4 nodes and then requesting any of the migration policy to optimize their placement for energy saving. Here 4 correspond to the maximum permitted VM's on any node which is defined by the physical memory present on that node (factoring out storage space which resides on NFS). A staggering 75% less power consumption is reported by consolidating VM's in this private cloud. In Table 2.3.1P_{base} and P_{io} are factored out. An experiment with treating these values as a constants (P_{base} + P_{io} = 60W) is shown below in Table 2.3.2

Policy	Placement of nodes	Power consump-tion in cloud	Power saving
Load Balanced Start	2 - 2 - 2 - 2 - 2	1150.92 W	- (base case)
Greedy Energy Saving Migration	0 - 0 - 2 - 4 - 4	627.63 W	45.46%
Load Balanced En-ergy Saving Migra-Tion	0 - 0 - 3 - 3 - 4	612.22 W	46.8%
Load Balanced Start	1 - 1 - 1 - 1 - 1	966.72 W	- (base case)
Greedy Energy Saving	0 - 0 - 0 - 1 - 4	369.32 W	61.79%
Load Balanced En-ergy Saving Migra-Tion	0 - 0 - 0 - 2 - 3	369.20 W	61.81%
Load Balanced Start	1 - 1 - 1 - 1 - 0	747.38 W	- (base case)
Any Energy Saving Migration Policy	0 - 0 - 0 - 4 - 0	184.83 W	75.26%

Table 2.3.1: Energy Saving Analysis in a private cloud with 5 nodes and 10 VM's

Policy	Placement of nodes	Power consump-tion in cloud	Power saving
Load Balanced Start	1 - 1 - 2 - 2	999.36 W	- (base case)
Greedy Energy Saving Migration	0 - 0 - 2 - 4	558.17 W	44.14%%
Load Balanced En-ergy Saving Migra-Tion	0 - 0 - 3 - 3	536.12 W	46.35%

Table 2.3.2: Energy Saving Analysis in a private cloud with 4 nodes and 6 VM's

3.4 Results and Discussion

The results are represented by using plots generated from rrd tool in each case. Latency and Packet loss % of a VM before migration is in Figure 5.2 and with migration is in Figure 5.3.

In base case with no migration there is no packet loss and the latency is 1ms. While the VM is in migration, It has been observed from several runs that there has been a 6-9ms network latency. A 25% packet loss is also reported during this time. The migration process typically takes 250ms every time. These values are strictly dependent on factors like available network bandwidth, load on the physical system which hosts the VM under test etc. Please make a note that the values reported here are only for this private cloud and might be significantly smaller with faster test bed's.

4. CONCLUSION & FUTURE

Energy efficient systems research has gained tremendous interest in the recent times in order to solve increasing power requirements issues in large Clouds and Data centers. Apart from saving huge power bills these energy efficient systems also helps reduce global carbon footprint by supporting Green

Computing initiative.The work presented in this document is a direction towards creating energy efficient cloud system. With the energy saving cloud polices designed in this project power savings of upto 75% are possible in comparison with default load balanced approaches. Energy savings of greater than 50% are achievable using single threshold policy according to CloudSim, a non-homogeneous private cloud policy simulator.Developing energy efficient systems takes a multidisciplinary approach and since this project is in its infancy there is scope for a lot of future enhancements some of which are mentioned below.

3.1 Future Work

Integrate with Energy monitoring framework developed by YiYu and develop an energy aware cloud. Develop robust Energy triggered migration policies. This work will require concepts from control systems and experimental analysis and is targeted for non homogeneous VM workloads. Modify architecture such that keep alive control plane is separate from virtual Box commands control plane. This change will delegate authority of periodically sending VM state information, power etc to compute nodes and will result in efficient use of agent based architecture. Apart from improvements in performance of cloud control software, this change would simplify troubleshooting issues at a later stage.

In order to scale this solution for larger clouds decentralization has to be added to the Cloud control centre architecture by sharing cloud centre responsibilities among multiple Controller nodes. Apart from scalability decentralization also comes with added benefit of improved resiliency from single point of failure. Evaluate use of cluster le system instead of NFS as network storage by analyzing I/O performance in both cases. Using clustered le systems will increase storage capacity for VM resources and eventually results in a scalable cloud system.

5. REFERENCES

[1] Prof. Saleem Bhatti Green cloud. andrews.ac.uk/node/1723#green_cloud, 2009.

[2] G. E Moore et al. cramming more components onto integrated circuits. Proceedings of the IEEE, 86(1):8285, 1998.

[3] Saurabh Kumar Garg, Chee Shin Yeo, Arun Anandasivam, and Rajkumar Buyya. Energy-e cient scheduling of hpc applications in cloud computing environments. CoRR, abs/0909.1146, 2009.

[4] Gartner estimates ict industry accounts for 2 percent of global co2 emissions. http://www.gartner.com/it/page.jsp?id=503867, April 2007

[5] Summary of the Copenhagen climate change conference - 7-19 December 2009 – Copenhagen Denmark. http://www.iisd.ca/vol12/enb12459e.html, 2009.

[6] Michael Armbrust, Armando Fox, Rean Grith, Anthony D. Joseph, Randy Katz, Andy Konwinski, Gunho Lee, David Patterson, Ariel Rabkin, Ion Stoica, and MateiZaharia. Above the clouds: A berkeley view of cloud computing. Technical report, University of Berkeley, California, February 2009.

[7] Energy consumption in data centre (silicon). http://www.podtech.net/home/5116/energy consumption-in-the-data-center,April 2008.

[8] L. Minas and B. Ellison. The problem of power consumption in servers, Intel-lab.Dr. Dobbs Journal May 2009.

[9] Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, and Andrew War eld. Xen and the art of visualization. In SOSP '03: Proceedings of the nineteenth ACM symposium on Operating systems principles, pages 164{177, New York, NY, USA, 2003. ACM.ISBN 1-58113-757-5.doi: http://doi.acm.org/10.1145/945445.945462.

[10] N.M. Mosharaf Kabir Chowdhury and Raouf Boutaba. A survey of network virtualization. Computer Networks, 54(5):862 {876, 2010. ISSN 1389-1286. doi: DOI:10.1016/j.comnet.2009.10.017.URLhttp://www.sciencedirect.com/science/article/B6VRG-4XKBYY4-4/2/91ac30e6e1c4d24c53d95798a6148f03.

[11] P Reisner and L Ellenberg, Drbd v8replicated storage with shared disk semantics. In Proceedings of the 12th International Linux System, 2005, 2005.

[12] Christopher Clark, Keir Fraser, Steven Hand, Jacob Gorm Hansen, Eric Jul, Chris-tian Limpach, Ian Pratt, and Andrew War eld. Live migration of virtual machines. In NSDI'05: Proceedings of the 2nd conference on Symposium on Networked Sys-tems Design & Implementation, pages 273{286, Berkeley, CA, USA, 2005. USENIX Association.

[13] Agentless monitoring | agentless server monitoring | http://www.eginnovations.com/web/egagentless.htm

[14] SerafeimZanikolas and RizosSakellariou. Taxonomy of grid monitoring systems. Future Generation Comp. Syst., 21(1):163{188, 2005.

[15] Anton Beloglazov, Rajkumar Buyya, Young Choon Lee, and Albert Y. Zomaya. A taxonomy and survey of energy-e cient data centers and cloud computing systems. CoRR, abs/1007.0066, 2010.

[16] Wattsup. Watts up? products: Meter https://www.wattsupmeters.com/secure/products.php?pn=0, 2009.

[17] Aman Kansal, Feng Zhao, Jie Liu, Nupur Kothari, and Arka Bhattacharya. Virtual machine power metering and provisioning. In ACM Symposium on Cloud Com-puting (SOCC), June 2010. URL http://research.microsoft.com/apps/pubs/default.aspx?id=120435.

[18] Timothy Wood, Prashant Shenoy, and Arun. Black-box and gray-box strategies for virtual machine migration. In 4th USENIX Symposium on Networked System Design and Implementation, pages 229{242, 2007}. http://www.usenix.org/events/nsdi07/tech/wood.html.

[19] Paul Marshall, Kate Keahey, and Timothy Freeman. Elastic site: Using clouds to elastically extend site resources. In CCGRID, pages 43{52. IEEE, 2010}.