Heightening the Performance Efficiency of Virtual Machines in Cloud Data Centres using Dynamic consolidation

Mrs. Saritha
Department of Information Science and Engineering
Sahyadri College of Engineering and Management

Mr. Rithesh Pakkala P.
Department of Computer Science and Engineering
Sahyadri College of Engineering and Management

Abstract

Cloud computing is a set of IT services that are provided to a customer over a network on a leased based manner and with the ability to scale up or down their service requirements. Cloud computing model have led to the establishment of large-scale virtualized data centers. Such data centers consume enormous amounts of electrical energy resulting in high operating costs. In addition, high energy consumption by the infrastructure leads to substantial carbon dioxide (CO2) emissions contributing to the greenhouse effect. By switching the idle node off and by using Dynamic consolidation of virtual machines the Cloud providers can optimize resource usage and reduce energy consumption. Virtualization technologies which are heavily relied on by the Cloud computing environments provide the ability to transfer virtual machines (VM) between the physical systems using the technique of live migration mainly for improving the energy efficiency. Dynamic server consolidation through live migration is an efficient way towards energy conservation in Cloud data centers. The focus of this work is on energy and performance efficient resource management strategies that can be applied in a virtualized data center by a Cloud provider (e.g. Amazon EC2). Performance characteristics of online algorithms for the problem of energy and performance efficient dynamic VM consolidation are investigated.

***Keywords*-- Green IT; Cloud computing; resource management; virtualization; dynamic consolidation**

I INTRODUCTION

Cloud computing is transforming the way people use computers. Cloud computing [1] is also transforming how networked services are run. The provider of a service (the cloud customer) is able to dynamically provide infrastructure to meet the current demand by leasing resources from a hosting company (the cloud infrastructure provider). The cloud infrastructure provider can leverage economies of scale to provide dynamic, on-demand infrastructure at a favorable cost. This proliferation of Cloud computing has resulted in the establishment of large-scale data centers containing thousands of computing nodes and consuming enormous amounts of electrical energy.

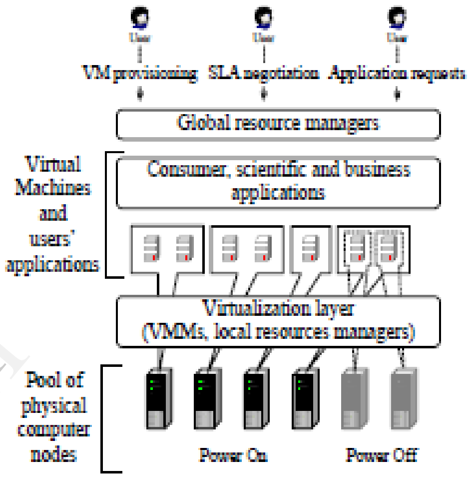

Figure 1. The system view

One of the ways to address the energy inefficiency is to leverage the capabilities of the virtualization technology. The virtualization technology allows Cloud providers to create multiple Virtual Machine (VMs) instances on a single physical server, thus improving the utilization of resources and increasing the Return On Investment (ROI). The reduction in energy consumption can be achieved by switching idle nodes to low-power [5] modes (i.e., sleep, hibernation), thus eliminating the idle power [2][3] consumption**.**

The focus of this work is on energy [4] and performance efficient resource management strategies that can be applied in a virtualized data center by a Cloud provider. First, simplified problem of determining the time to migrate a VM from an oversubscribed host to minimize the cost consisting of the cost of energy [6] consumption and the cost incurred by cloud provider due to SLA violation. We determine and prove the cost of the optimal offline algorithm for this problem, as well as the competitive ratio of the optimal online deterministic algorithm. Second more complex problem of dynamic consolidation of VMs considering multiple hosts and multiple VMs. We find and prove the competitive ratio of the optimal online deterministic algorithm for this problem.

It is widely known that randomized online algorithms which usually provide better performance than deterministic algorithms designed for the same problems. Therefore, we enhance deterministic proposed algorithms is evaluated by extensive simulation using the CloudSim toolkit .The algorithms significantly reduce energy consumption, while providing a high level of adherence to the SLAs.

II. RELATED WORK

One of the first works, in which power management has been applied in the context of virtualized data centers, has been done by Nathuji and Schwan [12]. The authors have proposed architecture of a data center's resource management system where resource management is divided into local and global policies.

Verma et al. [15] have formulated the problem of power-aware dynamic placement of applications in virtualized heterogeneous systems as continuous optimization: at each time frame, the placement of VMs is optimized to minimize power consumption and maximize performance. Like in [14], the authors have applied a heuristic for the bin packing problem with variable bin sizes and costs. Similarly to [12], live migration of VMs is used to achieve a new placement at each time frame. The proposed algorithms, on the contrary to our approach, do not support SLAs: the performance of applications can be degraded due to the workload variability. In their more recent work [16], Verma et al. have proposed dividing VM consolidation strategies into static (monthly, yearly), semistatic (days, weeks) and dynamic (minutes, hours) consolidation. In the paper, the authors have focused on static and semi-static consolidation techniques, as these types of consolidation are easier to implement in an enterprise environment. In contrast, in this work we investigate the problem of dynamic consolidation to take advantage of fine-grained optimization.

Jung et al. [17], [18] have investigated the problem of dynamic consolidation of VMs running a multi-tier web-application using live migration, while meeting SLA requirements. The SLA requirements are modeled as the response time precomputed for each type of transactions specific to the web-application. A new VM placement is produced using bin packing and gradient search techniques. The migration controller decides whether there is a reconfiguration that is effective according to the utility function that accounts for the SLA fulfillment. However, this approach can be applied only to a single web-application setup and, therefore, cannot be utilized for a multitenant IaaS environment. Zhu et al [19] have studied a similar problem of automated resource allocation and capacity planning. They have proposed three individual controllers each operating at a different time scale: longest time scale (hours to days); shorter time scale (minutes); and shortest time scale (seconds). These three controllers place compatible workloads onto groups of servers, react to changing conditions by reallocating VMs, and allocate resources to VMs within the servers to satisfy the SLAs. The middle-scale controller is the closest to the scope of our work. This approach is in line with our previous work and applies an approach based on the idea of setting fixed utilization thresholds. However, fixed utilization thresholds are not efficient for IaaS environments with mixed workloads that exhibit non-stationary resource usage patterns.

In contrast to the discussed studies, we propose efficient adaptive heuristics for dynamic adaption of VM allocation at run-time according to the current utilization of resources applying live migration, switching idle nodes to the sleep mode, and thus minimizing energy consumption. The proposed approach can effectively handle strict QoS requirements, multi-core CPU architectures, heterogeneous infrastructure and heterogeneous VMs.

III. THE SINGLE VM MIGRATION PROBLEM

In this section competitive analysis is applied to analyze a subproblem of the problem of energy [7] and Performance efficient dynamic consolidation of VMs. There is a single physical server, or host, and M VMs allocated to that host. In this problem the time is discrete and can be split into N time frames, where each time frame is 1 second. The resource provider pays the cost of energy consumed by the physical server. The resource capacity of the host and resource usage by VMs is characterized by a single parameter, the CPU performance. The VMs experience dynamic workloads, which mean that the CPU usage by a VM Arbitrarily varies over time. The host is oversubscribed, i.e. if all the VMs request their maximum allowed CPU performance, the total CPU demand will exceed the capacity of the CPU. We define that when the demand of the CPU performance exceeds the available capacity, a violation of the SLAs established between the resource provider and customers occurs. An SLA violation results in a penalty incurred by the provider. At some point in time v, an SLA violation occurs and continues until N time. In other words, due to the over-subscription and variability of the workload experienced by VMs, at the time v the overall demand for the CPU performance exceeds the available CPU capacity and does not decrease until N. It is assumed that according to the problem definition, a single VM can be migrated out of the host. This migration leads to a decrease of the demand for the CPU performance and makes it lower than the CPU capacity.

IV. THE DYNAMIC VM CONSOLIDATION PROBLEM

In this section we analyze a more complex problem of dynamic VM consolidation [8] considering multiple hosts and multiple VMs. For this problem, we define that there are n homogeneous hosts, and the capacity of each host is Ah. Although VMs experience variable workloads, the maximum CPU capacity that can be allocated to a VM is Av. Therefore, the maximum number of VMs allocated to a host when they demand their maximum CPU capacity is m = Ah Av . The total number of VMs is nm. VMs can be migrated between hosts using live migration. As for the single VM migration problem defined in Section III, an SLA violation occurs when the total demand for the CPU performance exceeds the available CPU capacity. We assume that when a host is idle, i.e. there is no allocated VMs, it is switched off and no power, or switched to the sleep mode with negligible power consumption. We call non-idle hosts active. The problem is to determine what time, which VMs and where should be migrated to minimize the total cost C.

V. THE SYSTEM MODEL

In this work, the targeted system is an IaaS environment, represented by a large-scale data center consisting of N heterogeneous physical nodes. Each node is characterized by the CPU [9] performance defined in Millions Instructions Per Second (MIPS), amount of RAM and network bandwidth. The servers do not have local disks, the storage is provided as a Network Attached Storage (NAS) to enable live migration of VMs. The type of the environment implies no knowledge of application workloads and time for which VMs [10] are provisioned. Multiple independent users submit requests for provisioning of M heterogeneous VMs characterized by requirements to processing power defined in MIPS, amount of RAM and network bandwidth. The fact that the VMs are managed by independent users implies that the resulting workload created due to combining multiple VMs on a single physical node is mixed. The mixed workload is formed by various types of applications, such as HPC and web-applications, which utilize the resources simultaneously. The users establish SLAs with the resource provider to formalize the QoS delivered. The provider pays a penalty to the users in cases of SLA violations. The software layer of the system is tiered comprising local and global managers (Figure 2). The local managers reside on each node as a module of the VMM. Their objective is the continuous monitoring of the node's CPU utilization, resizing the VMs [11] according to their resource needs, and deciding when and which VMs should to be migrated from the node (4). The global manager resides on the master node and collects information from the local managers to maintain the overall view of the utilization of resources (2). The global manager issues commands for the optimization of the VM placement (3). VMMs perform actual resizing and migration of VMs as well as changes in power modes of the nodes (5).

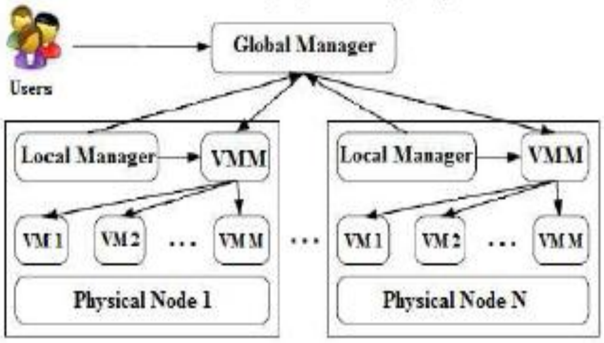

Figure 2. The system model

VI. ADAPTIVE HEURISTICS FOR DYNAMIC VM CONSOLIDATION

According to the analysis presented in Sections III and IV, in this section we propose several heuristics for dynamic consolidation of VMs based on an analysis of historical data of the resource usage by VMs. We split the problem of dynamic VM consolidation into four parts: (1) determining when a host is considered as being overloaded requiring migration of one or more VMs from this host; (2) determining when a host is considered as being underloaded leading to a decision to migrate all VMs from this host and switch the host to the sleep mode; (3) selection of VMs that should be migrated from an overloaded host; and (4) finding a new placement of the VMs selected for migration from the overloaded and underloaded hosts. We discuss the defined subproblems in the following sections.

The general algorithm of VM placement optimization is shown in Algorithm 1. First, the algorithm looks through the list of hosts and by applying the overloading detection algorithm checks whether a host is overloaded. If the host is overloaded, the algorithm applies the VM selection policy to select VMs that need to be migrated from the host. Once the list of VMs to be migrated from the overloaded hosts is built, the VM placement algorithm is invoked to find a new placement for the VMs to be migrated. The second phase of the algorithm is finding overloaded hosts and a placement of the VMs from these hosts. The algorithm returns the combined migration map that contains the information on the new VM placement of the VM selected to be migrated from both overloaded and underloaded hosts.

Algorithm 1: VM placement Optimization

```
Input: hostList Output: migrationMap
foreach host in hostList do
    if isHostOverloaded (host) then
        vmsToMigrate.add
            (getVmsToMigrateFrom
             OverloadedHost (host)
migrationMap.add
(getNewVmPlacement(vmsToMigrate))
vmsToMigrate.clear()
foreach host in hostList do
    if isHostUnderloaded (host) then
        vmsToMigrate.add(host.getVmList()
        migrationMap.add
(getNewVmPlacement(vmsToMigrate))
return migrationMap
```

VM Placement Power Aware Best Fit Decreasing (PABFD) is used which is presented in Algorithm 2.

Algorithm 2: Power Aware Best Fit Decreasing (PABFD)

```
Input: hostList, vmList
Output: allocation of VMs
vmList.sortDecreasingUtilization ()
foreach vm in vmList do
    minPower =MAX
        allocatedHost= NULL
        foreach host in hostList do
            if host has enough resources for vm then
                power = estimatePower(host, vm)
                if power < minPower then
                    allocatedHost = host
                    minPower = power
        if allocatedHost = NULL then
            allocation.add (vm, allocatedHost)
return allocation
```

VII.EXPERIMENTAL RESULTS

CloudSim toolkit has been preferred as a simulation platform, as it is a modern simulation framework targeted at Cloud computing environments. In contrast to another simulation toolkits (e.g. SimGrid, GangSim), it permits the modelling of virtualized environments, supporting on demand resource provisioning, and their management. It has been extended to enable energy aware simulations, as the core framework does not provide this ability. Apart from the energy consumption modelling and accounting to the capability to simulate service applications with dynamic workloads has been derived. Firstly, we tried to work on analysis of concept of Adaptive Migration Threshold and its implementation on Cloudsim Toolkit. Then we went on studying online deterministic energy efficient algorithms Throughout this experiment, the proposed policy(Adaptive Migration Thresholds) is compared with other dynamic policies to reflect the goal of energy saving and reducing number of Migrations and SLA violations.

From the figure 3, we can see that the energy consumption for Adaptive Migration Thresholds (ADT) based Policy is significantly less than the other dynamic policies. Figure 4, shows the Energy utilization graph Here also proposed ADT policy works better than others in reducing No. of Migrations.

Following figures shows the comparative results of all algorithms represented as graphs.

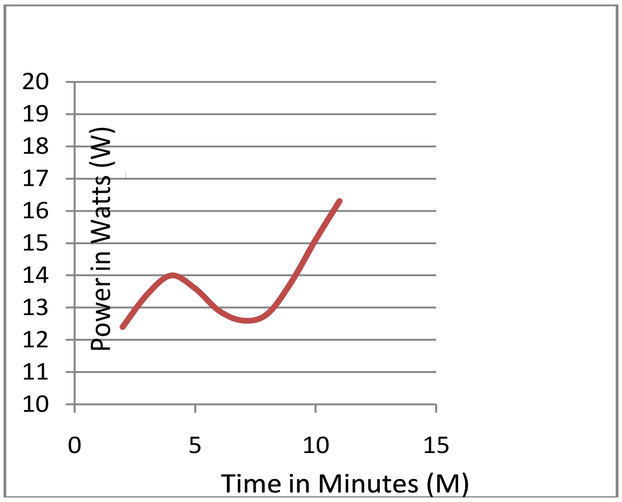

Figure 3:Energy consumption graph

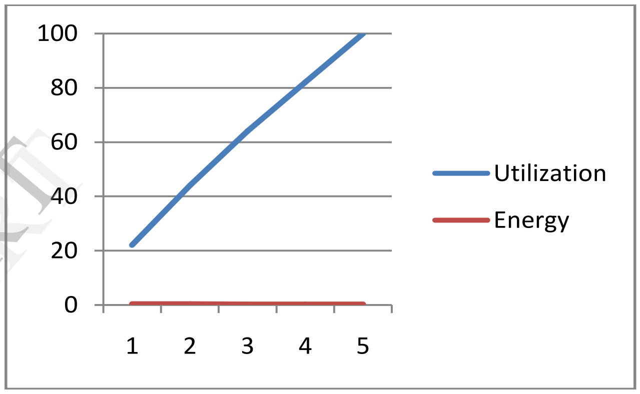

Figure 4: energy utilization graph

From the results, we can make several conclusions: (1) dynamic VM consolidation algorithms significantly outperforms static allocation policies, such as NPA and DVFS;

(2) heuristic-based dynamic VM consolidation algorithms substantially outperform the optimal online deterministic algorithm due to a vastly reduced level of SLA violations;

(3) dynamic VM consolidation algorithms based on local regression outperform the threshold-based and adaptive-threshold based algorithms due to better predictions of host overloading, and therefore decreased SLA violations due to host overloading (SLATAH) and the number of VM migrations.

VIII. CONCLUSION AND FUTURE WORKS

Applying the dynamic consolidation to VM and switching idle servers to maximize the ROI can result in violation of SLA negotiated with customers. In this paper we have concluded that it is necessary to develop randomized or adaptive algorithms to improve the performance of optimal deterministic algorithms.

According to the results of the analysis, another direction for future research is the investigation of more complex workload models, e.g. models based on Markov chains, and development of algorithms that will leverage these workload models. Besides the reduction in infrastructure and on-going operating costs, this work also has social significance as it decreases carbon dioxide footprints and energy consumption by modern IT infrastructures.

REFERENCES

1. Buyya R, Yeo CS, Venugopal S, Broberg J, Brandic I. Cloud computing and emerging IT platforms: Vision, hype,and reality for delivering computing as the 5th utility. Future Generation Computer Systems 2009; 25(6):599–616.
2. ASHRAE Technical Committee 99. Datacom equipment power trends and cooling applications 2005.
3. Belady C. In the data center, power and cooling costs more than the it equipment it supports 2007. URL http://www.electronics-cooling.com/articles/2007/feb/a3/.
4. Barroso LA, Holzle U. The case for energy-proportional computing. Computer 2007; 40(12):33–37.
5. Fan X, Weber WD, Barroso LA. Power provisioning for a warehouse-sized computer. Proceedings of the 34th Annual International Symposium on Computer Architecture (ISCA 2007), ACM New York, NY, USA, 2007; 13–23.
6. de Assunao MD, Gelas JP, Lefevre L, Orgerie AC. The Green Grid'5000: Instrumenting and using a Grid with energy sensors. Proceedings of the 5th International Workshop on Distributed
Cooperative Laboratories: Instrumenting the Grid (INGRID 2010), Poznan, Poland, 2010.
7. Ranganathan P, Leech P, Irwin D, Chase J. Ensemble-level power management for dense blade servers. Proceedings of the 33rd International Symposium on Computer Architecture (ISCA 2006), Boston, MA, USA, 2006; 66–77.
8. The green grid consortium 2011. URL http://www.thegreengrid.org.
9. Xen and the art of virtualization. Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP
2003), Bolton Landing, NY, USA.
10. Clark C, Fraser K, Hand S, Hansen JG, Jul E, Limpach C, Pratt I, Warfield A. Live migration of virtual machines.Proceedings of the 2nd Symposium on Networked Systems Design and Implementation (NSDI 2005), USENIX,Boston, MA, USA, 2005.
11. Ben-David S, Borodin A, Karp R, Tardos G, Wigderson A. On the power of randomization in on-line algorithms.Algorithmica 1994.
12. Nathuji R, Schwan K. Virtualpower: Coordinated power management in virtualized enterprise systems. ACM SIGOPS Operating Systems Review 2007.
13. Kusic D, Kephart JO, Hanson JE, Kandasamy N, Jiang G. Power and performance management of virtualized computing environments via lookahead control. Cluster Computing 2009; 12(1):1–15.
14. Srikantaiah S, Kansal A, Zhao F. Energy aware consolidation for cloud computing. Cluster Computing 2009; 12:1–15.
15. Verma A, Ahuja P, Neogi A. pMapper: Power and migration cost aware application placement in virtualized systems. Proceedings of the 9th ACM/IFIP/USENIX International Conference on Middleware (Middleware 2008),Springer, Leuven, Belgium, 2008; 243–264.
16. Verma A, Dasgupta G, Nayak TK, De P, Kothari R. Server workload analysis for power minimization using consolidation. Proceedings of the 2009 USENIX Annual Technical Conference, San Diego, CA, USA, 2009; 28–28.
17. Jung G, Joshi KR, Hiltunen MA, Schlichting RD, Pu C. Generating adaptation policies for multi-tier applications in consolidated server environments. Proceedings of the 5th IEEE International Conference on Autonomic Computing (ICAC 2008), Chicago, IL, USA, 2008; 23–32.
18. Jung G, Joshi KR, Hiltunen MA, Schlichting RD, Pu C. A cost-sensitive adaptation engine for server consolidation
of multitier applications. Proceedings of the 10th ACM/IFIP/USENIX International Conference on Middleware (Middleware 2009), Urbana Champaign, IL, USA, 2009; 1–20.
19. Zhu X, Young D, Watson BJ, Wang Z, Rolia J, Singhal S, McKee B, Hyser C, Gmach D, Gardner R, et al.. 1000
islands: Integrated capacity and workload management for the next generation data center. Proceedings of the 5th
International Conference on Autonomic Computing (ICAC 2008), Chicago, IL, USA, 2008; 172–181.