International Journal of Research in Advent Technology, Vol.4, No.1, January 2016
E-ISSN: 2321-9637
Available online at www.ijrat.org

Energy-aware Datacenter Resource Allocation with Minimized Virtual Machine Migrations

Saravana M K[1], Harish K[2]
Assistant Professor, Department of Computer Science & Engg, Jyothy Institue of Technology, Bangalore.[1, 2]
Email: mksaravanamk1@gmail.com[1], *kharish1987gmail.com*[2]

Abstract- In recent years, IT infrastructures continue to grow rapidly driven by the demand for computational power created by modern compute-intensive business and scientific applications. However, a large-scale computing infrastructure consumes enormous amounts of electrical power leading to operational costs that exceed the cost of the infrastructure in few years. Except for overwhelming operational costs, high power consumption results in reduced system reliability and devices lifetime due to overheating. Another problem is significant CO2 emissions that contribute to the greenhouse effect. One of the way to reduce power consumption by a data center is to apply virtualization technology. This technology allows one to consolidate several servers to one physical node as Virtual Machines (VMs) reducing the amount of the hardware in use. Recently emerged Cloud computing paradigm leverages virtualization and provides on-demand resource provisioning over the Internet on a pay-as-you-go basis. This allows enterprises to drop the costs of maintenance of their own computing environment and out-source the computational needs to the Cloud. It is essential for Cloud providers to offer reliable Quality of Service (QoS) for the customers that is negotiated in terms of Service Level Agreements (SLA), e.g. throughput, response time. Therefore, to ensure efficient resource management and provide higher utilization of resources, Cloud providers (e.g. Amazon EC2) have to deal with power-performance trade-off, as aggressive consolidation of VMs can lead to performance loss. In this work we leverage live migration of VMs and propose heuristics for dynamic reallocation of VMs according to current resources requirements, while ensuring reliable QoS. The objective of the reallocation is to minimize the number of physical nodes serving current workload, whereas idle nodes are switched off in order to decrease power consumption. A lot of research has been done in power efficient resource management in data centers. In contrast to previous studies, the proposed approach can effectively handle strict QoS requirements, heterogeneous infrastructure and heterogeneous VMs. The algorithms are implemented as fast heuristics, they do not depend on a par-ticular type of workload and do not require any knowledge about applications executing on VMs.

Index Terms- VM Migration, Cloud Datacenter, VM Allocation, QoS.

1. INTRODUCTION

Virtual machines consolidation aims at reducing the number of active physical servers in a data centre with the goal to reduce the total power consumption. In this context, most of the existing solutions rely on aggressive virtual machine migration, thus resulting in unnecessary overhead and energy wastage. This article presents a virtual machine consolidation algorithm with usage prediction (VMCPU) for improving the energy efficiency of cloud data centers. Our algorithm is executed during the virtual machine consolidation process to estimate the short term future CPU utilization based on the local history of the the considered servers. The joint use of current and predicted CPU utilization metrics allows a reliable characterization of overloaded and underloaded servers, thereby reducing both the load and the power consumption after consolidation. We evaluate our proposed solution through simulations on Cloudsim Simulator. In comparison with the state of the art, the obtained results show that consolidation with usage prediction reduces the total migrations and the power consumption of the servers while complying with the service level agreement. Minimizing the use of energy/network communication overhead with maximum resource utilization in large DCs is a challenging task as computing applications and data are growing so quickly for which increasing larger servers and disks are required to process them fast enough within the required time period. In order to minimize network load and maximize resource utilization in cloud DCs, we simulated our proposed network load-aware scheduling algorithm that ensures minimum VMs migration while delivering the negotiated Quality-of-Service (QoS).

2. LITERATURE SURVEY

Farahnakian, Ashraf , Pahikkala [1]. presented a novel dynamic Virtual Machine consolidation approach called ACS-based VMConsolidation. It reduces the energy consumption of datacenters by consolidating VMs into a reduced number of active Physical Machines while preserving Quality of Service requirements. Since the VM consolidation problem is strictly NP-hard, they used the Ant Colony System to find a near-optimal solution. We defined a multi-objective function that considers both the number of dormant PMs and the number of migrations. When compared to the existing dynamic VM consolidation approaches, ACS-VMC not only reduced the energy consumption, but also minimized SLA violations and the number of migrations. They evaluated the performance of proposed approach by conducting experiments with ten different real workload traces.

Shahzad,.Umer, Nazir, [2] developed an efficient load balancing algorithm by using scheduling to minimize VMs migration, which avoid different network performance parameters like congestion, latency etc. There work categories VMs by showing priorities of the VMs, which will help to reduce network overhead.

Zhao, Lu [3] developed online VM placement algorithms to increase cloud provider's revenue by reducing SLA violation cost. First-Fit and Harmonic algorithms are devised without considering VM migrations, while LRF and DDG are devised for VM migration considering VM migrations. There analysis shows that First Fit and Harmonic perform the same in their worst case, and is comparable to the ratio of lower bound to upper bound. They have conduct experiments to evaluate the algorithms using synthetic and real data, respectively. It is found that Harmonic could create more revenue than First-Fit by more than 10 percent when job arriving rate is greater than 1.0. DDG algorithm is applicable in scenarios when SLA penalty is high, job arriving rate is low and the migration cost is low, while LRF performs better in the opposite situations. Through evaluation on the real trace from TJUHPCC, they find that First-Fit could yield more revenue than Harmonic by 1.7 percent if migration is not allowed, and DDG creates more revenue than LRF by 1.23percent if migration is allowed. By comparing proposals against the algorithms adopted in Open-stack and Cloudstack, they find that FirstFit and DDGcould create more revenue in the TJUHPCC case. There algorithms currently could take effect in IaaS systems, where VM requests are not constrained by their communications.

Dhanoa & Khurmi, [4] have analyzed the impact of VM size and network bandwidth on VM migration time and energy consumption of the source system. Variation in VM size and network bandwidth results a significant impact on energy consumption of source system during VM Live migration. Further we can reduce energy consumption and migration time of subsystems by selecting VM with least memory size for migration and increased network bandwidth. Results of this study would help to design algorithm to optimize energy requirements in live migration of VMs. Live migration feature of Virtualization has great potential to optimize energy efficiency during live migration.

Vahora & Patel [5] Presents VM management technique for efficient utilization of resources which leads to reduce energy consumption and number of VM migration in virtualized data centers. There are number of research has been carried out in this subject, some are practical based and some are simulator-based. By analyzing related work on this subject, they found that it is critical and essential to handle three main things: first, how to allocate VM on host such that it is not over- loaded, second, which VM should be selected for migration from overloading hosts, third where to place (Reallocate) VM which is selected for migration. For this we have to balance the load such that host is not over loaded or under loaded with the requests. If host is over loaded one or more number of VM migrated from the host and if host is under loaded resources are not properly utilized which leads to unnecessary energy consumed by them. So, proper management of VM is necessary which is done by the algorithm and experiments shown that VM management techniques proposed by them performs greatly better than previous work on the whole.

3. PRELIMINARIES

Recent developments in virtualization have resulted in its proliferation across datacenters. By supporting the movement of VMs between physical nodes, it enables dynamic migration of VMs according to the performance requirements. When VMs do not use all the provided resources, they can be logically resized and consolidated to the minimum number of physical nodes, while idle nodes can be switched to the sleep mode to eliminate the idle Power consumption and reduce the total energy consumption by the datacenter. Currently, resource allocation in a Cloud data center aims to provide high performance while meeting SLAs, without focusing on allocating VMs to minimize energy consumption. To explore both performance and energy efficiency, three crucial issues must be addressed. First, excessive power cycling of a server could reduce its reliability. Second, turning resources off in a dynamic environment is risky from the QoS perspective. Due to the variability of the workload and aggressive consolidation, some VMs may not obtain required resources under peak load, and fail to meet the desired QoS. Third, ensuring SLAs brings challenges to accurate application performance management in virtualized environments. All these issues require effective consolidation policies that can minimize energy consumption without compromising the user-specified QoS requirements.

4. VM PLACEMENT

The problem of VM allocation can be divided in two parts: the first part is the admission of new requests for VM provisioning and placing the VMs on hosts, whereas the second part is the optimization of the current VM allocation. The first part can be seen as a bin packing problem with variable bin sizes and prices.

To solve it we apply a variant of the Best Fit Decreasing (BFD) algorithm that is shown to use no more than 11/9· OPT+ 1 bins (where OPT is the number of bins given by the optimal solution) [31]. In our modification, the Modified Best Fit Decreasing (MBFD) algorithms, we sort all VMs in decreasing order of their current CPU utilizations, and allocate each VM to a host that provides the least increase of power consumption due to this allocation. This allows leveraging the heterogeneity of resources by choosing the most power-efficient nodes first. The pseudo-code for the algorithm is presented in Algorithm1.The complexity of the allocation part of the algorithm is n . m, where n is the number of VMs that have to be allocated and m is the number of hosts.

***Algorithm 1*:** Improved Best Fit Decreasing (IBFD)

```
Input: hostList, vmList Output: allocation of VMs
vmList.sortDecreasingUtilization()
foreach vm in vmList do
    minPower←MAX
    allocated Host←NULL
    for each host in hostList do
    if  host has enough resource for vm then
            power←estimatePower(host, vm)
            if power < minpower then
                allocatedHost←host
                minPower←power
    if  allocated Host ≠ NULL then
          allocate vm to allocated Host
return allocation
```

4.1. *VM Selection*

The optimization of the current VM allocation is carried out in two steps: at the first step we select VMs that need to be migrated, at the second step the chosen VMs are placed on the hosts using the IBFD algorithm. To determine when and which VMs should be migrated, we introduce three double-threshold VM selection policies. The basic idea is to set upper and lower utilization.

Thresholds for hosts and keep the total utilization of the CPU by all the VMs allocated to the host between these thresholds. If the CPU utilization of a host falls below the lower threshold ,all VMs have to be migrated from this host and the host has to be switched to the sleep mode in order to eliminate the idle power consumption.

If the utilization exceeds the upper threshold, some VMs have to be migrated from the host to reduce the utilization. The aim is to preserve free resources in order to prevent SLA violations due to the consolidation in cases when the utilization by VMs increases.

The difference between the old and new placements forms a set of VMs that have to be reallocated. The new placement is achieved using live migration of VMs. In the following sections we discuss the proposed VM selection policies.

4.1.1 *The minimization of migrations policy*

The Minimum Migrations (MM) policy selects the minimum number of VMs needed to migrate from a host to lower the CPU utilization below the upper utilization threshold if the upper Threshold is violated. Let V_j be a set of VMs currently allocated to the host j.Then P(Vj) is the powerset of Vj. The MM policy finds a set R∈P (Vj) defined in (3).

$$R=\begin{cases}\begin{cases}S|S\mathcal{E}P\,(v_j)u_j-\sum_{v\mathcal{E}S}u_a(v)<T_{u,}\\ \quad |S|\rightarrow min\}\end{cases} & if\, u_j>T_u\\ V_j, & if\ u_j>T_l\\ \phi, & otherwise\end{cases}\qquad(3)$$

Where Tu is the upper utilization threshold; T_l is the lower utilization threshold, u_j is the current CPU utilization of the host j; and $u_a(v)$ is the fraction of the CPU utilization allocated to the VMv. The pseudo-code for the MM algorithm for the over-utilization case is presented in Algorithm 2. The algorithm sorts the list of VMs in the decreasing order of the CPU utilization. Then, it repeatedly looks through the list of VMs and finds a VM that is the best to migrate from the host. The best VM is the one that satisfies two conditions. First, the VM should have the utilization higher than the difference between the host's overall utilization and the upper utilization threshold. Second, if the VM is migrated from the host, the difference between the upper threshold and the new utilization is the minimum across the values provided by all the

VMs. If there is no such a VM, the algorithm selects the VM with the highest utilization, removes it from the list of VMs, and proceeds to a new iteration. The algorithm stops when the new utilization of the host is below the upper utilization threshold. The complexity of the algorithm is proportional to the product of the number of over-utilized hosts and the number of VMs allocated to these hosts.

Algorithm 2: Min Migrations (MM)

```
Input: hos tList Output: migration List
foreach h in host List do
    vmList←h.get VmList()
    vmList.sort DecreasingUtilization()
    hUtil←h.getUtil()
    bestFitUtil←MAX
    while hUtil>THRESH_UP do
      foreach vm in vmList do
        if vm.getUtil()>hUtil−THRESH_UP then
  t←vm.getUtil()−hUtil + THRESH_UP
          if t < bestFitUtil then
            bestFitUtil←t
            bestFitVm←vm
          else
            if bestFitUtil=MAX then
                  bestFitVm←vm
            break
        hUtil←hUtil−bestFitVm.getUtil()
        migrationList.add(bestFitVm)
        vmList.remove(bestFitVm)
        if hUtil<THRESH_LOW then
        migrationList.add(h.getVmList())
        vmList.remove(h.getVmList())
return migrationList
```

4.1.2. *The highest potential growth policy*

When the upper threshold is violated, the Highest Potential Growth (HPG) policy migrates VMs that have the lowest usage of the CPU relatively to the CPU capacity defined by the VM parameters in order to minimize the potential increase of the host's utilization and prevent an SLA violation, as formalized in (4)

$$R=\begin{cases}\begin{cases} S \mid S\ \mathcal{E}\rho\ (v_j) u_j - \sum_{v\mathcal{E}S} u_a < vT_u \\ \sum_{v\varepsilon S} \frac{u_a(v)}{u_r(v)} \rightarrow min\}, if\ u_j > T_u; \end{cases} & (4) \\ V_j, & if\ u_j > T_l \\ \phi, & otherwise \end{cases}$$

where $u_r(v)$ is the fraction of the CPU capacity initially requested for the VM_v and defined as the VM's parameter. we do not provide the pseudo-code for the HPG algorithm, as it is similar to the MM algorithm presented earlier.

4.1.3. *The random choice policy*

The Random Choice (RC) policy relies on ar and om selection of a number of VMs needed to decrease the CPU utilization by a host below the upper utilization threshold. According to a uniformly distributed discrete random variable(X), whose values index subsets of Vj, the policy selects a set R∈P (Vj),as shown in (5).

$$R=\begin{cases}\begin{cases} S \mid S\ \mathcal{E}P\ (v_j), u_j - \sum_{v\mathcal{E}S} u_a(v) < T_u \\ X \overset{d}{\cong} U(0, |P(0, |P(V_j)| - 1)\}, if\ u_j > T_u \end{cases} \\ V_j, & if\ u_j > T_l \\ \phi, & otherwise \end{cases} \quad (5)$$

Where X is a uniformly distributed discrete random variable used to select a subset of Vj. The results of a simulation-based evaluation of the proposed algorithms in terms of power consumption, SLA violations and the number of VM migrations are presented in Section 5.

5. EXPERIMENTS AND RESULTS

In this section, we discuss a performance analysis of the energy-aware allocation heuristics presented in Section 4. We have conducted our experiments on cloudSim Simulator , we calculate the time needed to perform a live migration of a VM as the size of its memory divided by the available network bandwidth. For the simulations, the utilization of the CPU by a VM is generated as a uniformly distributed random variable. This is appropriate due to unknown types of applications running on VMs, and as it is not possible to build the exact model of such a mixed workload.

We have simulated a data center comprising 100 heterogeneous physical nodes. Each node is modeled to have one CPU core with the performance equivalent to1000, 2000 or 3000MIPS, 8GB Of RAM and 1TB of storage.

A host consumes from 175W with 0% CPU utilization, up to 250W with 100% CPU utilization. Each VM requires one CPU core with 250, 500, 750 or 1000 MIPS, 128 MB of RAM and 1GB of storage. The users submit requests for provisioning of 290 Heterogeneous VMs that fill the full capacity of the simulated data center. Each VM runs a web-application or any kind of application with variable workload, which is modeled to generate the utilization of CPU according to a uniformly distributed random variable.

The application runsfor150,000 MI that is equal to 10 min of the execution on 250MIPS CPU with 100% utilization. Initially, the VMs are allocated according to the requested characteristics assuming 100% CPU utilization. Each experiment has been run 10 times.

5.1. *Performance metrics*

In order to compare the efficiency of the algorithms we use several metrics to evaluate their performance. The first metric is the total energy consumption by the physical resources of a data center caused by the application workloads. The second performance metric is called the SLA violation percentage, or simply the SLA violations, which is defined as the percentage of SLA violation events relatively to the total number of the processed timeframes. We define that an SLA violation occurs when a given VM cannot get the amount of Million Instructions Per Second (MIPS) that are requested.

5.2. *Simulation Results*

For the benchmark experimental results we have used a VM migration aware policy called Single Threshold (ST). It is based on the idea of setting the upper utilization threshold for hosts and placing VMs, while keeping the total utilization of CPU below this threshold. At each time frame all VMs are reallocated using the IBFD algorithm with additional condition of keeping the upper utilization threshold not violated. To evaluate the ST policy we have conducted several experiments with different values of the utilization threshold. The simulation results are presented in Fig.2.

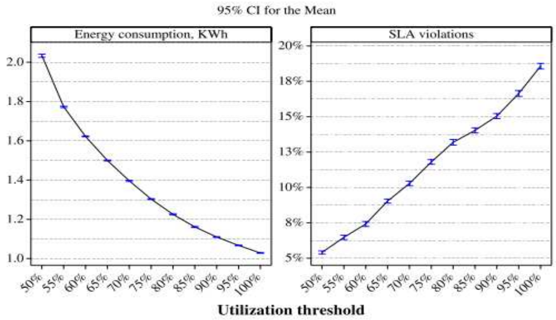

Fig.2. The energy consumption and SLA violations by the ST plicy

The results show that energy consumption can be significantly reduced relatively to the NPA and DVFS policies by77% and 53% respectively with 5.4% of SLA violations. The results show that with the growth of the utilization threshold energy consumption decreases, whereas the percentage of SLA violations increases. This is due to the fact that a higher utilization threshold allows more aggressive consolidation of VMs by the cost of the increased risk of SLA violations.

To evaluate the double-threshold policies it is necessary to determine the best values for the thresholds in terms of the energy consumption and QoS delivered. We have chosen the MM policy to conduct the analysis of the utilization thresholds. We have simulated the MM policy varying the absolute values of the lower and upper thresholds as well as the interval between them. The results showing the mean energy consumption achieved using the MM policy for different values of the lower utilization.

Threshold and the interval between the thresholds are presented in Fig.3. The graph shows that an increase of the lower utilization threshold leads to decreased energy consumption. However, the low level of energy consumption can be achieved with different intervals between the thresholds. Therefore, to determine the best interval we have to consider another factor, the level of SLA violations.

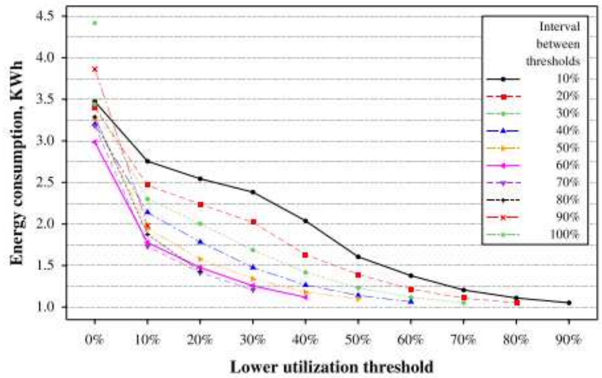

Fig 3.The mean energy consumption by the MM policy for different values of the utilization thresholds.

6. CONCLUSION AND DISCUSSION

In this article, we address the VM consolidation problem by adopting CPU usage prediction. Our aim was to reduce the frequency of the number of VM migrations and the number of server switches in order to save energy. To this end, based on a resource prediction scheme, we proposed a consolidation with usage prediction algorithm for energy efficient cloud data centers. The proposed algorithm effectively reduces not only the number of migrations, the number of power state changes and the energy consumption of the servers, but also the average number of SLA violations. The simulation results has shown that the proposed approach can significantly decrease the energy consumption that results from VM migrations and host switches with a better compliance with the SLA. As a future work, we seek to evaluate the performance of the proposed algorithm across multiple resource dimensions.

REFERENCES

[1] Using Ant Colony System to Consolidate VMs for Green Cloud Computing

[2] Reduced VM Migration in bandwidth over subscribed cloud data Centers

[3] Online Virtual Machine Placement for Increasing Cloud Provider's Revenue.

[4] Efficient Virtual Machine management based on dynamic workload in cloud computing environment.
[5] The green grid consortium, 2011. URL: http: // www.thegreengrid.org.
[6] R. Buyya, A. Beloglazov, J. Abawajy, Energy-efficient management of data center resources for cloud computing: a vision, architectural elements, and open challenges, in: Proceedings of the 2010 International Conference on Parallel and Distributed Processing Techniques and Applications, PDPTA 2010, LasVegas, USA, 2010.
[7] E.Pinheiro, R.Bianchini, E.V.Carrera, T. Heath, Load balancing and unbalancing for power and performance in cluster-based systems , in :Proceedings of the Workshop on Compilers and Operating Systems for Low Power, 2001, pp. 182–195.
[8] 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 Symposiumon Operating Systems Principles, ACM, NewYork, NY, USA, 2001, pp.103–116.
[9] E.Elnozahy, M.Kistler, R.Rajamony, Energy efficient server clusters, Power-Aware Computer Systems (2003) 179 – 197.
[10] R. Nathuji, K. Schwan, Virtualpower: coordinated power management in Virtualized enterprise systems, ACMSIGOPS Operating Systems Review 41 (6) (2007)265–278.
[11] R. Raghavendra, P. Ranganathan, V. Talwar, Z. Wang, X. Zhu, No ''power'' struggles: coordinated multi-level power management for the data center, SIGARCHComputerArchitectureNews36(1)(2008)48–59.
[12] D. Kusic, J.O. Kephart, J.E. Hanson, N. Kandasamy, G. Jiang, Power and Performance management of virtualized computing environments via look ahead control, Cluster Computing12 (1) (2009) 1–15.
[13] S. Srikantaiah, A. Kansal, F. Zhao, Energy aware consolidation for cloud computing, Cluster Computing 12 (2009)1–15.