Implementing Virtual Machines for Dynamic Resource Allocation in Cloud Computing Environment

s. prasanna laxmi

Assistant professor, CSE department Christu jyothi institute of technology & science Yeshwanthapur, jangaon

Abstract— To scale up and down the resource usage of stake holders such as customers, the cloud computing environment is used. In this paper, we present a system that uses virtualization technology to allocate data center resources dynamically based on application demands, the green computing is supported by optimizing the number of servers in use. We introduce the concept of "skewness" to measure the unevenness in the multidimensional resource utilization of a server. By minimizing skewness, we can combine different types of workloads nicely and improve the overall utilization of server resources. We develop a set of heuristics that prevent overload in the system effectively while saving energy used. Trace driven simulation and experiment results demonstrate that our algorithm achieves good performance.

Index Terms—Cloud computing, resource management, virtualization, green computing.

1. Introduction

There are great discussion how to move legacy applications onto the cloud platform and here we study how a cloud service provider best can multiplex its virtual resources onto the physical hardware. This is important because much of the touted gains in the cloud model come from multiplexing. It is observed that in many existing data centers the servers are underutilized due to over provisioning for the peak demand. [1], [2]. The cloud model is expected to make such practice unnecessary by offering automatic scale up and down in response to load variation. Besides reducing the hardware cost, it also saves on electricity which contributes to significant portion of the operational expenses in large data centers.

Virtual machine monitors (VMMs) like Xen provide a mechanism for mapping virtual machines (VMs) to physical resources [3]. This mapping is largely hidden from the cloud users. It is up to the cloud provider to make sure the underlying physical machines (PMs) have sufficient re- sources to meet their needs. VM live migration technology makes it possible to change the mapping between VMs and PMs While applications are running [5], [6]. However, a policy issue remains as how to decide the mapping adaptively so that the resource demands of VMs are met while the number of PMs used is minimized. This is challenging when the resource needs of VMs are hetero-geneous due to the diverse set of applications they run and vary with time as the workloads grow and shrink.

We aim to achieve two goals in our algorithm:

- Avoiding overloading: The capacity of a PM should be sufficient to satisfy the resource needs of all VMs running on it. Otherwise, the PM is overloaded and can lead to degraded performance of its VMs.
- Green Computing: The number of PMs used should be minimized as long as they can still satisfy the needs of all VMs. Idle PMs can be turned off to save energy.

For overload avoidance, we should keep the utilization of PMs low to reduce the possibility of overload in case the resource needs of VMs increase later.

For green computing, we should keep the utilization of PMs reasonably high to make efficient use of their energy.

In this paper, we present the design and implementation of an automated resource management system that achieves a good balance between the two goals. We make the following contributions:

- We develop a resource allocation system that can avoid overload in the system effectively while minimizing the number of servers used.
- We introduce the concept of “skewness” to measure the uneven utilization of a server. By minimizing skewness, we can improve the overall utilization of servers in the face of multidimensional resource constraints.
- We design a load prediction algorithm that can capture the future resource usages of applications accurately without looking inside the VMs. The algorithm can capture the rising trend of resource usage patterns and help reduce the placement churn significantly.

The rest of the paper is organized as follows. Section-2 provides an overview of our system and Section -3 describes our algorithm to predict resource usage. Section 4 provides simulation and Section-5 presents experiment results, respectively. Section- 6 discusses related work and Section- 7 concludes.

2. System Overview

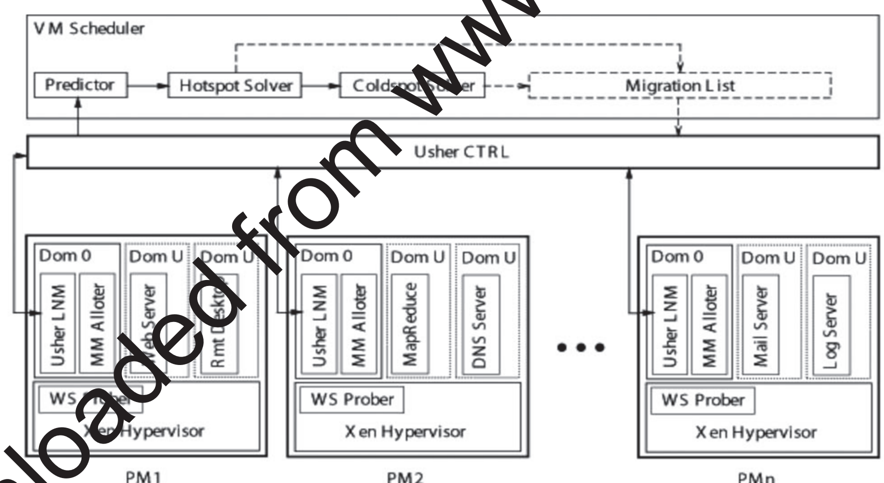

Figure. 1. System architecture.

The architecture of the system is presented in Fig. 1. Each PM runs the Xen hypervisor (VMM) which supports a privileged domain 0 and one or more domain U [3]. Each VM in domain U encapsulates one or more applications such as Web server, remote desktop, DNS, Mail, Map/ Reduce, etc. We assume all PMs Share back end storage.

The multiplexing of VMs to PMs is managed using the Usher framework [7]. The main logic of our system is implemented as a set of plug-ins to usher. Each node runs an Usher local node manager (LNM) on domain 0 which collects the usage statistics of resources for each VM on that node. The CPU and network usage can be calculated by monitoring the scheduling events in Xen. The memory usage with in a VM, however, is not visible to the hypervisor. One approach is to infer memory shortage of a VM by observing its swap activities [8]. Unfortunately, the guest OS is required to install a separate swap partition. Furthermore, it may be too late to adjust the memory allocation by the time swapping occurs. Instead we implemented a working set prober (WS Prober) on each hypervisor to estimate the working set sizes of VMs running on it. We use the random page sampling technique as in the VM ware ESX Server [9].The statistics collected at each PM are forwarded to the Usher central controller (Usher CTRL) where our VM scheduler runs. The VM Scheduler is invoked periodically and receives from the LNM the resource demand history of VMs, the capacity and the load history of PMs, and the current layout of VMs on PMs. The scheduler has several components. The predictor predicts the future resource demands of VMs and the future load of PMs based on past statistics. We compute the load of a PM by aggregating the resource usage of its VMs. The LNM at each node first attempts to satisfy the new demands locally by adjusting the resource allocation of VMs sharing the same VMM. Xen can change the CPU allocation among the VMs by adjusting their weights in its CPU scheduler. The MM Allotter on domain 0 of each node is responsible for adjusting the local memory allocation.

The hot spot solver in our VM Scheduler detects if the resource utilization of any PM is above the hot threshold (i.e., a hot spot). If so, some VMs running on them will be migrated away to reduce their load. The cold spot solver checks if the average utilization of actively used PMs (APMs) is below the green computing threshold. If so, some of those PMs could potentially be turned off to save energy. It identifies the set of PMs whose utilization is below the cold threshold (i.e., cold spots) and then attempts to migrate away all their VMs. It then compiles a migration list of VMs and passes it to Usher CTRL for execution.

3 The Skewness Algorithm

We introduce the concept of skewness to quantify the unevenness in the utilization of multiple resources on a server. Let n be the number of resources we consider r_i be the utilization of the ith resource. We define the resource skewness of a server p as

s n

$$\text{Skewness}(p) = \sum r_i^2$$

where r is the average utilization of all resources for server p. In practice, not all types of resources are performance critical and hence we only need to consider bottleneck resources in the above calculation. By minimizing the skewness, we can combine different types of workloads nicely and improve the overall utilization of server resources. In the following, we describe the details of our algorithm. Analysis of the algorithm is presented in Section 1..

3.1 Hot and Cold Spots

Our algorithm executes periodically to evaluate the resource allocation status based on the predicted future resource demands of VMs. We define a server as a hot spot if the utilization of any of its resources is above a hot threshold. This indicates that the server is overloaded and hence some VMs running on it should be migrated away. We define the temperature of a hot spot p as the square sum of its resource utilization beyond the hot threshold, move onto the next hot spot. Note that each run of the algorithm migrates away at most one VM from the overloaded server. This does not necessarily eliminate the hot spot, but at least reduces its temperature. If it remains a hot spot in the next decision run, the algorithm will repeat this process. It is possible to design the algorithm so that it can migrate away multiple VMs during each run. But this can add more load on the related servers during a period when they are already overloaded. We decide to use this more conservative approach and leave the system some time to react before initiating additional migrations.

3.2. Green Computing

When the resource utilization of active servers is too low, some of them can be turned off to save energy. This is where R is the set of overloaded resources in server p and rt is the hot threshold for resource r. (Note that only overloaded resources are considered in the calculation.) The temperature of a hot spot reflects its degree of overload. If a server is not a hot spot, its temperature is zero.

We define a server as a cold spot if the utilizations of all its resources are below a cold threshold. This indicatesthat the server is mostly idle and a potential candidate to turn off to save energy. However, we do so only when the average resource utilization of all actively used servers (i.e., APMs) in the system is below a green computing threshold. A server is actively used if it has at least one VM running. Otherwise, it is inactive. Finally, we define the warm threshold to be a level of resource utilization that is sufficiently high to justify having the server running but not so high as to risk becoming a hot spot in the face of temporary fluctuation of application resource demands.

Different types of resources can have different thresh- olds. For example, we can define the hot thresholds for CPU and memory resources to be 90 and 80 percent, respectively. Thus a server is a hot spot if either its CPU usage is above 90 percent or its memory usage is above 80 percent.

3.3 Hot Spot Mitigation

We sort the list of hot spots in the system in descending temperature (i.e., we handle the hottest one first). Our goal is to eliminate all hot spots if possible. Otherwise, keep their temperature as low as possible. For each server p, we first decide which of its VMs should be migrated away. We sort its list of VMs based on the resulting temperature of the server if that VM is migrated away. We aim to migrate away the VM that can reduce the server's temperature the most. In case of ties, we select the VM whose removal can reduce the skewness of the server the most. For each VM in the list, we see if we can find a destination server to accommodate it. The server must not become a hot spot after accepting this VM. Among all such servers, we select one whose skewness can be reduced the most by accepting this VM. Note that this reduction can be negative which means we select the server whose skewness increases the least. If a destination server is found, we record the migration of the VM to that server and update the predicted load of related servers. Otherwise, we move onto the next VM in the list and try to find a destination server for it. As long as we can find a destination server for any of its VMs, we consider this run of the algorithm a success and then handled in our green computing algorithm. The challenge here is to reduce the number of active servers during low load without sacrificing performance either now or in the future. We need to avoid oscillation in the system.

Our green computing algorithm is invoked when the average utilizations of all resources on active servers are below the green computing threshold. We sort the list of cold spots in the system based on the ascending order of their memory size. Since we need to migrate away all its VMs before we can shut down an underutilized server, we define the memory size of a cold spot as the aggregate memory size of all VMs running on it. Recall that our model assumes all VMs connect to share back-end storage. Hence, the cost of a VM live migration is determined mostly by its memory footprint. Section 7 in the supplementary file explains why the memory is a good measure in depth. We try to eliminate the cold spot with the lowest cost first.

For a cold spot p, we check if we can migrate all its VM somewhere else. For each VM on p, we try to find a destination server to accommodate it. The resource utilizations of the server after accepting the VM must be below the warm threshold. While we can save energy by consolidating underutilized servers, overdoing it may create hot spots in the future. The warm threshold is designed to prevent that. If multiple servers satisfy the above criterion, we prefer one that is not a current cold spot. This is because increasing load on a cold spot reduces the likelihood that it can be eliminated. However, we will accept a cold spot as the destination server if necessary. All things being equal, we select a destination server whose skewness can be reduced the most by accepting this VM. If we can find destination servers for all VMs on a cold spot, we record the sequence of migrations and update the predicted load of related servers. Otherwise, we do not migrate any of its VMs. The list of cold spots is also updated because some of them may no longer be cold due to the proposed VM migrations in the above process.

The above consolidation adds extra load onto the related servers. This is not as serious a problem as in the hot spot mitigation case because green computing is initiated only when the load in the system is low. Nevertheless, we want to bound the extra load due to server consolidation. We restrict the number of cold spots that can be eliminated in each run of the algorithm to be no more than a certain percentage of active servers in the system. This is called the consolidation limit.

Note that we eliminate cold spots in the system only when the average load of all active servers (APMs) is below the green computing threshold. Otherwise, we leave those cold spots there as potential destination machines for future offloading. This is consistent with our philosophy that green computing should be conducted conservatively.

4. Simulations

We evaluate the performance of our algorithm using trace driven simulation. Note that our simulation uses the same code base for the algorithm as the real implementation in the experiments. This ensures the fidelity of our simulation results. Traces are per minute server resource utilization, such as CPU rate, memory usage, and network traffic statistics, collected using tools like “perfmon” (Windows), the “/proc” file system (Linux), “pmstat/vmstat/netstat” commands (Solaris), etc.. The raw traces are pre-processed into “Usher” format so that the simulator can read them. We collected the traces from a variety of sources:

- Web InfoMall. The largest online Web archive in China (i.e., the counterpart of Internet Archive in the US) with more than three billion archived Web pages.
- Real course. The largest online distance learning system in China with servers distributed across 13 major cities.
- Amazing Store. The largest P2P storage system in China.

We also collected traces from servers and desktop computers in our university including one of our mail servers, the central DNS server, and desktops in our department. We post processed the traces based on days collected and use random sampling and linear combination of the data sets to generate the workloads needed. All simulation in this section uses the real trace workload unless otherwise specified.

We used the FUSD load prediction algorithm with " ¼ 0:2, # ¼ 0:7, and W ¼ 8. In a dynamic system, those parameters represent good knobs to tune the performance of the system adaptively. We choose the default parameter values based on empirical experience working with many Internet applications. In the future, we plan to explore using AI or control theoretic approach to find near optimal values automatically.

4.1 Effect of Thresholds on APMs

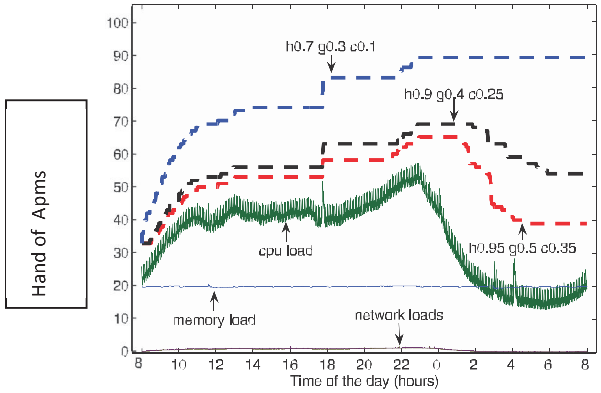

Figure 2 Impact of thresholds on the number of APMs

We first evaluate the effect of the various thresholds used in our algorithm. We simulate a system with 100 PMs and 1,000 VMs (selected randomly from the trace). We use random to PM mapping in the initial layout. The scheduler is invoked once per minute. The bottom part of Fig. 2 shows the daily load variation in the system. The x- axis is the time of the day starting at 8 am. The y-axis is overloaded with two meanings: the percentage of the load or the percentage of APMs (i.e., Active PMs) in the system. Recall that a PM is active (i.e., an APM) if it has at least one VM running. As can be seen from the figure, the CPU load demonstrates diurnal patterns which decrease substantially after midnight. The memory consumption is fairly stable over the time. The network utilization stays very low.

The top part of Fig. 2 shows how the percentages of APMs vary with the load for different thresholds in our algorithm. For example, "h0.7 g0.3 c0.1" means that the hot, the green computing, and the cold thresholds are 70, 30, and 10 percent, respectively. Our algorithm can be made more or less aggressive in its migration decision by tuning the thresholds. The figure shows that lower hot thresholds cause more aggressive migrations to mitigate hot spots in the system and increases the number of APMs, and higher cold and green computing thresholds cause more aggressive consolidation which leads to a smaller number of APMs. The percentage of APMs in our algorithm follows the load pattern closely.

To examine the performance of our algorithm in more extreme situations, we also create a synthetic workload which mimics the shape of a sine function (only the positive part) and ranges from 15 to 95 percent with a 20 percent random fluctuation. It has a much larger peak-to-mean ratio than the real trace. The results are shown in Section 2 of the supplementary file, which can be found on the Computer Society Digital Library.

4.2 Scalability of the Algorithm

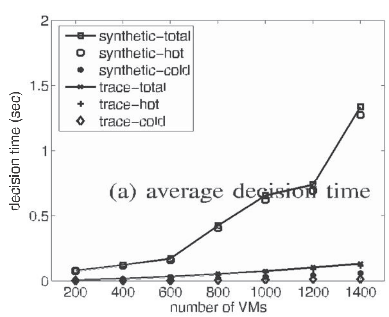

 ISBN : 978-81-925233-3-0

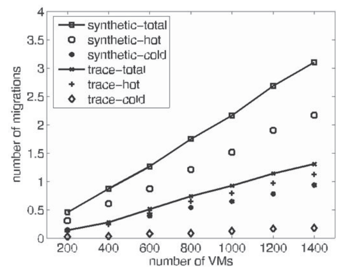

(b) average number of migrations

Figure 3. Scalability of the algorithm with system size.

We evaluate the scalability of our algorithm by varying the number of VMs in the simulation between 200 and 1,400. The ratio of VM to PM is 10:1. The results are shown in

Fig. 3. Fig. 3a shows that the average decision time of our algorithm increases with the system size. The speed of increase is between linear and quadratic. We break down the decision time into two parts: hot spot mitigation (marked as "hot") and green computing (marked as "cold"). We find that hot spot mitigation contributes more to the decision time. We also find that the decision time for the synthetic workload is higher than that for the real trace due to the large variation in the synthetic workload. With 140 PMs and 1,400 VMs, the decision time is about 1.3 seconds for the synthetic workload and 0.2 second for the real trace.

Fig. 3b shows the average number of migrations in the whole system during each decision. The number of migrations is small and increases roughly linearly with the system size. We find that hot spot contributes more to the number of migrations. We also find that the number of migrations in the synthetic workload is higher than that in the real trace. With 140 PMs and 1,400 VMs, on average each run of our algorithm incurs about three migrations in the whole system for the synthetic workload and only 1.3 migrations for the real trace.

4.3 Effect of Load Prediction

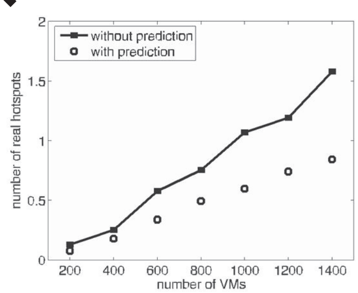

(a) number of hot spots

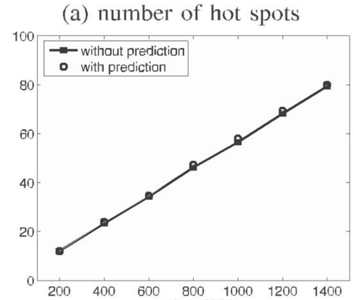

(b) number of APMs

 ISBN : 978-81-925233-3-0

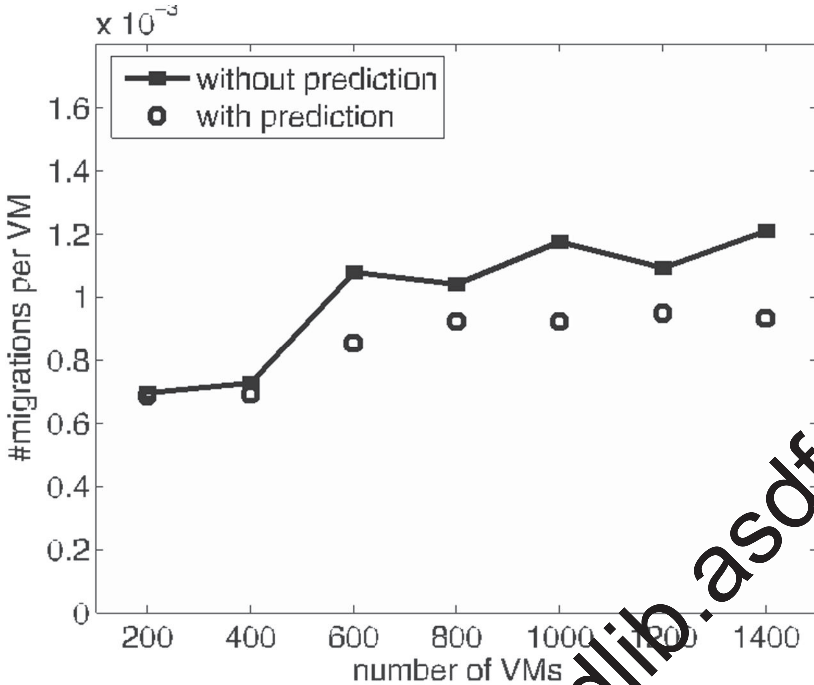

(c) number of migration

Figure. 4. Effect of load prediction

We compare the execution of our algorithm with and without load prediction in Fig. 4. When load prediction is disabled, the algorithm simply uses the last observed load in its decision making. Fig. 4a shows that load prediction significantly reduces the average number of hot spots in the system during a decision run. Notably, prediction prevents over 46 percent hot spots in the simulation with 1,400 VMs. This demonstrates its high effectiveness in preventing server overload proactively. Without prediction, the algo- rithm tries to consolidate a PM as soon as its load drops below the threshold. With prediction, the algorithm correctly foresees that the load of the PM will increase above the threshold shortly and hence takes no action. This leaves the PM in the "cold spot" state for a while. However, it also reduces placement churns by avoiding unnecessary migrations due to temporary load fluctuation. Consequently, the number of migrations in the system with load prediction is smaller than that without prediction as shown in Fig. 4c. We can adjust the conservativeness of load prediction by tuning its parameters, but the current configuration largely serves our purpose (i.e., error on the side of caution). The only downside of having more cold spots in the system is that it may increase the number of APMs. This is investigated in Fig. 4b which shows that the average numbers of APMs remain essentially the same with or without load prediction (the difference is less than 1 percent). This is appealing because significant overload protection can be achieved without sacrificing resources efficiency. Fig. 6c compares the average number of migrations per VM in each decision with and without load prediction. It shows that each VM experiences 17 percent fewer migrations with load prediction.

5 Experiments

Our experiments are conducted using a group of 30 Dell Power Edge blade servers with Intel E5620 CPU and 24 GB of RAM. The servers run Xen-3.3 and Linux 2.6.18. We deploy 8 VMs on each server at the beginning. Each VM is configured with one virtual CPU and two gigabyte memory. Self-ballooning is enabled to allow the hypervisor to reclaim un used memory. Each VM runs the server side of the TPC-W benchmark corresponding to various types of the workloads: browsing, shopping, hybrid workloads, etc. Our algorithm is invoked every 10 minutes.

 ISBN : 978-81-925233-3-0

5.1 Algorithm Effectiveness

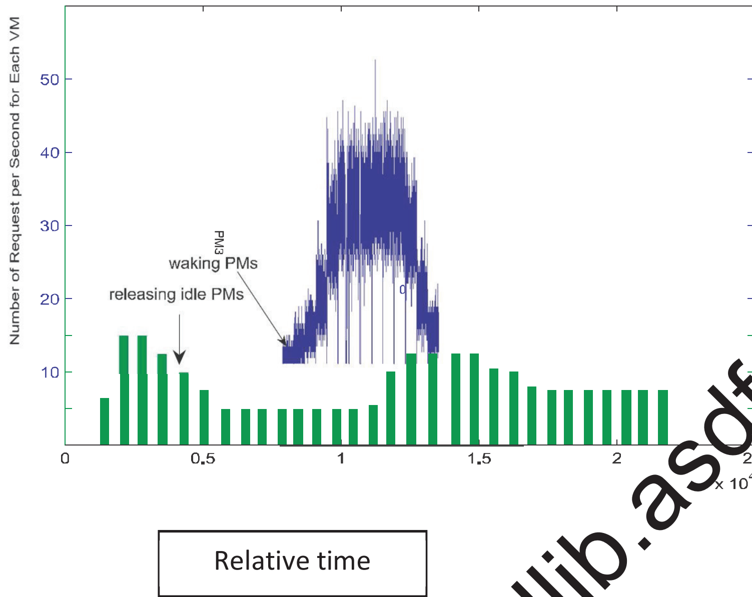

Figure 5. #APMs varies with TPC-W load.

We evaluate the effectiveness of our algorithm in overload mitigation and green computing. We start with a small scale experiment consisting of three PMs and five VMs so that we can present the results for all servers in Fig. 5. Different shades are used for each VM. All VMs are configured with 128 MB of RAM. An Apache server runs on each VM. We use httperf to invoke CPU intensive PHP scripts on the Apache server. This allows us to subject the VMs to different degrees of CPU load by adjusting the client request rates. The utilization of other resources are kept low. We first increase the CPU load of the three VMs on PM_1 to create an overload. Our algorithm resolves the overload by migrating VM_3 to PM_3. It reaches a stable state under high load around 420 seconds. Around 890 seconds, we decrease the CPU load of all VMs gradually. Because the FUSD prediction algorithm is conservative when the load decreases, it takes a while before green computing takes effect. Around 1,700 seconds, VM_3 is migrated from PM_3 to PM_2 so that PM_3 can be put into the standby mode. Around 2,200 seconds, the two VMs on PM_1 are migrated to PM_2 so that PM_1 can be released as well. As the load goes up and down, our algorithm will repeat the above process: spread over or consolidate the VMs as needed.

Next we extend the scale of the experiment to 30 servers. We use the TPC-W benchmark for this experiment. TPC-W is an industry standard benchmark for e-commerce even when idle, consumes several hundred megabytes of memory. After two hours, we increase the load dramatically to emulate a “flash crowd” event. The algorithm wakes up the stand-by servers to offload the hot spot servers. The figure shows that the number of APMs increases accord- ingly. After the request rates peak for about one hour, we reduce the load gradually to emulate that the flash crowd is over. This triggers green computing again to consolidate the underutilized servers. Fig. 5 shows that over the course of the experiment, the number of APM rises much faster than it falls. This is due to the effect of our FUSD load prediction. The figure also shows that the number of APMs remains at a slightly elevated level after the flash crowd. This is because the TPC-W servers maintain some data in cache and hence its memory usage never goes back to its original level.

5.2 Impact of Live Migration

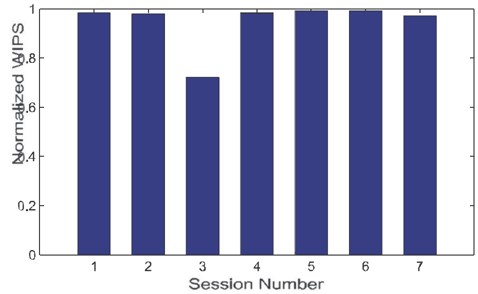

Figure 6. Impact of live migration on TPC-W performance.

One concern about the use of VM live migration is its impact on application performance. Previous studies have found this impact to be small [5]. We investigate this impact in our own experiment. We extract the data on the 340 live migrations in our 30 server experiment above. We find that 139 of them are for hot spot mitigation. We focus on these migrations because that is when the potential impact on application performance is the most. Among the 139 migrations, we randomly pick seven corresponding TPC-W sessions undergoing live migration. All these sessions run the “shopping mix” workload with 200 emu- lated browsers. As a target for comparison, we rerun the session with the same parameters but perform no migra- tion and use the resulting performance as the baseline. Fig. 6 shows the normalized Web interactions per second (WIPS) for the 7 sessions. WIPS is the performance metric used by TPC-W. The figure shows that most live migration sessions exhibit no noticeable degradation in performance compared to the baseline: the normalized WIPS is close to the only exception is session 3 whose degraded performance is caused by an extremely busy server in the original experiment. Next we take a closer look at one of the sessions in Fig. 6 and show how their performances vary over time. The figure verifies that live migration causes no noticeable performance degradation. The duration of the migration is under 10 seconds. Recall that our algorithm is invoked every 10 minutes.

5.3 Resource Balance

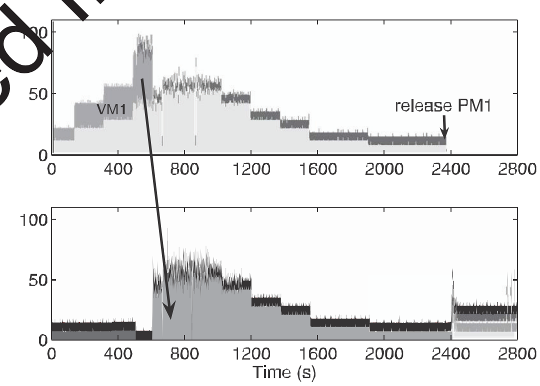

Figure 7 Resource balance for mixed workloads

Recall that the goal of the skewness algorithm is to mix workloads with different resource requirements together so that the overall utilization of server capacity is improved. In this experiment, we see how our algorithm handles a mix of CPU, memory, and network intensive workloads. We vary the CPU load as before. We inject the network load by sending the VMs a series of network packets. The memory intensive applications are created by allocating memory on demand. Again we start with a small scale experiment consisting of two PMs and four VMs so that we can present the results for all servers in Fig. 7. The two rows represent the two PMs. The two columns represent the CPU and network dimensions, respectively. The memory consumption is kept low for this experiment. Initially, the two VMs on PM_1 are CPU intensive while the two VMs on PM_2 are network intensive. We increase the load of their bottleneck resources gradually. Around 500 seconds, VM_4 is migrated from PM_2 to PM_1 due to the network overload in PM_2. Then around 600 seconds, VM_1 is migrated from PM_1 to PM_2 due to the CPU overload in PM_1. Now the system reaches a stable state with a balanced resource utilization for both PMs—each with a CPU intensive VM and a network intensive VM. Later we decrease the load of all VMs gradually so that both PMs become cold spots. We can see that the two VMs on PM_1 are consolidated to PM_2 by green computing.

Next we extend the scale of the experiment to a group of 72 VMs running over 8 PMs. Half of the VMs are CPU intensive, while the other half is memory intensive. Initially, we keep the load of all VMs low and deploy all CPU intensive VMs on PM_4 and PM_5 while all memory intensive VMs on PM_6 and PM_7. Then we increase the load on all VMs gradually to make the underlying PMs hot spots. Fig. 12 shows how the algorithm spreads the VMs to other PMs over time. As we can see from the figure, the algorithm balances the two types of VMs appropriately. The figure also shows that the load across the set of PMs becomes well balanced as we increase the load.

6 Related Work

Automatic scaling of Web applications was previously studied in [14] and [15] for data center environments. In Muse [14], each server has replicas of all web applications running in the system. The dispatch algorithm in a frontend L7-switch makes sure requests are reasonably served while minimizing the number of underutilized servers. Work [15] uses network flow algorithms to allocate the load of an application among its running instances.

6.1 Resource Allocation by Live VM Migration

VM live migration is a widely used technique for dynamic resource allocation in a virtualized environment [8], [12],

Our work also belongs to this category. Sandpiper combines multidimensional load information into a single Volume metric [8]. It sorts the list of PMs based on their volumes and the VMs in each PM in their volume-to-size ratio (VSR). This unfortunately abstracts away critical information needed when making the migration decision. It then considers the PMs and the Ms in the presorted order. The results are analyzed in Section 5 of the supplementary file, which is available online, to show how they behave differently. In addition, their work has no support for green computing and differs from ours in many other aspects such as load prediction. Dynamic placement of virtual servers to minimize SLA violations is studied in [12]. They model it as a bin packing problem and use the well-known first-fit approximation algorithm to calculate the VM to PM layout periodically. That algorithm, however, is designed mostly for offline use. It is likely to incur a large number of migrations when applied in online environment where the resource needs of VMs change dynamically.

6.2 Green Computing

Many efforts have been made to curtail energy consumption in data centers. Hardware-based approaches include novel thermal design for lower cooling power, or adopting power-proportional and low-power hardware.

Our work belongs to the category of pure-software low- cost solutions [10], [12], [14].

7. Conclusion

We have presented the design, implementation, and evalua- tion of a resource management system for cloud computing services. Our system multiplexes virtual to physical resources adaptively based on the changing demand. We use the skewness metric to combine VMs with different resource characteristics appropriately so that the capacities of servers are well utilized. Our algorithm achieves both overload avoidance and green computing for systems with multi-resource constraints.

References

1. M. Armbrust et al., "Above the Clouds: A Berkeley View of Cloud Computing," technical report, Univ. of California, Berkeley, Feb.2009.
2. L. Siegele, "Let It Rise: A Special Report on Corporate IT," The Economist, vol. 389, pp. 3-16, Oct. 2008.
3. P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield, "Xen and the Art of Virtualization," Proc. ACM Symp. Operating Systems Principles (SOSP '03), Oct. 2003.
4. "Amazon elastic compute cloud (Amazon EC2)," http://aws. amazon.com/ec2/, 2012.
5. C. Clark, K. Fraser, S. Hand, J.G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield, "Live Migration of Virtual Machines," Proc. Symp. Networked Systems Design and Implementation (NSDI '05), May 2005.
6. M. Nelson, B.-H. Lim, and G. Hutchins, "Fast Transparent Migration for Virtual Machines," Proc. USENIX Ann. Technical Conf., 2005.
7. M. McNett, D. Gupta, A. Vahdat, and G.M. Voelker, "Usher: An Extensible Framework for Managing Clusters of Virtual Machines," Proc. Large Installation System Administration Conf. (LISA '07), Nov. 2007.
8. T. Wood, P. Shenoy, A. Venkataramani, and M. Yousif, "Black-Box and Gray-Box Strategies for Virtual Machine Migration," Proc. Symp. Networked Systems Design and Implementation (NSDI '07), Apr. 2007.
9. C.A. Waldspurger, "Memory Resource Management in VMware ESX Server," Proc. Symp. Operating Systems Design and Implementa- tion (OSDI '02), Aug. 2002.
10. G. Chen, H. Wenbo, J. Liu, S. Nath, L. Rigas, L. Xiao, and F. Zhao, "Energy-Aware Server Provisioning and Load Dispatching for Connection-Intensive Internet Services," Proc. USENIX Symp. Networked Systems Design and Implementation (NSDI '08), Apr. 2008.
11. P. Padala, K.-Y. Hou, K.G. Shin, X. Zhu, M. Uysal, Z. Wang, S. Singhal, and A. Merchant, "Automated Control of Multiple Virtualized Resources," Proc. ACM European conf. Computer Systems (EuroSys '09), 2009.
12. N. Bobroff, A. Kochut, and K. Beaty, "Dynamic Placement of Virtual Machines for Managing SLA Violations," Proc. IFIP/IEEE Int'l Symp. Integrated Network Management (IM '07), 2007.
13. "TPC-W: Transaction Processing Performance Council," http://www.tpc.org/tpcw/,2012.
14. J.S. Chase, D.C. Anderson, P.N. Thakar, A.M. Vahdat, and R.P.Doyle, "Managing Energy and Server Resources in Hosting Centers," Proc. ACM Symp. Operating System Principles (SOSP '01), Oct. 2001.
15. C. Tang, M. Steinder, M. Spreitzer, and G. Pacifici, "A Scalable Application Placement Controller for Enterprise Data Centers,"Proc. Int'l World Wide Web Conf. (WWW '07), May 2007.