

Network and Energy Aware Load-Balancing via Parallel VM Migration for Data Centers

Kun-Ting Chen, Chien Chen, Po-Hsiang Wang

Department of Computer Science

National Chiao Tung University, Hsin-Chu, Taiwan

quentin2007.cs96g@g2.nctu.edu.tw, chienchen@cs.nctu.edu.tw, btc889874.cs98g@g2.nctu.edu.tw

Abstract—It becomes a challenge to design an efficient load balancing method via live virtual machine (VM) migration without degrading application performance. Two major performance impacts on hosted applications that run on a VM are the system load balancing degree and the total time till a balanced state is reached. Existing load balancing methods usually ignore the VM migration time overhead. In contrast to sequential migration-based load balancing, this paper proposes using a network-topology aware parallel migration to speed up the load balancing process in a data center. We transform the VM migration-based multi-resource load-balancing problem into a minimum weighted matching problem over a weighted bipartite graph. By obtaining the minimum weighted matching pairs from the Hungarian method, we parallel migrate multiple VMs from overloaded hosts to underutilized hosts to reduce the time reaching a load balanced state. In addition, since a major cause of energy inefficiency in a cloud data center is the idle power wasted when servers run at low utilization, cloud operators can turn off those low utilized servers after using VM migration to empty out their workloads. Therefore, we can extend our load balancing algorithm to force low utilized servers to empty out their VMs simultaneously. The experimental results show that our algorithm not only obtains a compatible multi-resource load balancing performance but also improves the balanced time which results in at most a 10% throughput gain by assuming a large batch application running on all VMs. With energy aware load balancing, the power saving is up to 38% without degrading the applications' completion time.

I. INTRODUCTION

Cloud data centers employ virtualization-based technology to consolidate hardware resource usage to provide application hosting for multiple service providers. In a cloud data center, thousands of commodity computers work in parallel to host virtual machines (VMs) that support different applications. The CPU speed, memory size, and network bandwidth of different commodity computers are widely heterogeneous. Besides, a physical host may run multiple services with different types of resource demands. Without proper allocation, the loads of different resources may become unbalanced among different physical hosts. Even with careful resource provisioning at the beginning, due to dynamic arriving and leaving of running workload, the cloud system could still become a load unbalanced state later. Thus, a cloud system may have quite a few overloaded hosts while lots of underutilized hosts are still available.

Load-balancing mechanisms improve system performance by reducing resource contention (e.g., CPU, network

bandwidth, and memory) on overloaded hosts. More importantly, it can prevent higher hardware failure rate in overloaded hosts.

Load-balancing mechanisms are effectively studied in many research areas, including distributed systems, web servers, and cloud systems, to maximize resource utilization and minimize the variance of the load of multiple servers [1] [2] [3] [4] [5] [6] [7] [8] [9] [10] [11]. While most recent research focuses on the VM migration as a mean to achieve load balancing in cloud systems [4] [5] [6] [7] [8] [9] [10] [11], few of them focus on the impact of the total time required to reach system balance. Normally, the existing load balancing algorithms search for a VM to instantiate a VM migration from overloaded hosts to under-loaded hosts. They usually select and start the next VM migration when the previous migration completely terminates. These schemes aim at achieving an egalitarian state without considering the waiting time for other hosts to offload their workload. This paper calls such schemes “sequential migration-based load balancing methods.” These methods could lead to all overloaded hosts taking a long time to offload their workload. Consequently, the applications running on those overload hosts will contend for the resources for a longer time. Therefore, users could experience application performance degradation.

In this work, we study the multi-resource load balancing problem for a cloud data center with heterogeneous hardware capacity. Compared with sequential migration-based load balancing methods, this paper proposes a parallel VM migration approach that focuses on minimizing the joint multi-resource imbalance and the time it takes to reach a balanced state. In addition, the migration delay due to the data center networks' architecture is considered, especially if the migration pairs are between machines with a large hop distance. In order to improve the application performances, this paper minimizes the time till a balanced state is reached by migrating VMs between independent pairs of hosts while considering the underlying network topology. To be specific, this paper models the load balancing problem with a weighted bipartite graph (WBG). The over- and underutilized hosts in a data center are selected as two disjoint sets of vertices: Trigger (*TR*) node and Non-Trigger (*NTR*) node sets. A physical host is said to be a trigger node if its resource utilization exceeds a system threshold on any dimensions of resources; otherwise, it is a non-trigger node. Each edge between the two sets is designated with a weight according to the multi-resource requirements of VMs, the various resource capacities of the hosts, and the network cost between two hosts. By solving the minimum

weighted matching problem using the Hungarian method [12], this paper could migrate a set of VMs between over- and underutilized hosts in parallel.

In addition, to maintain a large-scale computational environment, a cloud data center needs a lot of hosts running persistently and a cooling system to maintain a working temperature. In general, the increased number of running hosts and the cooling system consume large amounts of power. However, a major cause of energy inefficiency in a cloud data center is the idle power wasted when some servers run at low utilization. Recent researches such as [13] [14] are dedicated to saving energy by moving VMs from underutilized hosts to other hosts and turning off those hosts without any workloads. These researches usually transform the energy saving issue to a classical bin-packing problem. Physical hosts are viewed as individual bins and the workload as an item. The energy saving problem is then modeled as the problem of filling all items into bins while the goal is to minimize the number of bins.

In such a straightforward modeling, the load balancing problem still exists. First, due to heterogeneous machine capacity and various VM workloads, the bin-packing could still involve migrating VMs out of overloaded hosts. Since item size (i.e. size of VM resource requirement) is varied, it may still involve selecting an appropriate host to receive a migrating item (i.e. VM) instead of simply migrating any VM into any host to yield a better packing result. Therefore, as mentioned earlier, migrating many VMs out of hosts could exhibit a sequential migration. As a result, it may take a long time for selected hosts to empty out their workload before they can be turned off. This paper augments our parallel migration-based load balancing algorithm to migrate VMs simultaneously to quickly turn as many hosts off as possible while preserving the system load balancing state. In order to turn off selected hosts, our algorithm marks these hosts as triggered nodes and they remain triggered nodes until all tenant VMs have been migrated.

The experimental results show that our algorithm achieves a compatible multi-resource load balancing while improving the balancing time, which results in at most a 10% running application's throughput gain compared to conventional algorithms which perform load balancing in a sequential manner. With energy aware load balancing, the amount of power saved is up to 38% without degrading the applications' completion time.

The rest of this paper is organized as follows. Section II describes related works. Section III elaborates the motivation behind this work. Section IV describes the network-aware bipartite matching load balancing algorithm in detail. Section V describes the network-and-energy-aware load-balancing algorithm. Experimental results of the simulation and an evaluation will be presented in section VI. Finally this paper concludes with some remarks in section VII.

II. RELATED WORKS

A. Load Balancing Methods

A number of researches have proposed load balancing methods [1] [2] [3] [4] [5] [6] [7] [8] [9] [10] [11] for cloud data center. Most of these works focus on balancing the resource requirement of VMs on physical machines. Based on

their algorithms, they migrate either single or multiple VMs in each round of the load balancing procedure. These load-balancing schemes for cloud computing measure the load on physical hosts in different ways. However, none of their efforts address the multi-resource load balancing issue [1] [2] [3]. Zhao and Huang [2] use the number of VMs of a host as their load measurement. Recent study [3] mimics the animal behavior of honeybees foraging and scouting to harvest foods for load-balancing. Though the scheme works for a dynamic workload, they assume each host is homogeneous and deal with single resource only. And yet [4] considers the multi-resource demands of VMs and heterogeneous capacity in each host. They propose a load balancing method called *VectorDot (VD)*. Their goal is to bring overloaded hosts below their threshold by migrating one or more VMs sequentially. They extend Toyoda heuristic for a single knapsack problem to solve a multidimensional knapsack problem. *VD* treats a VM in an overloaded host as an item to be placed into an appropriate node. Then, this VM is migrated to the selected destination node. This procedure is repeated until no more overloaded hosts. From a practical perspective VM migration takes a while to finish. Thus migrating VMs in a one-by-one manner could take a long time till system balancing, and thereby exacerbating the running applications that are hosted on overloaded machines.

Recent works that consider VM migration overhead for load balancing can be found in [7] [8] [9] [10] [11]. Since a cloud system involves cloud operators and network operators with different goals in mind, [7] studies the problem of making both objects meet. At one hand, the cloud operators wish to assign a number of VMs to some selected hosts for the purpose of meeting the VM scheduling deadline. Thus, the cloud operators wish to pair for each VM with a destination host that minimizes the migration time. On the other hand, the network operators wish to minimize the network bandwidth. To make both ends meet, [7] models a balanced stable matching problem for both requirements of cloud operators and network operators. VM migration overhead is derived from the network topology and the transfer memory volume. However, their matching does not consider the bandwidth of NIC sharing at the destination host. Thus they could migrate several VMs from multiple distinct hosts to an identical host which could degrade into the sequential migration. In contrast, we consider finding independent pairs between the end hosts for load balancing via parallel migration.

Recent works [8] [9] conduct a number of experiments on the cloud testbed to study the factors that influence VM migration time and impact of live migration on application performances. In [8], they show that VM migration could contend resource with the running workload on the source and destination host, especially as the number of concurrently migrating VMs between one source and one destination host increases. Hence they study the problem of selecting a number of VMs to migrate that minimize the resource contention overhead. [8] models a VM migration as processes that run at source and destination hosts. They show that by carefully assigning VM they could reduce VM migration overhead. Though, it is difficult to formulate the impact of migration time given different combination of jobs and VM migrations running on machines with various resource configurations. As

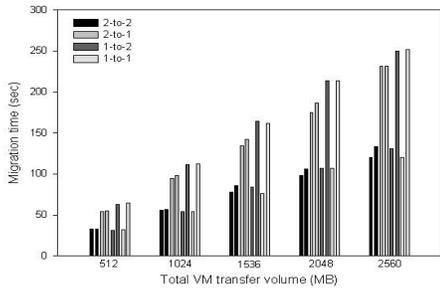


Fig. 1. Migration time vs. VM memory size

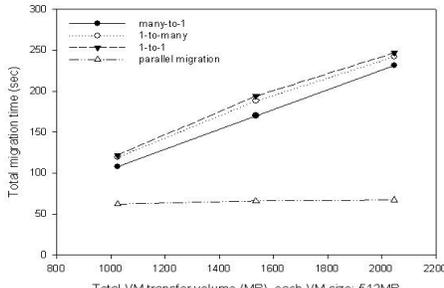


Fig. 2. Total elapsed migration time vs. number of VMs

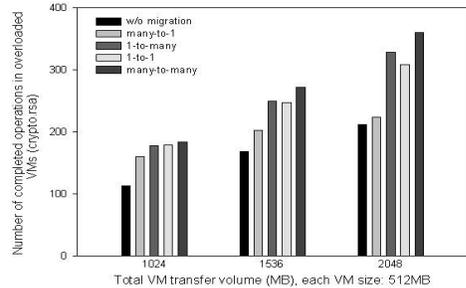


Fig. 3. Number of completed operations vs. number of VMs

the cloud system gets larger, the resource contention model would become more difficult to build. [9] shows that the destination host CPU reserved for the migration process has little to do with migration time. The source host CPU impacts the migration time when the host is overloaded, especially when its memory dirty rate is high. However, they only suggest a direction to balance multi-resource load without relating it to the performance overhead incurred by concurrent VM migrations. In contrast to [8] [9], this paper performs an experimental study of live migrations and application throughputs with multiple, concurrent VM migrations (in next section). These experiment results provide a motivation to a new parallel migration-based load balancing algorithm.

Recent studies [10] [11] formulate the load balancing problem with migration overhead being the distance measured by the number of hops in the Internet. Nevertheless these schemes often do not consider the heterogeneous capacity of multi-resource in each host. In addition, the concurrent migration approach they adopt still has a chance to migrate multiple VMs to an identical destination host or from a single source to multiple distinct destinations.

B. Energy Issues

Considering the fact that a large percentage of power consumption comes from the static power from those underutilized machines, this paper considers migrating all VMs out from those hosts. Thus those hosts can be switched off to reduce total power consumption in a cloud system. The power consumption of a running host consists of two parts: static power and dynamic power. Studies show that on average an idle server consumes approximately 70% of the total power consumption. To deal with energy issues, this paper adopts the power model described in [14] in equation (1).

$$P(l) = k \times P_{\max} + (1-k) \times P_{\max} \times l \quad (1)$$

P_{\max} denotes the maximum power consumed in a host when its CPU utilization is 100%; k is the fraction of static power and is approximately 0.7; and l is the current CPU utilization which determine the dynamic power consumption.

In [13], the energy issue is transformed to a traditional bin-packing problem. Physical hosts are viewed as individual bins and the workload as an item. The energy saving problem then is modeled as the problem of filling all items into bins while minimizing the number of bins. Although this paper takes a similar approach to reducing the number of running hosts, it further focuses on preserving the system load balanced state.

III. MOTIVATION

Even through the modern cloud administrators migrate multiple VMs concurrently, if the hosts of migration pairs do not choose carefully, concurrent migration still has a chance to degrade into sequential migration. The worst case circumstance in load balancing is to migrate many VMs from different hosts to one under-utilized host (many-to-1) concurrently. In order to turn off a host in energy saving, one common example occurs when we migrate all VMs from one host to many hosts (1-to-many) concurrently. In this section, we measure the real application performance and compare parallel migration with a range of many-to-1 and 1-to-many concurrent migration. The experiment is conducted on a real testbed based on Xen [16]. In this system, we reserve the network bandwidth as 100 Mbps at each host NIC for VM migration. Each VM runs crypto.rs application in the SPECjvm [17]. VMs serving the same application type are allocated with identical resources. During live migrations, applications running on VMs persist in execution with small interruption time compared with non-live migration [16]. We first measure the time taken for each individual VM during concurrent migration by observing different combinations of VM pairs. Then we observe the application performance for an increasing number of concurrently migrating VMs. A conclusion that leads to the work in this paper follows thereafter. Each evaluation is an averaged result for 10 times.

Fig. 1 shows the individual VM migration time for migrating two VMs with increasing total migrating memory volume. Consistently with [8], as each VM memory size increases from 256 to 1280 MB (where the total migrated size increases from 512MB to 2048MB, respectively,) VM migration time takes a linearly increasing proportional to transfer memory volume. As two distinct hosts migrate to a single host (i.e., 2-to-1), it suffers from a linearly increasing time for both VM migrations since the network capacity at the destination host's NIC is shared by the two migrations. When migrating from a single host to two distinct hosts (i.e., 1-to-2), one VM does not start migrating until another finishes migrating. This is the mechanism implemented by the Xen platform¹ which would start VM migration after the previous one finishes, that behaves the same as sequential migration in

¹Xen provides migrating modules for synchronous and asynchronous modes. The only difference is whether the migrating thread immediately returns to user thread or returns until migration completes. Both implementations migrate VMs in sequential manner even they are issued concurrently.

the case of migrating two VMs from one host to the other host. As opposed to the above cases, the parallel VM migration scheme (i.e., 2-to-2) achieves the smallest migration time without suffering from additional network delays.

Given that migrating two VMs could degrade into sequential migration in the 2-to-1 and 1-to-2 cases, we would like to know the performance impact of migrating multiple VMs concurrently in a larger scale, especially in an overloaded system environment. We assume that an overloaded environment consists of overloaded hosts and idle hosts. In the following experiments we assume that all the VMs run on those overloaded hosts. In such an overloaded environment, the workload on one VM contends resources with all other VMs resident on the same host. Thus the time until a VM is migrated to an idle host could affect the application performance running on that VM.

We study the impact of total migration time of different number of migrating VMs on the performance of workload on the migrating VMs, where the total migration time denotes the longest completion time for migrating a number of VMs. In the many-to-1 case, we concurrently migrate 2, 3 and 4 VMs each from 2, 3 and 4 overloaded hosts to one idle host. On the contrary, the 1-to-many case exhibits that 2, 3 and 4 VMs in one overloaded host are migrated to 2, 3 and 4 idle hosts respectively. The 1-to-1 case let 2, 3 and 4 VMs migrate from an overloaded host to an idle host. Parallel migration independently migrate 2, 3, and 4 VMs each from 2, 3 and 4 overloaded hosts to 2, 3 and 4 idle hosts. Each VM migration is instantiated by our controller 1 minute after the workload starts running. During VM migration, we measure the application performance obtained on migrating VMs. In order to measure the impact on the running application, the iteration duration parameter of the crypto.rs application is set to last after all VMs finish migrating. This parameter ensures that the VM continues to demand resources during the VM migration. The baseline performance is approximately 108 operations under baseline single core CPU at 2.67GHz.

Fig. 2 shows the total migration time under 2, 3 and 4VM migrations. As the number of VM migrations increases, the completion time of migration increases as in many-to-1, 1-to-many and 1-to-1 cases. However, the total migration time of parallel migration remains almost the same. It's the smallest among the other cases. Consistently with [8], as the total transfer memory size of concurrent VM migrations increases, the total migration time also increases. Given that the total transfer volume of memory between two hosts is fixed, they observe that the total migration time increases as the number of concurrent VM migrations increases. This is due to the resource contention between those migration processes and running workload on the source and destination host. In order to improve application performance, they study the VM assignment problem by selecting different pairs of VM migrations to lower the impact of resource contention.

Fig. 3 shows the total number of completed operations of crypto.rs observed on migrating VMs. The case of a single host to multiple distinct hosts (1-to-many) completes 4%, 9%, and 10% fewer operations than parallel migration, as the number of concurrently migrating VMs increase from 2, 3 to 4. Although the many-to-1 case experiences a smaller total migration time than that of the 1-to-many as shown in Fig. 2,

the number of completed operations of many-to-1 decreases more severely than that of 1-to-many with 15%, 34%, and 61% fewer operations done than parallel migration. Further, the many-to-1 case is also close to the worst case without migration (i.e. no load balancing). The 1-to-many case completes more operations than the case of many-to-1 where it suffers from a longer migration delay for each individual VM as shown in the 2-to-1 case in Fig. 1. Among them, parallel migration shows the highest workload completion rate. These observations suggest that maneuvering parallel VM migrations is feasible to improve application performances.

IV. NETWORK-AWARE BIPARTITE MATCHING LOAD-BALANCING ALGORITHM

Load balancing for multi-resource requirements while considering heterogeneous resource capacity has been a challenging problem, especially in a large scale cloud hosting environment. In this paper, we study the load balancing problem for a cloud data center. Especially, we try to reduce the time taken for a cloud data center to reach its load balance state. Since a live migration of VM takes some time to accomplish, sequentially migrating VMs from trigger nodes to non-trigger nodes may take a long time to offload workload for overload hosts. Further, the migration delay between two hosts varies according to the underlying network architecture and the transfer volume of VM memory. Thus, we propose a *Network-Aware Bipartite matching (NABM)* load balancing algorithm. This paper first transforms the load-balancing problem into a minimum weighted matching problem. According to the minimum weighted matching obtained from the Hungarian method, this paper migrates VMs from overloaded hosts to underutilized hosts in parallel. Furthermore, this paper adds the hop distance between source and destination hosts to reflect impact of the network to the VM migration time. By reducing the VM migration time, it can shorten the time for a cloud system to reach its load balanced state.

A. Preliminary

The notations for use in *NABM* will be defined as follows. Let $H = \{h_1, h_2, \dots, h_m\}$ and $V = \{vm_1, vm_2, \dots, vm_v\}$ denote a set of hosts and VMs in a cloud system, respectively, with an existing allocation $A: V \rightarrow H$ where $A(vm_i)$ denotes the host where vm_i resides. The resource consumption for a VM vm_β is $S_\beta = \{s_\beta^i | i \in \{1, 2, \dots, n\}\}$. The input to our load balancing algorithm includes a capacity value, a current utilization value and a suggested threshold fraction between 0 and 1. Let $C_a = \{c_a^i | i \in \{1, 2, \dots, n\}\}$ and $U_a = \{u_a^i | u_a^i \in [0, 1], i \in \{1, 2, \dots, n\}\}$ denote the vectors of capacity and resource utilization along n dimensions of resources for a Host h_a , respectively. $T = \{t_i | t_i \in [0, 1], i \in \{1, 2, \dots, n\}\}$ denotes the vector of system threshold along each dimension of resources. Our load balancer will maintain the usage of the resources below those thresholds. Moreover, the available capacity for host h_a in any dimension of resources can be represented as in (2).

$$F_a = \{f_a^i | f_a^i = (1 - u_a^i) \cdot c_a^i, i \in \{1, 2, \dots, n\}\}. \quad (2)$$

B. Bipartite Matching Load-Balancing

Similar to [4], *NABM* collects multi-dimensional resource utilization information of hosts and resource requirements of VMs as input. In the real testbed deployed with Xen, we send XML-RPC requests to collect load information from any host in the same server farm through a control domain. Xen implements multiple control domains in charge of sending and receiving end hosts' control messages. To construct a weighted bipartite graph (*WBG*), *NABM* leverages three decision policies: participation, candidate selection, and location, as defined in [17]. The participation policy decides which hosts are involved in the load balancing process. In this paper, we specify two node sets for over- and under-loaded hosts. They can be defined as a triggered set: $TR = \{h_a \mid h_a \in H, \exists i(u_a^i > t_i), i \in \{1, 2, \dots, n\}\}$ and a non-triggered set: $NTR = \{h_a \mid h_a \in H, h_a \notin TR\}$, respectively. The triggered set is a set of hosts which resource utilization exceeds the system threshold in any dimension of resources. The hosts that are not marked as the triggered nodes belong to the non-triggered set. These two sets together form two vertex sets for the weighted bipartite graph used in *NABM*.

Both the candidate selection policy and the location policy play the intermediate roles of generating edges and edges' weight for the *WBG*. The candidate selection policy chooses the VMs to transfer to alleviate an overloaded hot spot. A VM is a candidate if the removal of this VM turns a triggered host into a non-triggered host. However, if the removal of any VM in a trigger node couldn't turn the trigger host into a non-trigger host, then all tenant VMs in this triggered host are chosen for the candidate set. For a triggered host h_a , we identify the candidate set as (3).

$$CE_a = \begin{cases} \left\{ \begin{array}{l} vm_\beta \mid A(vm_\beta) = h_a, h_a \in TR \\ \forall i(\max\{u_a^i - s_\beta^i, 0\} \leq t_i), i \in \{1, 2, \dots, n\} \end{array} \right\} \\ \text{if } \exists \beta, \forall i(\max\{u_a^i - s_\beta^i, 0\} \leq t_i), i \in \{1, 2, \dots, n\}, \\ \left\{ \begin{array}{l} vm_\beta \mid A(vm_\beta) = h_a \end{array} \right\}, h_a \in TR \\ \text{if } \forall \beta, \exists i(\max\{u_a^i - s_\beta^i, 0\} > t_i), i \in \{1, 2, \dots, n\} \end{cases} \quad (3)$$

The location policy selects the destination host to which the candidate VM previously selected is migrated. A Host h_a in *NTR* is said to be feasible if after migration its capacity constraint C_a is not violated. For any $vm_\beta \in CE_a$ the location policy determines the set of potential destination hosts from the feasible hosts in *NTR* set. Since we attempt not to increase the number of triggered nodes after the migration of VMs, a feasible node can be a potential destination node only if the VM migration wouldn't turn it into a triggered node. The set of potential destination hosts for any candidate VM $vm_\beta \in CE_a$ is shown in (4).

$$EP_\beta = \left\{ \begin{array}{l} (vm_\beta, h_a) \mid \forall i(\max\{u_a^i + s_\beta^i / c_a^i, 0\} \leq t_i), \\ i \in \{1, 2, \dots, n\}, A(vm_\beta) \neq h_a, \\ h_a \in NTR. \end{array} \right\}. \quad (4)$$

This policy then associates a cost with each (vm_β, h_a) . The cost is denoted as the sum of the resource requirements of vm_β divided by the residual capacity on n dimensions of the potential destination host h_a . It indicates that a host with scarce remaining resources exhibits a higher cost. The formal load balancing cost function for a pair of migration $(h_{src}, h_{dst}, vm_\beta)$ can be defined in equation (5).

$$Cost_{LB}(h_{src}, h_{dst}, vm_\beta) = \sum_{i=1}^n \frac{s_\beta^i}{f_{dst}^i} \quad (5)$$

where vm_β resides on h_{src} which belongs to *TR*, h_{dst} belongs to *NTR*, s_β^i denotes resource demands of vm_β over each resource dimension i , f_{dst}^i denotes the available resource of h_{dst} , and n denotes the total number of machine resources. For the candidate VM $vm_\beta \in CE_a$, the location policy selects a destination host out of the potential destination host set according to cost. We can apply one of the strategies in the best-fit, first-fit, worst-fit and the relaxed-best-fit [4] to determine the destination host. With the best-fit, we select a destination host with the smallest cost. Thus, the potential destination host with higher cost is less likely to be selected as a destination for the candidate VM. However, it requires a linear search from a set of potential destination hosts. Thus, the relax-best-fit exhibits a better search time by investigating the smallest cost from a much small set of the potential destination hosts which are randomly chosen from the original potential destination hosts. With the relaxed-best-fit, it also reduces the chance of picking the same best potential hosts among different VMs and thereby increases the number of valid edges in weighted bipartite graph. *NABM* applies only relaxed-best-fit (*RBF*) as in [4].

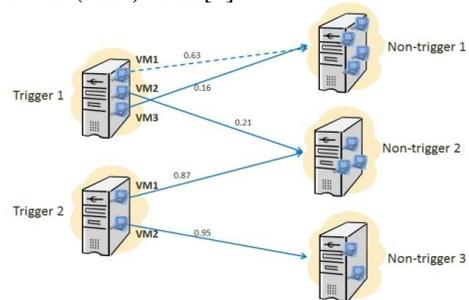


Fig. 4. An illustration of our WBG

For a candidate VM in CE_a , relax-best-fit select a destination host to form an edge in EP_a . Since multiple candidate VMs residing at the same host could select the same destination host, it causes multiple edges between triggered and non-triggered nodes. We discard all such edges except for the edge with the smallest weight. Fig. 4 shows an illustration of *WBG* with two triggered nodes and three non-triggered

nodes. Since candidate VM 1 and VM 2 in the Trigger node 1 select the Non-Trigger node 1 as their potential destination host, two potential edges associated with the weights are formed between the Trigger node 1 and Non-Trigger node 1. However, the solid edge with the less weight than the dotted edge, the dotted edge is discarded from our final weighted bipartite graph. For any node pair, the edge set is computed as (6).

$$E = \left\{ \begin{array}{l} (h_{src}, h_{dst}) \mid h_{src} \in TR, h_{dst} \in NTR, \\ \exists vm_\beta (\forall vm_j (\beta \neq j, \\ Cost_{LB}(h_{src}, h_{dst}, vm_j) \geq Cost_{LB}(h_{src}, h_{dst}, vm_\beta))), \\ vm_\beta, vm_j \in CE_{src} \end{array} \right\}. \quad (6)$$

Finally, we construct a weighted bipartite graph: $G=(V,E,W)$ where $V=TR \cup NTR$ and E is the edge set obtained from (6), W is the cost function defined in (5).

After constructing the *WBG*, *NABM* applies the Hungarian algorithm [12] to solve the minimum weighted matching problem. Since the Hungarian algorithm solves the instance of *WBG* with perfect matching, *NABM* must ensure that the *WBG* created above has an equal number of nodes for *TR* and *NTR*. Thus *NABM* algorithm adds pseudo nodes to either *TR* or *NTR*, whichever one has the smaller number of nodes, until they have the same number of nodes. Pseudo edges associated with a significantly large value also are added from a pseudo node to all nodes in the other set. Based on the outcome of each match from the Hungarian algorithm except for the pairs connected with pseudo edges, *NABM* migrate VMs from triggered nodes to their corresponding non-triggered nodes in parallel for balancing load. Compared with sequential migrations described earlier, *NABM* not only migrates VMs concurrently but also prevents from exhibiting the migration overhead due to contention at the end hosts' NIC.

NABM iterates the steps of constructing a weighted bipartite graph, finding its minimum weighted bipartite match, and parallel migrating the corresponding VMs until either there's no triggered node left or the match consists of nothing but the pseudo edges. We omit the pseudo-code of *NABM* due to space limit.

C. Network-aware Extension

In addition to the weighted bipartite match that increases the total number of VM migration pairs in each round, the length of the duration between each parallel migration round also plays an important role in reducing total time till balance load. The time between each round depends on the longest migration time among all migrations. Since the network hop distance between two hosts will affect the migration time [7], we further consider network hop distance in the location policy. Like [6], we assume that the traffic patterns rarely change in the production data center. Even though the traffic distribution could be highly uneven, a balanced traffic distribution could still be obtained in the data center network. For example, the cloud operators could reassign traffic flows among communicating VMs via equal-cost-multiple-path (ECMP) [18] or a centralized network controller such as NOX [19].

Under even traffic distribution assumption, VM migration between end hosts with long network path exhibits a higher probability of long network delay. According to (7), it indicates that the migration time is proportional to the size of VM transfer memory and the communication distance while being disproportional to the bottleneck of end hosts' NIC bandwidth. *NABM* extends the previous cost function in (5) to account for migration cost which is related to the network hop distance and the size of VM transfer memory as defined in (7).

$$Cost_{mig}(h_{src}, h_{dst}, vm_\beta) = \frac{s_\beta^{mem}}{\min(f_{src}^{bw}, f_{dst}^{bw})} \cdot D(h_{src}, h_{dst}). \quad (7)$$

where vm_β resides on h_{src} which belongs to *TR*, h_{dst} belongs to *NTR*, s_β^{mem} is memory volume of vm_β , f_{src}^{bw} and f_{dst}^{bw} denote the available bandwidth for VM migration in source and destination NIC, respectively, and $D(i, j)$ is network hop distance between Host i and j . An example of $D(i, j)$ is to account for the hop count between end hosts in some data center network architectures such as fat-tree [20] and VL2 [18]. This paper implements *fat-tree* and *VL2* as the underlying network architectures. The extension of *NABM* takes a normalized cost of (5) and (7) as in (8).

$$Cost(h_{src}, h_{dst}, vm_\beta) = (1-\alpha) \cdot Cost_{LB}(h_{src}, h_{dst}, vm_\beta) + \alpha \cdot (\rho \cdot Cost_{mig}(h_{src}, h_{dst}, vm_\beta)) \quad (8)$$

where ρ is an adjustable parameter for normalizing the migration cost with load balancing cost; and the α is the parameter for adjusting the relative importance between the two costs. For example, it considers no network cost if α is zero. During constructing the *WBG*, the edge weight is substituted with $Cost(h_{src}, h_{dst}, vm_\beta)$ in (8) in the location policy. This paper conducts a number of event-driven simulation analyses to study how parameter α affects the time till system is balanced, the mean VM migration time between each round, and the application performances in section VI.

V. ENERGY-AWARE BIPARTITE MATCHING LOAD-BALANCING

Power consumption constitutes one of the major costs for the data center. Thus it is imperative to save energy by turning off some unnecessary machines. To turn off those machines, all tenant VMs need to migrate to other active machines. In order to achieve energy saving while maintaining performances of running jobs, one has to carefully deal with both the impact of migration time and uneven workload distribution which will degrade application performance. This section first defines energy saving as a machine consolidation problem. Then we extend the previous *NABM* load balancing algorithm to turn off machines.

The problem of machine consolidation is defined as follows. Given the sets of physical hosts H , VMs V , host capacity C , available capacity F , and the average target system utilization $TU = \{tu_1, tu_2, \dots, tu_n\}$, where n is the number of dimensions of resources, this paper wants to estimate the minimum number of hosts to provision for the total resource requirement of VMs such that the average system utilization is closer to TU . Aside

from aiming to turn off as many machines as possible like first-fit-decreasing (FFD) bin-packing algorithm in [21], we further focus on maintaining the load-balanced state. We extend the *NABM* algorithm by adding an energy aware policy and name it as *Network- and Energy-aware Bipartite Matching (NEABM)* method.

NEABM first searches for potential hosts to turn off. Then *NEABM* applies parallel migration to empty VMs out of those hosts simultaneously. Specifically, like *NABM*, *NEABM* transforms the machine consolidation problem to a minimum weight matching problem over a *WBG*. In addition to choosing the hosts whose resource utilization exceeds the system threshold of any dimension of the resources as triggered nodes, *NEABM* would force some target machines to remain as triggered nodes until they empty out all resident VMs in order to turn off them.

According to the target system utilization set by the operators, *NEABM* approaches this target utilization by finding as many hosts as possible to turn off. Similar to other multi-capacity bin-packing algorithms which sort the bins in either increasing or decreasing order according to the squared sum, maximum load or product of load [21], *NEABM* searches hosts to turn off in either increasing or decreasing order of resource utilization of hosts. For our multi-resource case, the sorting order will be determined by one resource at a time. If the first resource utilizations of two hosts are the same, second resource utilizations are compared and so on. The increasing order of utilization for multi-resource can be defined as in (9):

$$h_a \prec h_b \triangleq (u_a^1 < u_b^1) \vee ((u_a^1 = u_b^1) \wedge (u_a^2 < u_b^2)) \vee \dots \vee (u_a^k = u_b^k) \wedge (u_a^{k+1} < u_b^{k+1}), k \leq n \quad (9)$$

where the notation \prec denotes the turning off priority order for any two hosts.

Given target utilization constraint TU , let PM^i denote the total available capacity for all VMs in each dimension of resources i as in (10).

$$PM^i \leftarrow \sum_{\forall h_a \in H} tu^i \cdot c_a^i, i \in \{1, 2, \dots, n\} \quad (10)$$

Let VM^i denote the system requirements of resources i as the sum over ith resource requirement of all VMs as in (11).

$$VM^i \leftarrow \sum_{\forall vm_p \in V} s_{\beta}^i, i \in \{1, 2, \dots, n\} \quad (11)$$

Let SL denote a set of candidate hosts to turn off, if they can be removed without violating the total available capacity under target utilization constraint. Let PM_{imp}^i denote a variable. Initially the set SL is empty and PM_{imp}^i is set to PM^i . Depending on either the increasing or decreasing order of host's utilization to turn off the hosts, a Host $h_a \in H$ which is either least utilized or greatly utilized will be tested to see if turning off h_a would satisfy $PM_{imp}^i - c_a^i > VM^i$ for each dimension of resource i . If it's true, the Host h_a will be added to SL and update the vector as $PM_{imp}^i \leftarrow PM_{imp}^i - c_a^i$. This iteration stops when removing a host will violate the total available

capacity under target utilization constraint for any dimension of resources.

Although *NEABM* could meet the target utilization by finding a number of potential hosts to turn off, the number of actual hosts being turned off could decrease as the target utilization gets larger. When the target utilization gets larger, the number of potential hosts to turn off also gets larger. It results in a much smaller room to accommodate the variance of multiple resources. Thus, it becomes difficult for the VMs that are resident on the triggered nodes to select a destination host without violating the system threshold if the system threshold equal to target utilization. Therefore, we take the system threshold t_i as $tu_i + \varepsilon$, where ε is a small number to give more room for *NEABM* to accommodate multi-resource variance. The value of ε is determined by the standard derivation of resource usage under *NABM*.

Given a SL , *NEABM* defines a new trigger node set TR as

$$TR' = \{h_a \mid (h_a \in SL) \vee (h_a \in TR)\} \quad (12)$$

Hosts not appearing in the TR set are in the NTR set. Throughout the load balancing process, *NEABM* ensures that those triggered nodes from the SL remain triggered nodes unless all resident VMs are removed. *NEABM* iteratively runs the parallel VM migration. It shuts down a machine only if there are no resident VMs. This process continues until it finds no valid migration edges from the match.

VI. EXPERIMENTAL EVALUATION

This section studies the effectiveness and scalability of the *NABM* and *NEABM* algorithms by comparing them with the *VectorDot (VD)* [4] using the *NetworkCloudSim* [22] simulation platform.

A. Experimental Settings

To vary the system scale we simulate with 50~1050 hosts and 157~3244 VMs. The baseline host is equipped with a 2.8GHz (approximately 12,000 MIPS) quad-core CPU, 4 GB memory, and 1024 Mbps NIC. The multi-resource and heterogeneous capacity of each host is generated as baseline capacity*(1±heterogeneous degree.) The heterogeneous degree is 0.2. Each VM consumes 12%~25% multi-resource demands of the baseline host capacity. Each VM runs a large-scale Bag-of-Task (*BoT*) application which is independent without need to communicate with other VMs. The workload consists of 287,712,000 million instructions which take approximately 9 hours to finish for a VM with 8880 MIPS CPU. Typical examples of such workload are biological computation, data mining, and scientific engineering applications. The 37% of hosts are over provisioned, which are considered as overloaded hosts. The average system utilization along each resource dimension is approximately 57%. In order to accommodate the variance of multi-resource load balancing, the triggered node threshold is set to 75% along each dimension of resources.

For the case of energy saving, the average target utilization for power saving is from 60% to 80%. In order to study the impact of the tolerance ratio on power saving, the tolerance

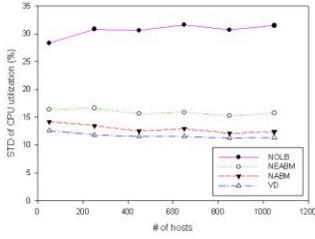


Fig. 5(a). STD of CPU utilization

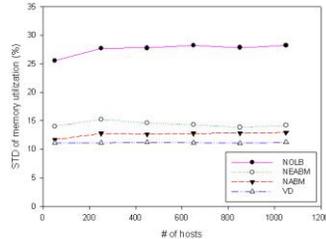


Fig. 5(b). STD of memory utilization

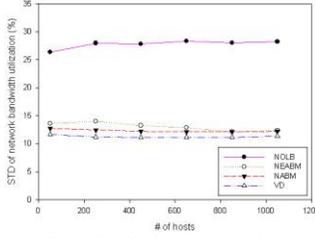


Fig. 5(c). STD of network utilization

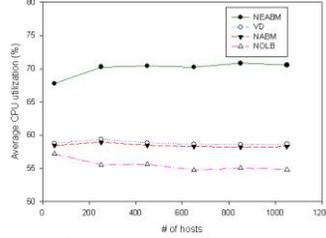


Fig. 6(a). Average CPU utilization

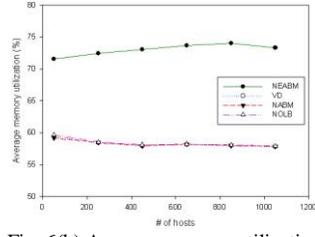


Fig. 6(b). Average memory utilization

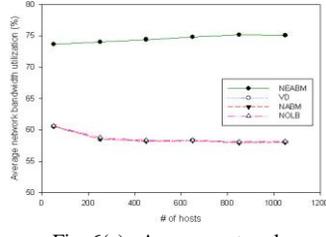


Fig. 6(c). Average network utilization

ratio ε for power saving is set to 0.15 and 0.2 where ε is a small number to give more room for *NEABM* to accommodate multi-resource variance. While $\varepsilon = 0.15$ and 0.2, the triggered node threshold will be set from 75% to 95% and from 80% to 100% for power saving, respectively. In order to study the effect of power saving, we take power equation in (1) to measure the power consumption of the host, where k is set to 0.7 and recurrent CPU utilization l is obtained at runtime. P_{max} is set to 250 W. The power consumption of the system is obtained by the sum of power consumption of those hosts not being turned off over a fixed number of hours (kwh).

For the network parameters, we assume that the workload in VM consumes on average 57% of host NIC bandwidth. This bandwidth is reserved and not used by VM migrations. The network topologies fat-tree and VL2 used in our simulation are built according to [6]. The network parameter α is set from 0 to 1 with 0.2 increment. While $\alpha = 0$, *NABM* would not consider any impact of the network topology. We take a modest value of network weight α equal to 0.6 for the simulation results in Fig. 5, 6, and 7. The available bandwidth of core and aggregate switch reserved for VM migration is set to 1184 Mbps in the following experiments. The parameter ρ is determined as follows. We observe that load balancing cost $Cost_{LB}$ in general is in the range (0, 1.8]. In order to combine the costs of load balancing and migration, we normalize the migration cost to be in the range (0, 1.8]. Thus we take ρ as 1.8 divided by the maximum value of migration cost which is in the range [1000, 2800] in our experiments. For the rest of simulation, we fix $\rho = 1.8/1500$.

All the performance results are an average of 10 repeated runs and obtained at the moment of system convergence. We use the following metrics to evaluate *NABM*, *NEABM*, and *VD*: the system utilization and standard deviation of multi-resources of the hosts, the time it takes for the load balancing process to reach a balanced state, the number of VM migrations taken during the load balancing process, the total application completion time, the number of hosts being turned off, and the power consumption of the system.

B. Experimental Results

Fig. 5(a), (b), and (c) and Fig. 6(a), (b), and (c) show the average standard deviation of resource usage of CPU, memory, network bandwidth, and their average resource utilization, respectively, for different numbers of hosts. All of *NABM*, *NEABM*, and *VD* exhibit a lower resource standard deviation than *NOLB*, which does not do any load balancing. For the case of energy saving, the average target utilization for power saving is 75%. Therefore, *NEABM* can turn some machines off to reduce power consumption of the system. As a result, *NEABM* has a higher utilization along each dimension of resources than *NABM* and *VD*. *NEABM* achieves a moderately higher resource standard deviation than *NABM* and *VD* while few numbers of hosts are operating. The differences of average standard deviation between *NABM* and *VD* are less than 5% which means that *NABM* still maintains a good degree of imbalance at balanced states compared with *VD*.

Fig. 7 (a) shows that the time for *NABM* and *NEABM* to get the system to a balanced state is much less than the time for *VD* especially in a data center with a large number of hosts. For *VD*, the balance time increases along with increases of the number of VM migrations. This is because *VD* sequentially migrates one VM at a time. In contrast, *NABM* and *NEABM* concurrently migrate VMs according to the number of independent matching pairs. *NEABM* takes more time than *NABM* to reach system balance. Fig. 7(b) shows that *NEABM* takes more migrations until all resident VMs of the selected hosts are emptied. While the average number of VM migrations per round for *VD* is constant (equal to 1), the average number of VM migrations per round for *NABM* and *NEABM* increases as the number of hosts increases as shown in Fig. 7(c). This is because as the number of hosts increases, the matching pairs in the minimum weighted matching also increases. Fig. 7(d) shows that the total application completion time of the *NABM* algorithm is better than *VD* by 1%, 3%, 5%, 7%, 9%, and 10%, respectively, under different numbers of hosts. This is because *NABM* handles multiple overloaded hosts at a time, decreasing the load on those hosts and thus decreasing the degree of resource contention amongst the applications running on the VMs. In contrast, *VD* handles only one trigger node at a time and the algorithm needs to spend a longer time to decrease the amount of trigger nodes. When the number of hosts is increased, there are more VMs in the hotspot competing for the resources. For *VD*, it will take more rounds (i.e. time) to alleviate the hot spot. Besides, *NABM* achieves a slightly better throughput than *NEABM* because *NABM* takes less time to reach a balanced state.

Fig. 8 shows the VM migration time and the time it takes to reach a load balancing state under different weight of network parameter α . As the weight of the network parameter α

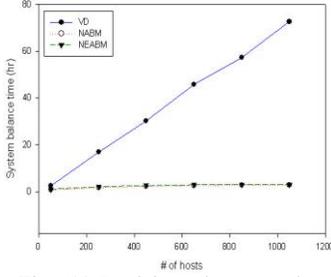


Fig. 7(a) Total time taken to reach a balanced state

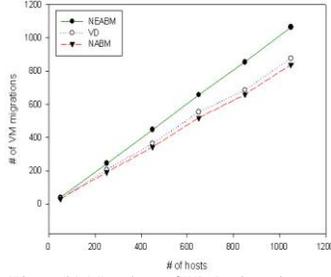


Fig. 7(b) Number of VM migrations to reach a balanced state

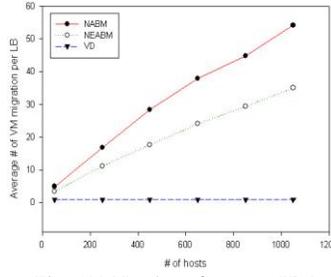


Fig. 7(c) Number of average VM migrations per round

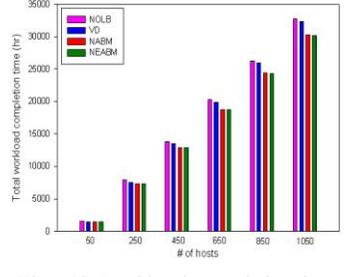


Fig. 7(d) Total batch completion time in different number of hosts

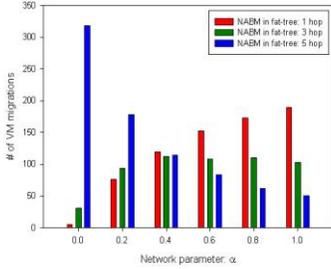


Fig. 8(a) Fat-tree: the number of VM migrations with different hop count

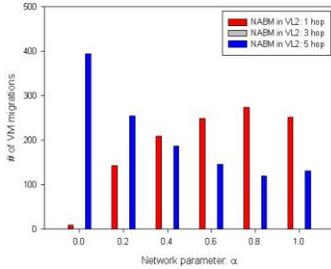


Fig. 8(b) VL2: the number of VM migrations with different hop count

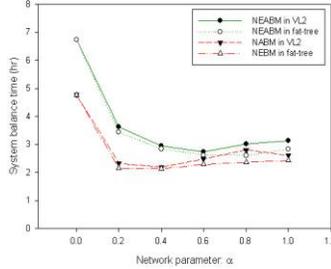


Fig. 8(c) Total time taken to reach a balanced state in fat-tree and VL2

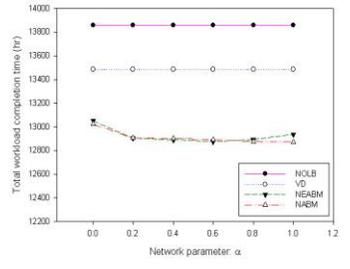


Fig. 8(d) Total batch completion time in fat-tree topology

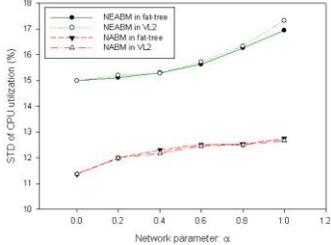


Fig. 8(e) STD of CPU utilization in different network parameters

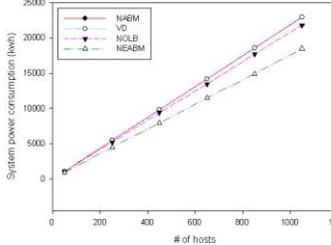


Fig. 9. Power consumption of system in different number of hosts

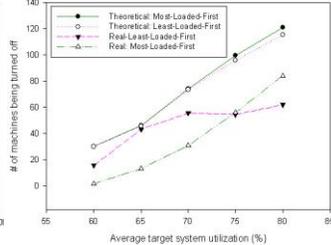


Fig. 10(a). Number of machines being turned off when ϵ is 0.15

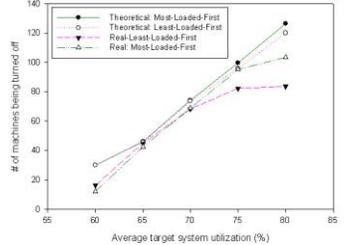


Fig. 10(b) Number of machines being turned off when ϵ is 0.2

increases, *NABM* can select more hosts with shorter paths to migrate VMs for both fat-tree and VL2 as shown in Fig. 8(a) and 8(b). Fig. 8(c) shows that the time till the system reaches the balanced state is reduced as the network parameter α increases. It's because migrating VMs along a shorter path could improve the mean migration time. Especially, compared with no network cost or very small network cost cases, considering the network hop distance ($\alpha > 0.2$) saves us more than half of the time to they take to reach a balancing state. With the decline of migration time, network-aware further improves application performances as shown in Fig. 8(d). However, large network parameter value could still hurt the time till the system reach balance and application performance. This is why we pick network weight α equal to 0.6 in our previous experiment setting. Fig. 8(e) shows the standard deviation of CPU utilization verse network weight while *NABM* and *NEABM* reach a balance state. When network weight α increases, the standard deviation of CPU utilization increases. It indicates the tradeoff between VM migration time and load balancing degree.

To compare the effect of power saving under a different system scale, we fix the power measurement interval to be 100 hours in all cases. During the measurement interval, we ensure that the batch running at each VM still consumes machine resources. Thus the dynamic power of power equation (1) will not become zero. Fig. 9 shows that the power consumption for

NEABM. *NEABM* saves, especially when the system scale gets large. This is because *NEABM* consumes the fewest static power as it turns off a large number of hosts.

Next, we study how average target utilization and the tolerance ratio ϵ of *NEABM* would affect the number of machines being turned off. In addition, since *NEABM* approaches the target utilization by first searching for potential hosts to turn off followed by applying parallel migration to empty VMs out of those hosts, we would like to know the ratio of selected potential machines to turn off over the exact number of turned off machines. In all the following experiments, the initial number of hosts is 450 hosts. In Fig. 10, we refer to the number of selected potential machines to turn off and the exact number of machines to turn off as *Theoretical* and *Real*, respectively. Fig. 10(a) and 10(b) show the number of machines to turn off with different tolerance ratio ϵ . The gap between *Theoretical* and *Real* number of turned off machines becomes closer as the tolerance ratio ϵ increases from 0.15 to 0.2. This is because increasing the room for accommodating the multi-resource variance also increases the chances for the VMs in potential turn off hosts (trigger hosts) to be migrated to non-trigger hosts without violating the system threshold. The selecting order from the least loaded machines first to turn off could actually turn off more machines than selecting from the most loaded machines first while the average target utilization is below 75% as shown in Fig. 10(a). This is because when the

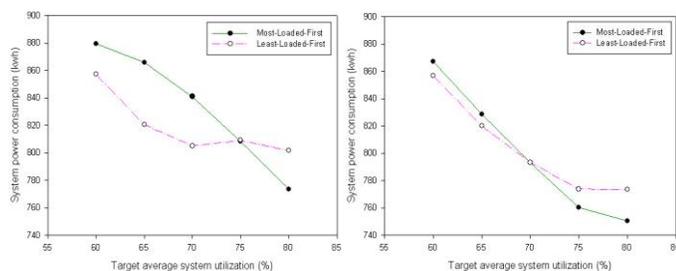


Fig. 11 (a). Total power consumption when tolerance ratio ε is 0.15

Fig. 11 (b). Total power consumption when tolerance ratio ε is 0.2

average target utilization is below 75%, the number of *Theoretical* machines to turn off is also small. Since there are fewer VMs to empty out from lightly utilized machines than from overloaded machines, the least load first doesn't take a large number of migrations to empty out machines compared to selecting order from the most load first. However, when the average target utilization is greater than or equal to 75%, the least load first perform worse than the most load first in term of number of machines to turn off. Although the least load first still takes few migrations to empty out each selected potential host, it also selects a large number of non-trigger nodes to be trigger nodes. Thus it could lead to a few number of trigger nodes not being turned off. Fig. 10(b) shows that with a higher tolerance ratio ε , the real number of turned off machines is close for both cases as the average target utilization is below or equal to 70%. As the average target utilization is greater than 70%, the exact number of turned off machines for selecting from the most load first is greater than selecting from the least load first. The reason is similar to the trend with Fig. 9(a) described earlier.

Fig. 11(a) and 11(b) show the respective power consumption for 10(a) and 10(b), respectively. The power consumption is measured over a power measurement interval of 9 hours. As the target utilization increases, each host in the system could afford more workloads and thus *NEABM* could turn off more machines, thereby the power consumption also decreases. As it was shown, deciding which sorting order to turn off machines also allows *NEABM* to achieve better power efficiency.

VII. CONCLUSION & FUTURE WORKS

We propose a network-aware multi-resource load-balancing scheme using a parallel VM migration. We transform the parallel VM migration to a minimum weighted matching problem of a weighted bipartite graph in a cloud system. Our algorithm migrates VMs parallel with each other to minimize the time to get the system to a balanced state and thus increases the throughput of overloaded hosts. Since the network hop distance between two hosts will affect the migration time, we further consider network hop distance into our cost function. In addition, since power wastes in the cloud data center are mainly due to the idle power at low utilized hosts, in order to save energy while preserving multi-resource load balancing, we extend our load balancing algorithm to turn off as many machines as possible. Simulation results show that our *NABM* algorithm improves the throughput on overloaded machines up to 10% compared with *VD*. Our energy-aware method saves up to 38% power consumption without degrading the application's completion time. In the future, we will take network bandwidth into consideration. In order to avoid multiple migrations

sharing a same network link, a greedy algorithm will be applied to select one VM migration pair at a time until no more migration pair. However, the VM migrations still can be performed in parallel.

REFERENCES

- [1] T. Chieu, A. Mohindra, A. Karve and A. Segal, "Dynamic Scaling of Web Applications in a Virtualized Cloud Computing Environment," in *ICEBE*, 2009.
- [2] Y. Zhao and W. Huang, "Adaptive distributed load balancing algorithm based on live migration of virtual machines in cloud," in Fifth International Joint Conference on INC, IMS and IDC, 2009, pp. 170-175.
- [3] M. Randles, et al., "A comparative study into distributed load balancing algorithms for cloud computing," in IEEE 24th International Conference on Advanced Information Networking and Applications Workshops, 2010, pp. 551-556.
- [4] A. Singh, et al., "Server-storage virtualization: integration and load balancing in data centers," in Proceedings of the 2008 ACM/IEEE conference on Supercomputing, 2008, p. 53.
- [5] V. Shrivastava, et al., "Application-aware virtual machine migration in data centers," in Proceedings IEEE INFOCOM, 2011, pp. 66-70.
- [6] X. Meng, et al., "Improving the scalability of data center networks with traffic-aware virtual machine placement," in Proceedings IEEE INFOCOM, 2010, pp. 1-9.
- [7] H. Xu and B. Li, "Egalitarian stable matching for VM migration in cloud computing," in IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), 2011, pp. 631-636.
- [8] S.-H. Lim, et al., "Migration, assignment, and scheduling of jobs in virtualized environment," *ACM USENIX workshop HotCloud*, 2011, vol. 40, p. 45, 2011.
- [9] K. Ye, et al., "Live migration of multiple virtual machines with resource reservation in cloud computing environments," in IEEE International Conference on Cloud Computing, 2011, pp. 267-274.
- [10] D. Arora, et al., "On the benefit of virtualization: Strategies for flexible server allocation," in Proc. USENIX Workshop on Hot Topics in Management of Internet, Cloud, and Enterprise Networks and Services (Hot-ICE), 2011.
- [11] M. Bienkowski, et al., "Competitive analysis for service migration in vnets," in Proceedings of the second ACM SIGCOMM workshop on Virtualized infrastructure systems and architectures, 2010, pp. 17-24.
- [12] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, pp. 83-97, 2006.
- [13] B. Li, et al., "EnaCloud: an energy-saving application live placement approach for cloud computing environments," in IEEE International Conference on Cloud Computing, 2009, pp. 17-24.
- [14] R. Buyya, et al., "Energy-efficient management of data center resources for cloud computing: A vision, architectural elements, and open challenges," *PDPTA*, 2010.
- [15] Xen. Available: <http://www.cl.cam.ac.uk/research/srg/netos/xen/>
- [16] SPECjvm. Available: <http://www.spec.org/>
- [17] K. P. Bubendorfer and J. H. Hine, "A compositional classification for load-balancing algorithms," technical report, No. CS-TR-99-9, 1998.
- [18] A. Greenberg, et al., "VL2: a scalable and flexible data center network," in ACM SIGCOMM Computer Communication Review, 2009, pp. 51-62.
- [19] NOX. Available: <http://www.noxrepo.org/>
- [20] M. Al-Fares, et al., "A scalable, commodity data center network architecture," in ACM SIGCOMM Computer Communication Review, 2008, pp. 63-74.
- [21] W. Leinberger, et al., "Multi-capacity bin packing algorithms with applications to job scheduling under multiple constraints," in International Conference on Parallel Processing, 1999, pp. 404-412.
- [22] S. K. Garg and R. Buyya, "NetworkCloudSim: Modelling Parallel Applications in Cloud Simulations," in Fourth IEEE International Conference on Utility and Cloud Computing, 2011, pp. 105-113.