



Virtual machine placement in cloud systems through migration process

Kangkang Li, Huanyang Zheng, Jie Wu & Xiaojiang Du

To cite this article: Kangkang Li, Huanyang Zheng, Jie Wu & Xiaojiang Du (2015) Virtual machine placement in cloud systems through migration process, International Journal of Parallel, Emergent and Distributed Systems, 30:5, 393-410, DOI: [10.1080/17445760.2014.931957](https://doi.org/10.1080/17445760.2014.931957)

To link to this article: <http://dx.doi.org/10.1080/17445760.2014.931957>



Published online: 31 Jul 2014.



Submit your article to this journal [↗](#)



Article views: 98



View related articles [↗](#)



View Crossmark data [↗](#)

Virtual machine placement in cloud systems through migration process

Kangkang Li*, Huanyang Zheng¹, Jie Wu² and Xiaojiang Du³

Department of Computer and Information Sciences, Temple University, Philadelphia, PA 19122, USA

(Received 17 March 2014; final version received 2 June 2014)

Cloud computing is an emerging technology that greatly shapes our lives, where users run their jobs on virtual machines (VMs) on physical machines (PMs) provided by a cloud service provider, saving the investment in upfront infrastructures. Due to the heterogeneity of various jobs, different VMs on the same PMs could have different job completion times. Meanwhile, the PMs are also heterogeneous. Therefore, different VM placements have different job completion times, and our objective is to minimise the total job completion time of the input VM requests through a reasonable VM placement schedule. This problem is NP-hard, since it can be reduced to a knapsack problem. We propose an off-line VM placement method through an emulated VM migration process, while the on-line VM placement is solved by a real VM migration process. The migration algorithm is a heuristic approach, in which we place the VM to its best PM directly, if this PM has enough capacity. Otherwise, we migrate another VM from this PM to accommodate the new VM, if a pre-specified migration constraint is satisfied. Furthermore, we study a hybrid scheme where a batch is employed to accept upcoming VMs for the on-line scenario. Evaluation results validate the high efficiency of the proposed algorithms.

Keywords: VM migration; job completion time; off-line; on-line; VM placement

1. Introduction

Nowadays, cloud computing is an emerging technology that greatly shapes our lives. With the large pools of computing and storage resources provided by cloud providers, many companies can rent these resources and run their jobs on virtual machines (VMs), saving the investment in upfront infrastructures. In particular, the technique of VM migration [1,2] enables us to move a VM from one physical host to another. This spatial flexibility is effective in server consolidation, power consumption saving and so forth.

Obviously, VMs are different among their resource demands and their running jobs. Due to this heterogeneity, different VMs running on the same physical machine (PM) could have different job completion times. Moreover, in most circumstances, the host PMs are also heterogeneous, such as different CPU architectures, different OS, differing amounts of memory, storage and so on. Hence, the same VM placed on different PMs would have different completion times, and our objective is to minimise the total job completion time of the input VM requests through a reasonable VM placement schedule. This problem is NP-hard, since it can be reduced to a knapsack problem.

VM migration is an efficient tool for resource provisioning, by dynamically rearranging the previous placement. For instance, when all PMs do not have enough

*Corresponding author. Email: kang.kang.li@temple.edu

capacities for the new VM, through migrating one running VM to another host, we could save space to accept the incoming VM. In this paper, we consider both off-line and on-line VM placement problems with the objective of minimising the total job completion time.

Under the off-line scenario, the information of input VMs is known a priori. Therefore, we propose a VM placement algorithm through an *emulated* VM migration process, in which we place the VMs one-by-one. Suppose that there is an incoming VM waiting for placement in a cloud system with three PMs: PM_1 , PM_2 , PM_3 . The job completion times of the incoming VM placed on PM_1 , PM_2 , PM_3 are, respectively, t_1 , t_2 , t_3 . As shown in Figure 1(a), for the incoming VM, we could directly place it in PM_1 , which has enough remaining capacity to support it. Then, the total completion time will increase by t_1 . However, this option might not be optimal. Suppose that placing the VM in PM_2 has the minimal completion time t_2 (among t_1 , t_2 and t_3), then PM_2 is the best choice for the incoming VM to minimise the total completion time. However, as shown in Figure 1(b), the capacity of PM_2 is not sufficient. Instead of placing the incoming VM in PM_1 , an alternative is to migrate one VM on PM_2 to PM_3 , which makes room for the incoming VM. In that case, the incoming VM could be accepted into PM_2 to minimise the total job completion time.

However, migration has a preliminary constraint, even if the remaining capacity of PM_3 could accept the migrated VM. On the one hand, by saving space to accept this incoming VM into PM_1 , we could have the minimal completion time t_2 for the incoming VM. On the other hand, since the completion time of the same VM on different PM hosts varies, for the migrated VM, there might be an increase in completion time due to the change of host. Since we aim to minimise the total completion time of the whole VMs set, if the time increase of the migrated VM is too large, we would rather select the first choice without migration. This is because migration would be meaningless in minimising the total completion time. Therefore, we conclude that the migration constraint is that, even with the increased completion time of the migrated VM, the migration-based VM placement (MBVMP) is still better than the direct placement in minimising the total completion time.

Therefore, in the emulated-MBVMP, we have two options, i.e. direct placement and migration-based placement. The insight behind the option selection is to compare the total job completion time of these two choices: we choose migration-based placement, if its gain of placing the incoming VM minus the switching loss of migrating the victim VM is larger than the gain of direct placement; otherwise, we choose direct placement. Again, note that the migration process is only emulated in the off-line VM placement scenario. We do not actually place the VMs until all VMs' final positions are determined.

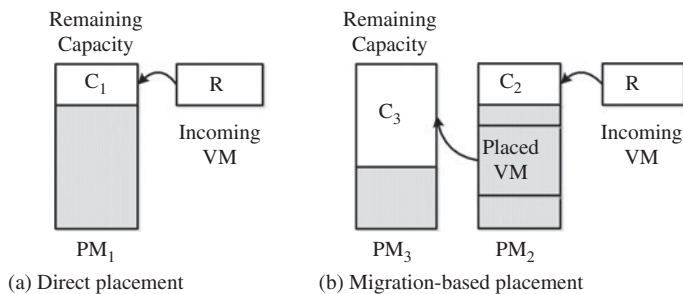


Figure 1. An illustration of the VM placement.

Regarding the on-line scenario where pre-information is not available, our VM placement algorithm is based on *real* VM migrations, which will inevitably introduce migration delay.[3] In fact, VM migration basically consists of transferring its memory image from the source host to the destination. Through experiments, Verma et al. [4] observed that there is a linear relationship between the active memory of a VM and the duration of migration. We refer to this delay as the *migration overhead*, since it will increase the total job completion time. Obviously, migrating a VM with a larger resource demand will lead to a longer delay, and thereby, a large migration overhead.

Furthermore, we conduct a study on a hybrid scheme, where we introduce a *batch* to help the VM placement in the on-line scenario. Instead of one-by-one placement, several on-line incoming VMs are reserved into a batch. Then, the information of the VMs in the batch would be known to us. In that case, we can simultaneously place them together through the off-line placement algorithm without the migration overhead. However, the batch will introduce the waiting time of the VM placement. Therefore, there exists a tradeoff between the on-line migration overhead and batch waiting time.

In this paper, we formulate the total completion time minimisation VM placement problem under both off-line and on-line scenarios. Due to the NP-hardness of this problem, we propose a heuristic MBVMP algorithm to give an efficient solution. Our main contributions are summarised in the following:

- To the best of our knowledge, our work is the first to adopt VM migration for the total completion time minimisation VM placement problem, considering both off-line and on-line scenarios.
- For the off-line VM placement, we also study the case of homogeneous resource demand, and compare our heuristic algorithm with the optimal solution of the maximal matching problem. The results show that the performance of our algorithm is very close to the optimal solution, verifying the high efficiency of our algorithm.
- We also conduct a study on a hybrid scheme, where a *batch* is introduced in the VM placement. In that case, we can integrate the off-line VM placement into the on-line scenario, which could avoid the migration overhead.

The remainder of the paper is organised as follows: in Section 2, we study the VM placement problem under the off-line scenario, and propose our off-line MBVMP algorithm. In Section 3, we perform research on the on-line scenario. Section 4 introduces a hybrid scheme of VM placement. In Section 5, we conduct experiments to evaluate the performance of our algorithms under both off-line and on-line scenarios. In Section 6, we introduce some previous work. Finally, we give the conclusions in Section 7.

2. Off-line VM placement

2.1 Off-line problem description

In this section, we study the VM placement problem under the off-line scenario, in which we know the information about the incoming VMs set a priori (e.g. each VM's resource demand, and the total number of VM requests). Suppose a cloud system has N PMs, where the capacity of the j th PM is C_j . There are M VMs waiting for the placement, and the resource demand of the i th VM is R_i . Due to the heterogeneity of PMs and VMs, different VMs on the same host have various job completion times. Therefore, the element of t_{ij} in completion time matrix $[t_{ij}]$ denotes the job completion time of the i th VM, which is placed into the j th PM. For simplicity, we assume t_{ij} is a constant, which is not related to

Table 1. Notations.

| Notation | Description |
|----------|---|
| VM_i | The i th VM |
| PM_j | The j th PM |
| R_i | The resource demand of VM_i |
| C_j | The remaining capacity of PM_j |
| x_{ij} | $\{0,1\}$ variable indicates whether VM_i is placed into PM_j |
| t_{ij} | The job completion time of VM_i in PM_j |

the current load in the j th PM. Table 1 presents the notations for the variables. The objective is to find an optimal placement that minimises total completion time of input jobs. Then, this optimisation problem is formulated as follows:

$$\text{Minimise } \sum_{i=1}^M \sum_{j=1}^N x_{ij} t_{ij}, \quad (1)$$

$$\text{s.t. } \sum_{j=1}^N x_{ij} \leq 1, \quad \sum_{i=1}^M x_{ij} R_i \leq C_j, \quad x_{ij} \in \{0, 1\}. \quad (2)$$

It can be seen that this problem is essentially a multi-knapsack problem, where VMs are the items selected for the PMs (equivalent to knapsacks). Resource demands of the VMs are the weights of the items, and the job completion time is equivalent to the values of the items. Note that here we are minimising the job completion time, rather than maximal values in the knapsack problem (but they are essentially the same). Therefore, this problem is NP-hard. Hence, we prepare to give a heuristic solution to solve it.

2.2 Off-line problem analysis

In this section, we analyse the off-line VM placement problem, which is based on the idea of VM migration. The VMs are placed one after another. As discussed before, for VM_i , there is an optimal PM_j with minimal t_{ij} . If the remaining capacity C_j is enough, we could finish placing VM_i and switch to the next VM. Otherwise, if C_j is not sufficient, we have two options below to place it.

- *Direct placement*: Among those PMs which have enough resources to accept the incoming VM, we select the one with the minimal completion time.
- *Migration-based placement (emulated)*: We try to migrate one VM placed on the optimal PM to another host for saving space to accept the incoming VM. In the off-line scenario, the migration process is emulated. Therefore, we call it (emulated) migration-based placement.

Obviously, direct placement is not the optimal choice for the incoming VM, since we should not miss the opportunity to place the incoming VM into the PM corresponding to the minimal completion time, especially under some circumstance that the completion time of the first-choice is much less than that of the others. Therefore, we want to take advantage of VM migration to do the placement.

The key for the feasibility of migration-based placement is to satisfy the migration constraint, i.e. the migration-based placement is better than the direct placement in terms of the total completion time. To achieve our migration-based placement, we need to find a *qualified* victim VM placed on the optimal PM to migrate, which is not always searchable. There are three qualifications that such a qualified victim VM must meet:

- (1) Could save enough resources for accepting a new VM.
- (2) Could find a new available host with enough current capacity to accept the victim VM.
- (3) With the increased completion time of the migrated victim VM, the completion time by migration-based placement is less than that of direct placement.

If such a qualified victim VM is found on the minimal completion time PM, then we would directly choose migration-based PM. If more than one qualified victim VM exists in the minimal completion time PM, we would choose the VM with the minimal completion time increase to migrate. Of course, there could exist some special circumstances for a particular incoming VM.

- We could not find a qualified victim VM on the minimal completion time PM. We then try to place the incoming VM in the next-minimal completion time PM. If the incoming VM still could not be placed, we continue to try the next-next-minimal completion time PM. The search process is terminated when the incoming VM is placed, or encounters the best available PM found by direct placement.
- After trying all the PMs, and if the incoming VM is still not placed, we have no choice but to reject it.

Algorithm 1. VM placement algorithm

```

1: for  $i = 1$  to  $M$  do
2:   Sort PMs increasingly by  $t_{ij}$ .
3:   for  $j = 1$  to  $N$  do
4:     if  $C_j$  is sufficient for  $VM_i$  then
5:       Place  $VM_i$  in  $PM_j$ 
6:       if migration-based placement ( $VM_i, PM_j$ ) then
7:         Do VM migration and place  $VM_i$  into  $PM_j$ 
8:       if  $VM_i$  is still not placed then
9:         Reject  $VM_i$ 

```

Algorithm 2. (emulated/real) Migration-based placement

```

1: for all the VMs ( $VM_k$ ) on  $PM_j$  do
2:   if  $VM_k$ 's migration makes enough room for  $VM_i$  then
3:     for all the PMs (except  $PM_j$ ) do
4:       Find the best available PM (except  $PM_j$ ) for  $VM_k$ 
5:       Select the best qualified victim VM running on  $PM_j$ 
6:   if migration constraint is satisfied then
7:     return true
8:   else
9:     return false

```

As we discussed before, under the off-line scenario, VM migration is not actually implemented. However, the emulated migration process shows great promise in dynamically improving the previous placement to lower the total completion time.

2.3 MBVMP algorithm

As we could observe from Algorithm 1, for each incoming VM_i , we first sort the completion time t_{ij} of each PM and accordingly rearrange the PMs set. After that, starting from the first PM (the best choice to place VM_i in minimising total completion time) in the rearranged PMs, we try to place the incoming VM_i . If PM_j 's remaining capacity C_j is sufficient, we could directly place VM_i and switch to place the next VM. Otherwise, by using the direct placement to test the migration constraint, we try to adopt the MBVMP. Considering the two loops in Algorithm 1, the total time complexity of our off-line MBVMP algorithm is $O(M^2N)$. Again, note that the migration-based placement is only emulated for the off-line scenario.

In Algorithm 2, we try to achieve the MBVMP. That is, we try to find the qualified victim VM among all VMs on PM_j to migrate. For each VM_k already placed in PM_j , if it could make enough room for the incoming VM_j , we then try to find it a best new host. That is, among all the other PMs (except PM_j itself) with enough remaining capacity to accept VM_k , we select the PM that could yield the minimal completion time increase as the best new host. Among these VMs that could make enough room for the incoming VM_j , we select the best victim VM that has the minimal increased completion time. After that, we verify whether the victim is qualified by testing the migration constraint. Through comparing the direct placement and migration-based placement in minimising the total completion time, we could know whether the selected best victim VM is qualified. If the migration constraint is satisfied, the victim VM is qualified. Then, we do the migration and place the incoming VM_j .

2.4 Case study

Now we consider a special case under the off-line scenario, in which the resource type is uniform. That is, each PM's resource capacity could be sliced into unit slots, and all of the VMs have the same resource demand of one slot. Each VM is placed into one slot in the PM, thus there exists a unique correspondence with each VM and each PM's unit slot. Given this, we could transfer the VM placement into a minimal matching problem in the graph theory. We would use Figure 2 to illustrate the mapping from VM placement to minimal matching.

In Figure 2, each VM and PM unit pair is regarded as the vertices in the bi-partite graph, and the weighted edge connecting each VM and PM unit pair refers to the completion time of placing VM in that unit of the PM. Thus, finding a minimal total completion time VM placement is equal to finding the minimal matching between VM and PM unit pairs in the weighted bi-partite graph. There is an optimal Kuhn–Munkras (KM) algorithm to solve this problem, and we use an example to compare our proposed MBVMP algorithm with the optimal solution.

We assume that each PM in the cloud system has the capacity of six VM slots, and we vary the number of PMs from 50 to 250. We set that the total resource demands of VMs are equal to the total capacity of PMs. In that case, given the N PMs, we have a set of $N \times 6$ VMs to be placed into the cloud system. Since all of the VMs have the same resource demand of one slot, each VM should have a corresponding slot in the PM to be placed into.




Figure 2. Minimal matching and VM placement.

We also randomly generated the set of each element in completion time matrix $[t_{ij}]$. From the comparison results shown in Figure 3, we could see that the performance of our proposed algorithm is very approximate to the optimal solution. Given M incoming VMs and N PM hosts, the time complexity of the KM algorithm is $O(M^2N)$, which is equal to our off-line MBVMP algorithm.

2.5 Multiple-hop migration

In the previous discussion, we only consider the one-hop migration, which has a low time complexity. However, multiple-hop migration might lead to a better performance. In Algorithm 3, we present the multi-hop algorithm. Suppose the average number of placed VMs on a PM is K , then the time complexity of Algorithm 3 is K times as much as that of

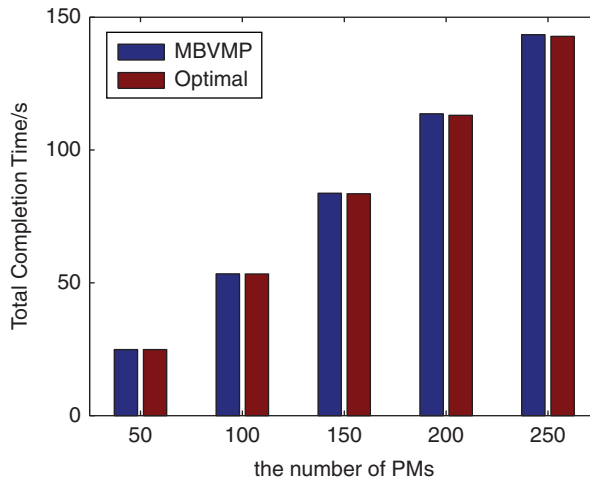


Figure 3. Comparison results.

Algorithm 3 Multiple-hop VM placement

```

1: for each PM in the set do
2:   Sort PMs set increasingly by completion time of the VM to be placed
3:   Select the best  $PM_j$  with the minimal completion time for the incoming VM
4: if  $C_j$  is sufficient for  $VM_i$  then
5:   Place the VM in  $PM_j$ 
6: for all the VMs running on the best PM do
7:   Select the best victim VM by comparing migration gain
8: Recursively call this algorithm to place the best victim VM among  $PM \in P - PM_j$ 

```

Algorithm 1. Obviously, Algorithm 3 introduces multiple-hop migrations with a better performance on minimising the total completion time, as the tradeoff of a higher time complexity.

3. On-line VM placement

Under the on-line scenario, the VM requests are coming one by one without knowing the information beforehand. Hence, we place VM one after another through an emulated process. Since our MBVMP algorithm places VMs one after another for the off-line scenario, it could still work in the same way for on-line VM placement.

The major difference between on-line and off-line VM placement is that the migration process is actually implemented. Hence, it is a real process which would introduce a migration delay, increasing the total completion time. Therefore, for the victim VM, not only do we have the completion time increase caused by migration, but we also have an extra migration overhead due to the on-line migration process.

Similar to the off-line scenario, we place the incoming VM one after another. When a new VM' comes, we first try to place it into the PM that corresponds to the minimal completion time value. If the corresponding minimal completion time PM does not have enough resources, we would have the same two options (direct placement and migration-based placement) as in the off-line scenario.

However, the migration constraint is different. Due to the migration overhead, the migration constraint would be that, even with both the migration overhead and the completion time increase, migration-based placement is still better than direct placement in minimising the total completion time. Thus, different with the off-line scenario, the third item of qualifications for a qualified victim VM under the on-line scenario should be changed into: even with both the completion time increase and migration overhead, the completion time by migration-based placement is still better than direct placement. We want to choose the qualified victim VM with the minimal sum of migration overhead and completion time increase. Besides, we could have similar special circumstances as in the off-line VM placement; constrained by the literature, here we do not discuss them further.

Besides, for on-line VM placement, we are not aware of the information of the incoming VMs, such as how many total VMs will arrive. Thus, when all the PMs are so fully loaded that even migration is not helpful to accept more VMs, we need to stop the placement process for this PMs set. This requires us to reserve a buffer to determine when we should directly reject the new VM request. In our evaluation, we set the size of a buffer to 10. That is to say that if 10 consecutive VMs are rejected, then we conclude the VM placement.

4. A hybrid scheme

Obviously, the major disadvantage of the on-line scenario is the migration delay, which will cause much overhead in the total completion time. However, the off-line scenario does not have such a drawback, because we know all the VMs' information beforehand, and they could all be placed at one time together. Thus, there is no migration overhead in the off-line scenario. This advantage of off-line scenario leads to the intuition that we could adopt a hybrid scheme that integrates the off-line scheme into the on-line scenario.

In fact, we could reserve a batch to store the on-line VM requests. In that case, the reserved VMs' information is known to us. Therefore, we could place the VMs in a batch according to the off-line scenario, in which case, we could avoid the migration overhead.

However, the batch will introduce a waiting time overhead, since we need to wait for the VMs in the batch to reach a certain number, and then place them together. Therefore, this waiting overhead will also increase the total completion time, although we avoid the migration overhead. Therefore, there is a tradeoff between the migration overhead and the waiting overhead. Due to this, there are two batch models: user-oriented and provider-oriented batch models.

4.1 Provider-oriented batch model

The provider-oriented batch model is that we place all the VMs in the batch until it is full, as shown in Figure 4. By this, the advantage of the batch could be fully utilised, and the migration overhead could be largely reduced, which is beneficial to the cloud providers. An extreme case is that, the batch size is very large, and we place the VMs together until the total resource demands of VMs are up to the total capacity of PMs. In that case, it is equivalent to the off-line scenario.

However, this waiting overhead will also largely increase the total completion time, since we have to wait for the batch to be full, and then, place the VMs together. Obviously, the larger the batch is, the larger the waiting overhead will cost. Under the on-line scenario, we do not know when a VM will arrive, and how long it will take to fulfil the batch. If the batch waiting overhead is too large, users will suffer the degradation of their applications' performance. Hence, we have another user-oriented batch model.

4.2 User-oriented batch model

The performance of users' applications needs to be guaranteed by the cloud provider, according to service-level agreement (SLA). Therefore, the batch overhead cannot be so large as to harm the user's applications' performance requirements. The user-oriented batch model is that, after waiting for a certain period of time, we place the VMs in the batch no matter whether the batch is full or not, as shown in Figure 5. In that case, the applications' performance could be guaranteed as the users' requirements.




Figure 4. Provider-oriented batch.




Figure 5. User-oriented batch.

4.3 Qualitative analysis on batch size

In the previous sections, we have discussed the implementation method of the batch, i.e. provider-oriented batch and user-oriented batch. In this section, we give a qualitative analysis on the batch size. According to [5], the time consumption of VM migrations usually varies from 30 to 150 s, due to the transmission time of the memory (usually several GBs). Therefore, the bandwidth has significant influence on the migration time. On the other hand, the acceptable waiting time of the user greatly depends on the application. For example, web services [6] generally respond in less than 1 s. In this case, migration-based placement is not practical, i.e. the batch is useless and we should place the VMs as soon as possible. Therefore, the batch model should be applied to user applications which takes a relatively long time (e.g. several hours or more).

For most cases, we can determine the batch size through a utility model. The user's utility monotonously decreases with respect to the batch size, while the provider's utility monotonously increases with respect to the batch size. The objective is to maximise the summation of the user's utility and the provider's utility. Moreover, the utility function should exhibit some certain properties, as shown in Figure 6. According to the experience of daily life, the utility function of the user should have three stages as follows. (1) In the first stage, the utility is high and decreases slowly, since the user is not sensitive to a small waiting time. For example, the waiting time of 2 and 10 ms does not make too much difference, since people are not aware of such a small time duration. (2) Generally speaking, the user has an expected maximal waiting time (no one wants to wait too long). Once this time is approaching, the utility decreases quickly in the second stage. (3) In the third stage, the user has been waiting for too long. The utility is very low, such that waiting for an addition time does not greatly turn down the utility. Meanwhile, the utility function of the provider is very simple, since it should have a *diminishing return* property. The utility gain should decrease with respect to the batch size.

The two utility functions should be designed according to real systems and applications, i.e. different systems and applications should have different function formats, although they should have the properties discussed above. If the user's utility dominates (usually in some real-time applications and user interactive applications), then a small batch size is desired. If the provider's utility dominates (usually in some large




Figure 6. The utility functions of the user (left) and the provider (right).

computation applications), then a large batch size is desired. This tradeoff always exists in the cloud system.

5. Performance evaluation

In this section, we evaluate our proposed MBVMP algorithms under both off-line and on-line scenarios. We made comparisons with first-fit and best-fit algorithms. Notably, for the off-line scenario, we adopt the first-fit-decreasing and best-fit-decreasing algorithms.[7] In order to compare the performance of one-hop and multiple-hop migration, we conduct extended simulation by adopting two-hop migration under off-line scenarios. As for the on-line scenario, multiple-hop migration would largely increase the migration cost. Therefore, we usually do not consider it as an alternative in the on-line VM placement.

5.1 Off-line MBVMP algorithm evaluation

We evaluate the performance of our off-line MBVMP algorithm under three groups of simulations on the total completion time. We use mathematical abstraction to describe the physical meaning of PM capacity and VM resource demand. The measuring unit of resources is the unit slot, which could be easily interpreted to real configuration. For instance, one unit slot for Amazon S3 could be interpreted as unit storage space for renting. The set of each element in completion time matrix $[t_{ij}]$ is randomly generated.

(1) *Simulation settings*

- *Group 1:* We deploy PMs' numbers $N = 100, 150$ and 200 . We set all PMs' resource capacity range $[150,150]$ and all the VMs' resource demand range $[10,100]$, which conforms to the real proportional relationship.
- *Group 2:* We deploy all VMs' resource demand ranges $[10,50]$, $[10,100]$ and $[10,150]$. We set PMs' number $N = 150$ and all PMs' current capacity range $[150,150]$.
- *Group 3:* We deploy all PMs' current capacity ranges $[50,150]$, $[100,150]$ and $[150,150]$. We set PMs' number $N = 100$ and all VMs' resource demand range $[10,100]$.

(2) *Simulations results:* Results for three groups of simulations are shown in Figures 7–9. From these figures, we can see that under all settings of the number of PMs, our MBVMP algorithm could achieve a significantly higher total completion time than the other two heuristic algorithms. Besides that, we could still have the following observations:

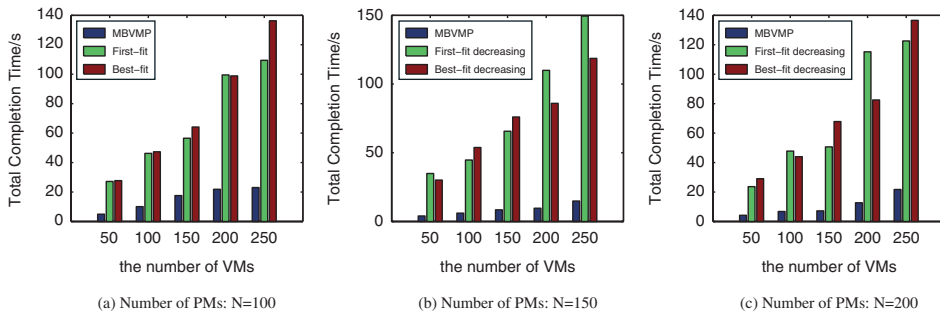


Figure 7. Performance comparisons of total completion time versus number of PMs.

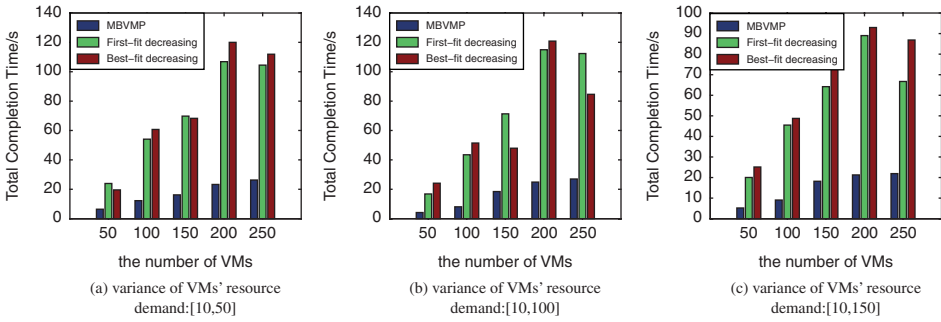


Figure 8. Performance comparisons of total completion time versus variance of VMs' resource demand.

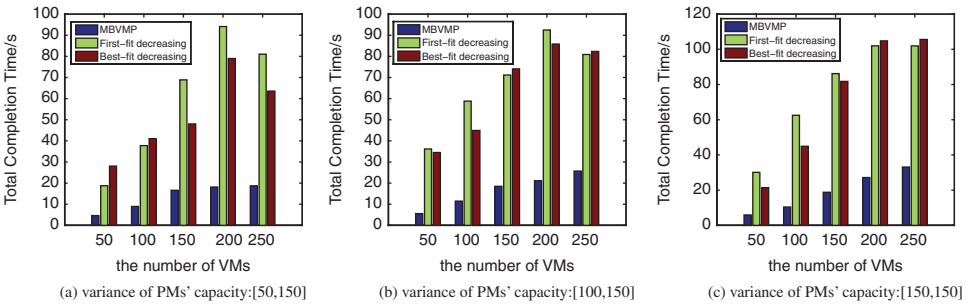


Figure 9. Performance comparisons of total completion time versus variance of PMs' capacity.

- In each of subfigures (a)–(c), with the increase of the number of VMs, the total completion time increases. This is quite straightforward, since more VMs are placed in the PMs, thus bringing more completion time.
- By comparison of (a)–(c) of Figure 7, we find that when the number of PMs increases, the gap of the total completion time between the 250 VMs and 200 VMs increases. This is due to the fact that, when number of PMs is few, given a 250 VMs input, many VMs actually cannot be accepted. Thus, there is not much difference between 200 VMs input and 250 VMs. However, due to the increase of PM hosts, and thereby more resources, more VMs could be accepted into the cloud system, and we could get an increase in total completion time. This also applies to comparisons between Figure 9(a)–(c), due to the increase of PM's capacity. On the other hand, Figure 8(a)–(c) shows the opposite trend, due to the increase of VMs' resource demand.

5.2 On-line MBVMP algorithm evaluation

For the on-line scenario, we could not know how many VM requests would come; therefore, we set it so that if 10 consecutive VM requests could not be placed, we then would stop the whole VM placement process. As discussed before, migration overhead has a positive correlation with the VM's resource demand. For analysis simplicity without loss of generality, we set a linear coefficient α to show that the migration overhead is α times the VM's resource demand, which is $c_i = \alpha \times R_i$. We also conducted three groups of

simulations with settings similar to those of the off-line scenario. In the first group, we change the number of PMs, while keeping the PMs' capacity and migration cost coefficient stable. In the second group, we vary the migration cost coefficient α to test its effect on the performance. At last, we vary the variance of PMs' capacities, while keeping the number of PMs and cost coefficient α stable. Simulation results are presented in Table 2. Table 2 shows that our MBVMP algorithm outperforms the best-fit and first-fit algorithms under the on-line scenario in all settings of variables.

5.3 Comparison of one-hop and two-hop migrations

We can see from the comparison in Figures 10–12, the performance is almost comparable between one-hop and two-hop migrations. This indicates that the two-hop migration will not significantly improve the performance. Therefore, our one-hop MBVMP algorithm will be good enough.

6. Related work

As a new paradigm of distributed computing, cloud computing has brought several key advantages into our lives, such as on-demand scaling and pay-as-you-go metred service. The development of virtualisation technologies has boosted the spread of cloud computing. Through management by VM monitor, [8–11] the physical resources of one PM could be sliced into multiple VMs. Such resource multiplexing largely reduces the total cost of ownership, and significantly improves the resource utilisation. As a contribution of virtualisation technology, VM migration [12–17] improves the resource rearrangement on the fly.

Much work has been done regarding VM placement in the cloud computing environment, which is a complicated task involving various constraints, including performance,[18] availability,[19] network [20,21] and cost.[22] Economic interests are one popular topic that shows up in research literature.[23,24] They tried to find an optimal VM placement that could either minimise the revenue for the cloud provider, or minimise the costs for the customers. In [24], for minimising the total economic completion time of the cloud provider, the author introduced an SLA-based dynamic resource allocation. The pricing mechanisms are related to the performance of quality of service (QoS) that the cloud provider could guarantee. The better performance the cloud provider could offer, the more revenue the cloud provider could obtain. Since different service requests have different pricing, higher priced services can get more resource provisioning from the cloud provider. The author also used a convex optimisation to present the optimal resource allocation. On the contrary, the author in [23] proposed an optimal VM placement with the objective of minimising the total renting of the users. In this paper, the author gives another pricing mechanism, referring to two payment plans: reservation plan and on-demand plan. Since the total resource demand is uncertain, and reservation plan has to be decided in advance, it may not meet the future demands of the user. For that reason, the author used the optimal solution of stochastic integer programming to give an optimal VM placement that could minimise total renting. However, the VM placement problem in [23,24] could achieve an optimal solution, which is not a common case. Many VM placement issues are NP-hard, thus we need to find a good heuristic algorithm to solve the problem, such as the first-fit and best-fit greedy algorithms used in [7].

Apart from the problems above aiming to get an optimal economic interest, network is another constraint needing consideration in the VM placement problem. In [20], a

Table 2. Performance comparison for on-line scenario.

| Total completion time (s) | Number of PMs | | | Coefficient of migration cost | | | Variance of PMs' capacity | | |
|---------------------------|---------------|----------|----------|-------------------------------|----------|----------|---------------------------|----------|----------|
| | 100 | 150 | 200 | 0.1 | 1 | 10 | [10,50] | [10,100] | [10,150] |
| MBVMP | 52.4475 | 55.2926 | 83.9606 | 56.1748 | 58.3576 | 53.1326 | 51.2983 | 53.9628 | 52.8962 |
| First-fit | 140.7237 | 214.7462 | 265.3581 | 150.1385 | 254.9847 | 284.9215 | 147.4251 | 221.2317 | 251.1223 |
| Best-fit | 130.2293 | 206.0019 | 233.3534 | 127.9724 | 215.9011 | 201.1946 | 125.9402 | 237.7812 | 263.9472 |

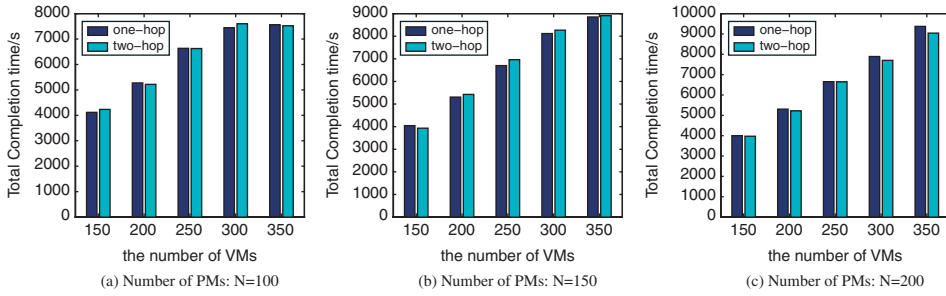


Figure 10. Performance comparisons of total completion time versus number of PMs.

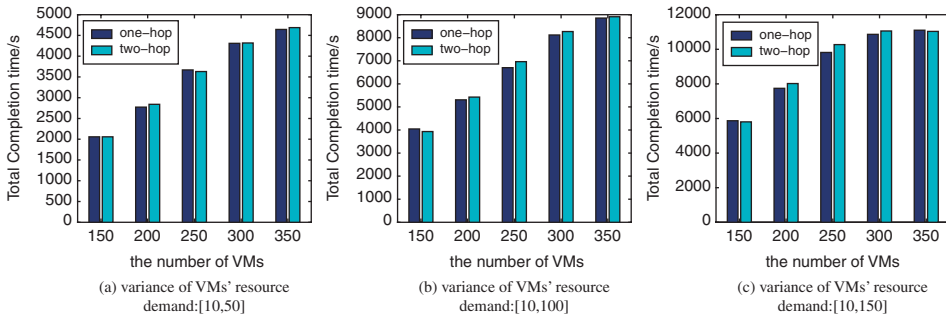


Figure 11. Performance comparisons of total completion time versus variance of VMs' resource demand.

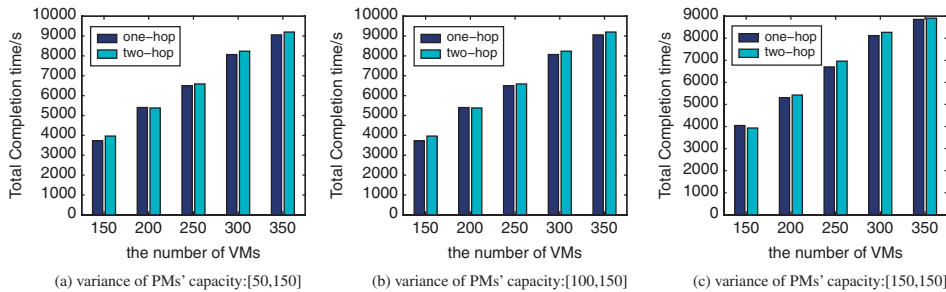


Figure 12. Performance comparisons of total completion time versus variance of PMs' capacity.

network-aware VM placement is proposed. In [20], when performing VM placement, not only are the physical resources (such as CPU and memory) considered, but the traffic demands between different VMs are also taken into account. The author gave a heuristic algorithm to allocate placement to satisfy both the communication demands and physical resource restrictions. In [25], Oktopus uses the hose model to abstract the tenant's bandwidth request, including both virtual cluster and oversubscribed virtual clusters. The virtual cluster provides tenants with guarantees on the network bandwidth they demand, which, according to [26], could be interpreted as minimal guarantee requirements. In [27], the author proposes to minimise the traffic cost through VM placement. Their objective is to place VMs that have large communication requirements close to each other, so as to

reduce network capacity needs in the data centre. A quadratic-assignment formulation of the traffic-aware placement problem is presented and solved with an approximation algorithm. However, their algorithm did not take into account the VM migration traffic, leading to a near complete shuffling of almost all VMs in each round. To alleviate this, VirtualKnotter [28] minimises the continuous congestion mainly in core and aggregation links with controllable migration traffic, which enables online VM replacement. In [29], a distributed cloud system is studied. The authors propose a network-aware VM placement algorithm, which considers the difference of latencies between inter-data-centre and inner-data-centre. They developed data-centre selection algorithms for VM placement that minimise the maximum distance between the selected data centres.

Moreover, our work is also potentially related to the process assignment problem in both symmetric and asymmetric multiprocessing.[30] In that technology, processes can also be migrated from one processor core to another. In the case of the non-uniform memory access (NUMA) model, the pages on the NUMA nodes are migrated. The idea of process migration is further applied into our VM placement strategy, where we emulate a VM migration process for the objective of off-line VM placement.[31]

7. Conclusion

In this paper, we studied the VM placement problem, focusing on minimising the total completion time of cloud providers under both off-line and on-line scenarios. Due to the NP-hardness of this problem, we propose our heuristic MBVMP algorithm. In particular, for the off-line scenario and homogeneous resource type, we compare our algorithm with the optimal maximal matching solution, which shows a high approximation of our algorithm with the optimal one. Furthermore, we study the hybrid scheme of integrating off-line placement into the on-line scenario, and propose two batch models considering users or providers separately. The simulation shows the high efficiency of our proposed heuristic algorithms, compared to the best-fit and first-fit heuristic algorithms.

Acknowledgement

This research was supported in part by NSF grants ECCS 1231461, ECCS 1128209, CNS 1138963, CNS 1065444 and CCF 1028167.

Notes

1. Email: huanyang.zheng@temple.edu
2. Email: jjiewu@temple.edu
3. Email: dux@temple.edu

References

- [1] Wang W, Zhang Y, Lin B, Wu X, Miao K. Secured and reliable VM migration in personal cloud. In: Computer Engineering and Technology (ICCET), 2010 2nd International Conference on, vol. 1; 2010 April. 2010. p. 705–709.
- [2] Xu H, Li B. Egalitarian stable matching for VM migration in cloud computing. In: Computer Communications Workshops (INFOCOM WKSHP), 2011 IEEE Conference on; 2011 April. p. 631–636.
- [3] Voorsluys W, Broberg J, Venugopal S, Buyya R. Cost of virtual machine live migration in clouds: a performance evaluation. In: Proceedings of the 1st International Conference on Cloud Computing, CloudCom '09; Berlin, Heidelberg: Springer-Verlag; 2009. p. 254–265.

- [4] Verma A, Kumar G, Koller R. The cost of reconfiguration in a cloud. In: Proceedings of the 11th International Middleware Conference Industrial Track, Middleware Industrial Track '10; New York (NY): ACM; 2010. p. 11–16.
- [5] Wood T, Ramakrishnan K, Shenoy P, Van der Merwe J. CloudNet: dynamic pooling of cloud resources by live WAN migration of virtual machines. In: Proceedings of the 7th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, VEE '11; New York (NY): ACM; 2011. p. 121–132.
- [6] Canfora G, Di Penta M, Esposito R, Villani ML. An approach for QoS-aware service composition based on genetic algorithms. In: Proceedings of the 7th Annual Conference on Genetic and Evolutionary Computation, GECCO '05; New York (NY): ACM; 2005. p. 1069–1075.
- [7] Xu J, Fortes J. Multi-objective virtual machine placement in virtualized data center environments. In: Green Computing and Communications (GreenCom), 2010 IEEE/ACM Int'l Conference on Int'l Conference on Cyber, Physical and Social Computing (CPSCom); 2010 December. p. 179–188.
- [8] Lange JR, Dinda P. Symcall: symbiotic virtualization through VMM-to-guest upcalls. In: Proceedings of the 7th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, VEE '11; New York (NY): ACM; 2011. p. 193–204.
- [9] Jones ST, Arpaci-Dusseau AC, Arpaci-Dusseau RH. VMM-based hidden process detection and identification using lycosid. In: Proceedings of the Fourth ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, VEE '08; New York (NY): ACM; 2008. p. 91–100.
- [10] Xu M, Jiang X, Sandhu R, Zhang X. Towards a VMM-based usage control framework for OS kernel integrity protection. In: Proceedings of the 12th ACM Symposium on Access Control Models and Technologies, SACMAT '07; New York (NY): ACM; 2007. p. 71–80.
- [11] Ranadive A, Gavrilovska A, Schwan K. IBMon: monitoring VMM-bypass capable infiniband devices using memory introspection. In: Proceedings of the 3rd ACM Workshop on System-level Virtualization for High Performance Computing, HPCVirt '09; New York (NY): ACM; 2009. p. 25–32.
- [12] Danev B, Masti RJ, Karame GO, Capkun S. Enabling secure VM-vTPM migration in private clouds. In: Proceedings of the 27th Annual Computer Security Applications Conference, ACSAC '11; New York (NY): ACM; 2011. p. 187–196.
- [13] Kumar S, Schwan K. Netchannel: a VMM-level mechanism for continuous, transparent device access during VM migration. In: Proceedings of the Fourth ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, VEE '08; New York (NY): ACM; 2008. p. 31–40.
- [14] Jain N, Menache I, Naor JS, Shepherd FB. Topology-aware VM migration in bandwidth oversubscribed datacenter networks. In: Proceedings of the 39th International Colloquium Conference on Automata, Languages, and Programming - Volume Part II, ICALP'12; Berlin, Heidelberg: Springer-Verlag; 2012. p. 586–597.
- [15] Surie A, Lagar-Cavilla HA, de Lara E, Satyanarayanan M. Low-bandwidth VM migration via opportunistic replay. In: Proceedings of the 9th Workshop on Mobile Computing Systems and Applications, HotMobile '08; New York (NY): ACM; 2008. p. 74–79.
- [16] Choi HW, Kwak H, Sohn A, Chung K. Autonomous learning for efficient resource utilization of dynamic VM migration. In: Proceedings of the 22Nd Annual International Conference on Supercomputing, ICS '08; New York (NY): ACM; 2008. p. 185–194.
- [17] Srinivasan K, Yuw S, Adelmeyer TJ. Dynamic VM migration: assessing its risks & rewards using a benchmark. SIGSOFT Softw Eng Notes. 2011;36(5):317–322.
- [18] Kusic D, Kephart J, Hanson J, Kandasamy N, Jiang G. Power and performance management of virtualized computing environments via lookahead control. In: Autonomic Computing, 2008. ICAC '08. International Conference on; 2008 June. p. 3–12.
- [19] Bin E, Biran O, Boni O, Hadad E, Kolodner E, Moatti Y, Lorenz D. Guaranteeing high availability goals for virtual machine placement. In: Distributed Computing Systems (ICDCS), 2011 31st International Conference on; 2011 June. p. 700–709.
- [20] Piao JT, Yan J. A network-aware virtual machine placement and migration approach in cloud computing. In: Grid and Cooperative Computing (GCC), 2010 9th International Conference on; 2010 November. p. 87–92.

- [21] Li K, Wu J, Blaisse A. Elasticity-aware virtual machine placement for cloud datacenters. In: Cloud Networking (CloudNet), 2013 IEEE 2nd International Conference on; 2013 November. p. 99–107.
- [22] Sharma U, Shenoy P, Sahu S, Shaikh A. A cost-aware elasticity provisioning system for the cloud. In: Distributed Computing Systems (ICDCS), 2011 31st International Conference on; 2011 June. p. 559–570.
- [23] Chaisiri S, Lee B-S, Niyato D. Optimal virtual machine placement across multiple cloud providers. In: Services Computing Conference, 2009. APSCC 2009. IEEE Asia-Pacific; 2009 December. p. 103–110.
- [24] Feng G, Garg S, Buyya R, Li W. Revenue maximization using adaptive resource provisioning in cloud computing environments. In: Proceedings of the 2012 ACM/IEEE 13th International Conference on Grid Computing, GRID '12; Washington (DC): IEEE Computer Society; 2012. p. 192–200.
- [25] Ballani H, Costa P, Karagiannis T, Rowstron A. Towards predictable datacenter networks. In: Proceedings of the ACM SIGCOMM 2011 Conference, SIGCOMM '11; New York (NY): ACM; 2011. p. 242–253.
- [26] Popa L, Krishnamurthy A, Ratnasamy S, Stoica I. Faircloud: sharing the network in cloud computing. In: Proceedings of the 10th ACM Workshop on Hot Topics in Networks, HotNets-X; New York (NY): ACM. p. 22:1–22:6.
- [27] Meng X, Pappas V, Zhang L. Improving the scalability of data center networks with traffic-aware virtual machine placement. In: INFOCOM, 2010 Proceedings IEEE; 2010 March. p. 1–9.
- [28] Wen X, Chen K, Chen Y, Liu Y, Xia Y, Hu C. VirtualKnotter: online virtual machine shuffling for congestion resolving in virtualized datacenter. In: Proceedings of the 2012 IEEE 32Nd International Conference on Distributed Computing Systems, ICDCS '12; Washington (DC): IEEE Computer Society; 2012. p. 12–21.
- [29] Alicherry M, Lakshman TV. Network aware resource allocation in distributed clouds. In: INFOCOM, 2012 Proceedings IEEE; 2012 March. p. 963–971.
- [30] Jiang Y, Shen X, Chen J, Tripathi R. Analysis and approximation of optimal co-scheduling on chip multiprocessors. In: Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques, PACT '08; New York (NY): ACM; 2008. p. 220–229.
- [31] Li K, Zheng H, Wu J. Migration-based virtual machine placement in cloud systems. In: Cloud Networking (CloudNet), 2013 IEEE 2nd International Conference on; 2013 November. pp. 83–90.