

Migration Decision for Hybrid Mobility in Reconfigurable Virtual Machines

Song Fu and Cheng-Zhong Xu
Department of Electrical & Computer Engg.
Wayne State University, Detroit, Michigan 48202
{oaksong, czxu}@wayne.edu

Abstract

Virtual machine (VM) is an important mechanism to multiplex computer resource. The increasing popularity of Grid computing and Peer-to-Peer systems has renewed research interests in dynamic and adaptive virtual machines. Service migration is a vital technique to construct reconfigurable VMs. By incorporating mobile agent technology, VM systems can improve their resource utilization, load balancing and fault tolerance significantly via hybrid mobility of services and computational agents. This paper focuses on the decision problem of hybrid mobility for load-balancing in reconfigurable VMs. We tackle this problem from three aspects: migration candidate determination, migration timing and destination server selection. The service migration timing and destination server selection are formulated as two optimization models. We derive the optimal migration policy for heterogenous distributed systems based on stochastic optimization theories. Renewal processes are applied to model the dynamics of migration. We model the agent migration problem by dynamic programming and extend the optimal service migration timing by considering dynamic composition of agents. Our decision policy is complementary to the existing service and agent migration techniques. Its accuracy is verified by simulations.

Keywords: Reconfigurable Virtual Machine, Load Balancing, Service Migration, Agent Migration, Migration Decision, Stochastic Optimization, Dynamic Programming.

1 Introduction

A virtual machine (VM) presents an abstract view of the underlying physical machine or system software to programs that run with it. It allows multiple applications and even different operating systems to run concurrently and multiplex computer resource. It also decouples the execution of applications from each other. The VM abstraction and isolation greatly simplify the development of applications and improve program portability. Since the classical VMs [13] were first designed as a solution to multiplexing shared mainframe resource in the early seventies, it has been an active research area; see [22] for an early review of this topic. With the explosive growth of network services, the Internet is evolved to enable the sharing and coordinated use of geographically distributed resource. To support network centric applications, research interests in VM are recently renewed. Systems like Denali [26], DVM [21], Terra [10] and Collective [20] develop virtual machine techniques for multiple network applications hosting on a single computer, or for inter-VM communication by standard network protocols, or for state transfer. Besides, the widely used Java VM (JVM) middleware and MPI software package facilitate the construction of distributed and parallel

applications. There are also recent efforts in the development of the agent-oriented resource brokerage infrastructure on JVM. Traveler [29] is an example of application layer virtual machines atop available servers to execute agent tasks.

Virtual machines provide a powerful layer of abstraction in distributed computing environments. It is certainly possible to deploy them as static computing units. However, this abstraction layer would not be fully exploited unless VMs are instantiated and managed dynamically. Virtual machine migration is one approach to their dynamic reconfiguration. Since VMs are usually created to provide services for programs atop them, we use the terms of VM migration and service migration interchangeably. Many recently proposed virtual machine systems incorporate service migration techniques for adaptation. Examples include the VM-based architecture for grid computing in [6], *capsule* in Collective [20], M-DSA [8] in Traveler, CFS [5] and [18] for Peer-to-Peer systems.

In addition to service migration, process/thread migration is another way to construct reconfigurable distributed systems [14]. Process/thread migration concerns the transfer of user-level processes or threads across different address spaces. JESSICA2 [31] and MigThread [11] are two examples that use thread migration to utilize idle cycles on other computers and to balance workload. In comparison with process/thread migration, agent migration is a more general approach to load balancing, fault tolerance, and data locality. An agent represents an autonomous computation that can migrate in a network to perform tasks on behalf of its creator. It is the autonomy of agent that makes high performance Grid computing different from cluster-wide process/thread migrations. Multi-hop mobility allows agent to move computations across a wide area network and negotiate with others. Agent technology has been applied to Grid computing for load scheduling in [4]. Load-balancing can be achieved by agent migration in Traveler and FLASH [17], which are able to react efficiently in response to the change of background load.

In a reconfigurable virtual machine system with hybrid mobility of services and agents, migration decision becomes a crucial performance issue. It deals with three aspects of the problem. First is migration candidate determination, concerning which server should transfer a service and/or which agent should exercise migration. Another one is migration timing, determining when a migration should be performed. Finally, a destination server should be selected so that the performance gain from a migration won't be outweighed by the migration overhead. Literature is littered with heuristic migration algorithms for either service or agent migrations and lacks formal analyses of the hybrid decision problem. This paper provides a rigid stochastic optimization model for this mobility decision problem. Our decision policy is complementary to the existing migration mechanisms.

The hybrid migration decision problem is non-trivial. A VM provides a dynamic environment for the execution of agents. Each agent will carry out its tasks during its life span (corresponding to a *time domain*) and its execution may be performed on different virtual servers due to migrations (referring to this as a *space domain*). Therefore, we have to consider both domains when deriving an optimal migration decision. Their interaction makes the decision process more complicated. Secondly, the capacities of physical servers in a distributed VM may be different, which is typical in Internet-based computing. This capacity heterogeneity results in different distributions of agent workload changes. Their distributions can even alter from one server to another during service migration. Besides, agent migration is hard to schedule in such a dynamic environment with server reconfiguration. The impact of agent migration on service migration also has to be carefully handled.

In this paper, we consider the mobility support for load balancing in heterogeneous distributed systems. We focus on the hybrid migration decision problem, by assuming the workload change of an agent is a random process with arbitrary probabilistic distributions. We relate the service/agent migration decision

in reconfigurable VM to a dynamic remapping scheduling problem in parallel computing. The migration issues are formulated as two stochastic optimization models, and optimal migration policies are derived in both time and space domains. Since general stochastic optimization approaches tend to reveal asymptotic or stationary properties of a random process, renewal processes are applied to model the dynamics of migration due to the fact that a migration operation may be invoked anytime over the course of computation. Our approaches are distributed in nature and scalable to large systems. The agent migration decision problem is modelled by dynamic programming and we integrate agent migration into the optimal service migration model. Simulations were conducted and the accuracy of our migration decisions was verified.

The rest of the paper is organized as follows. In Section 2, we present the related work. Section 3 introduces the system and workload evolution model of a target reconfigurable virtual machine. Section 4 and 5 describe the optimal decisions of service/agent migration, and their interplay. Simulation results for a verification of our decision policy are presented in Section 6. Section 7 concludes the paper with remarks on future work.

2 Related Work

Research interests in reconfigurable VM were renewed in recent years. Figueiredo *et al.* [6] proposed an architecture for grid computing based on dynamic virtual machines. It allows a VM to be created on any resource that is sufficient to support it. It enables service migrations among computers to pursue resource locality. Sapuntzakis *et al.* [20] tackled the problem of how to migrate operating system level VMs across the Internet. A virtual machine monitor is utilized to encapsulate the state of a running computer into a data type, *capsule*. By serializing and transferring capsules between physical machines, a VM can resume running after migration. To support virtual machine reconfiguration in Traveler [29], we designed a mobile distributed shared array (M-DSA) [8], which supports the migration of a DSA [3] service to a new physical server when the original one is overloaded or going to fail. Preliminary experimental results showed that service migrations at the application level with underlying middleware support could significantly reduce the migration overhead.

The virtual machine mechanism and service migration technique have also been applied to Peer-to-Peer systems. Chord [23] was one of the first that used the notion of virtual servers as a means of improving load balance. By allocating $\log N$ virtual servers on each real homogenous node, Chord ensures with high probability the number of objects per node is within a constant factor from optimal. CFS [5] accounts for node heterogeneity by ensuring the number of allocated virtual servers at each node is proportional to the node capacity. System adaptation is realized by deleting and creating virtual servers according to the load of each physical server. In [18], Rao *et al.* presented three heuristic algorithms to achieve virtual server migration in structured Peer-to-Peer systems. Their simulation results indicated that virtual server migration was able to balance the system load to some extent even with a simple migration scheme.

By utilizing mobile agents, a load scheduling and resource management scheme for Grid computing was proposed in [4]. In that scheme, agents cooperatively advertise and discover system services and distribute tasks accordingly, while the local grid load balancing is performed by individual task schedulers. Traveler presented a different approach, where mobile agents can autonomously roam the network to find virtual servers with appropriate capacities to run on. FLASH [17] established a two level mobile agent hierarchy to maintain system load information and utilized mobile agent migration to search for free resource through a cluster.

Although all of the works dealt with reconfigurable VMs, they focused on how to realize service and agent migration by providing different mechanisms. In this paper, we concentrate on migration decision in support of hybrid mobility. Although there are many reports in literature on mobile agent systems (see [25] for a review), there are few on agent scheduling for mobility. Moizumi and Cybenko [15] formulated a travelling agent problem to find optimal agent migration sequence which minimizes the expected completion time of a task with unreliable networks and unpredictable nodal services. In [32], Zhuang *et. al* proposed an approach to optimize the number of agent migrations for a given task by utilizing compiler optimization techniques. They both focused on finding the optimal agent itinerary to complete its tasks. However, we are interested in the decision problem as when and to which server an agent migration will be exercised. The objective is to balance system workload in a dynamic distributed environment supporting server reconfiguration.

In this paper, we relate the service/agent migration timing in reconfigurable VM to a dynamic remapping scheduling problem in parallel computing. Dynamic remapping is to redistribute the workload of processors at runtime efficient execution of dynamic applications [27]. Most recently, we studied the periodic remapping policy in [7], analyzing the impact of remapping frequencies. We formulated the problem as sequential stochastic optimization and derived optimal remapping frequencies with a priori known statistical behavior of the workload. We further modelled the optimal remapping problem as a vector-valued Markov chain and formulated it as a binary decision process in [30]. The optimal strategy was developed by employing optimal stopping rules in stochastic control. This approach can only be applied to small systems, because computation of the transition matrix increases dramatically as the size of a system becomes larger. Both approaches only consider the time domain in deriving the optimal solutions. The optimization approaches are based on global knowledge of system workload distributions for optimization, which is almost impossible in wide-area distributed computing.

3 Reconfigurable Virtual Machine Model

An agent-aware virtual machine is composed of a set of virtual servers spreading among distributed physical servers. Each virtual server can accommodate multiple agents to share resource and perform computation and communication. Each physical server can be run in two modes: *dedicated* or *multiprogrammed*. On a dedicated server, only virtual servers of the same application can reside, while a multiprogrammed server allows virtual servers of different applications to share its resource. On the other hand, a VM can be run on a single physical server or a cluster of servers. Correspondingly, on *dedicated* and *multiprogrammed single-server* systems, only service migration can be applied for their reconfigurability. In a *dedicated multi-server* VM, both service and agent migrations are useful to balance workload among multiple physical servers. The most general type is *multiprogrammed multi-server* VMs. But the asynchrony among different applications makes multiprogrammed servers hard to analyze. In this paper, we confine our discussion to the dedicated multi-server reconfigurable virtual machines.

We consider the bulk synchronous computation as our execution model because of its popularity in scientific and engineering applications. Its computations proceed in phases that are separated by global synchronization. During each phase, they perform calculations independently and then communicate new results with their data-dependent peers. Due to the need of synchronization between steps, the duration or execution time of a step is determined by the most heavily loaded server, on which the agent having the longest finish time resides. Phase-wise computation may exhibit varying requirements as the computation proceeds and even a static bulk synchronous computation may have non-deterministic computational characteristics. Fig. 1 illustrates an example of bulk synchronous computations over a VM with multiple agents on each server.

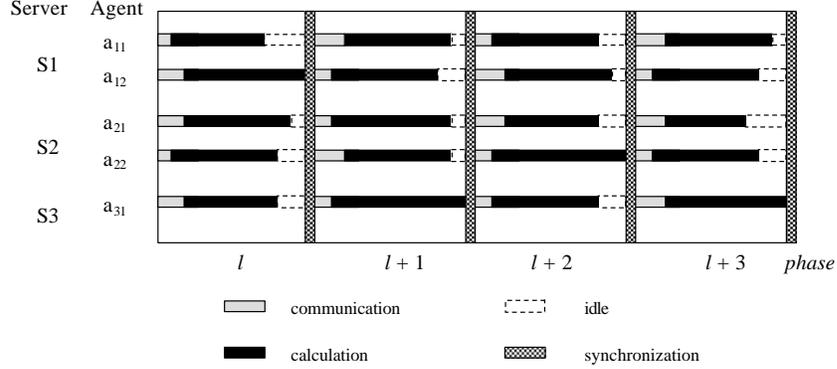


Figure 1: Bulk synchronous computations of virtual machine.

Consider a distributed application in a heterogeneous environment. A virtual machine $V = \langle s_1, s_2, \dots, s_N \rangle$, formed by N servers out of available physical servers, is allocated to the application. Let a sequence $C = \langle c_1, c_2, \dots, c_N \rangle$ denote their computational capacities and each $c_i, i = 1..N$, is a constant. The computation on the virtual machine comprises multiple agents, each of which executes a number of tasks that collectively determine its workload. We use a dynamic agent distribution sequence $M = \langle m_1, m_2, \dots, m_N \rangle$ to represent the number of agents assigned to different servers. Due to a many-to-one mapping between agents and servers, we use a sequence $A_i = \langle a_{i,1}, a_{i,2}, \dots, a_{i,m_i} \rangle$ to denote the set of agents residing on server s_i in the VM.

In the bulk synchronous computation, agents proceed in phases. Let $l = 0, 1, \dots$, be the phase index of the computation and $w_{i,k}(l)$ denote the workload of agent $a_{i,k}$ on server s_i at phase l . Due to the heterogeneity of server capacities, we use scaled workload, $\tilde{w}_{i,k}(l)$, to denote the workload of agent $a_{i,k}$ normalized by its quota of server capacity. We assume the proportional share scheduling [24] among mobile agents of the same server. Multiple agents share the resource and each of them has an equal quota of the server capacity. That is, for server s_i with capacity c_i , an residing agent $a_{i,k}$ is assigned a portion of capacity as $c_{i,k} = c_i/m_i$. Therefore, we have $\tilde{w}_{i,k}(l) = w_{i,k}(l)/c_{i,k} = m_i w_{i,k}(l)/c_i$. For simplicity in notation, we will use $w_{i,k}(l)$ to denote the scaled workload henceforth.

As the computation proceeds, the scaled workload of each agent may evolve dynamically. Let $\delta_{i,k}(l-1)$ denote the net change of scaled workload $w_{i,k}(l)$ from phase $l-1$ to l due to the workload generation and/or consumption by agent $a_{i,k}$ during the period. The overall scaled workload of server s_i equals to the summation of its agents' scaled workload, *i.e.* $w_i(l) = \sum_{k=1}^{m_i} w_{i,k}(l)$. Similarly, its scaled workload change is defined as $\delta_i(l-1) = \sum_{k=1}^{m_i} \delta_{i,k}(l-1)$. Let sequences $w(l) = \langle w_1(l), w_2(l), \dots, w_N(l) \rangle$ and $\delta(l-1) = \langle \delta_1(l-1), \delta_2(l-1), \dots, \delta_N(l-1) \rangle$ denote the scaled workload distribution at phase l and the scaled workload change distribution from phase $l-1$ to l of the virtual machine, respectively. Then the bulk synchronous computation can be characterized by an additive dynamic system:

$$w(l) = w(l-1) + \delta(l-1) \quad (3.1)$$

where the scaled workload change is independent of the server workload. The scaled workload is described by a Markov chain, which is assumed in many other studies on bulk synchronous computations [16, 7, 30]. As a result of scaled workload change between phases, the computation of each server becomes non-deterministic. Although this paper focuses on the additive scaled workload evolution model, our analysis could be extended to other dynamic systems as discussed in [28].

During a service migration, the residing virtual servers on a physical server can be transferred as a whole. We

can treat the virtual servers on each physical server as a set and the two terms, physical and virtual servers will be used interchangeably. When a server decides to perform a service migration to another available server during the computation, the correspondent entry in the capacity sequence C will be replaced by the destination server’s capacity. Similarly, agent migration leads to a modification of the agent distribution sequence M , as agents land on or leave from servers of the virtual machine. So the migration decision problem can be tackled based on renewal processes. We use sequences $C^0 = \langle c_1^0, c_2^0, \dots, c_N^0 \rangle$ and $M^0 = \langle m_1^0, m_2^0, \dots, m_N^0 \rangle$ to represent the capacities and agent distribution on the *original servers*, *i.e.* the servers allocated to the application at its initiation. The scaled workload of an agent a_i at phase l and its scaled workload change from phase $l-1$ to l on its original server are denoted by $w_{i,k}^0(l)$ and $\delta_{i,k}^0(l-1)$, respectively. We assume initially agents have equal scaled workload at phase $l = 0$.

To make the migration decision problem tractable, we make some simplifying assumptions.

Assumption 1: The initial scaled workload changes of different agents on a server, $\delta_{i,k}^0(\cdot)$, are i.i.d. random variables and they are distribution-free. Agent autonomy makes the construction of independent agents easier by utilizing data locality and migration.

Assumption 2: Service/agent migration keeps the probabilistic distribution of the scaled workload change of an agent unchanged. But it may have different mean after migration. This is reasonable because the scaled workload change equals to the amount of agent workload change relative to its capacity quota. This quota is a constant during the agent execution on a server with a fair-share scheduling scheme.

4 Service Migration Decision

A virtual machine is composed of a set of servers executing tasks of their residing agents and synchronizing with each other by performing barrier operations between phases. By the dynamic system in (3.1), it is expected that the servers’ scaled workload will change with time and cause overloaded states. Since the duration of a phase is determined by the heavily loaded servers, the overall system performance may deteriorate in phase. Service migrations aim to eliminate the performance bottleneck of a virtual machine by increasing capacities of the overloaded components. A service migration involves the tasks of transmitting service states and data along with the residing agents to a destination server with a higher capacity, continuing services and resuming agent execution there. Although we do not focus on the details of how to perform service migrations, it is clear that a migration operation would incur significant run-time overhead and the adaptive computation can not afford frequent migrations.

The objective of service migration decision is to obtain the optimal timing in phase to minimize the migration frequency, and the low bound of the destination server capacity to guide the server selection. In the following, we suppose there are only service migrations during the execution of a virtual machine. The agent migration decision problem and its interplay with service migration will be tackled in Section 5.1.

4.1 Migration Candidate Determination and Service Migration Timing

Although servers of a virtual machine need to perform global barrier synchronization between phases, their executions within a phase can be assumed to be independent. This is particular the case in agent-based systems, where autonomous agents perform their own tasks with local services and data sets. A service migration decision is made by a server based on its overall workload relative to its capacity. In the following, we use a set of simplified notations, in which s denotes any server in a virtual machine, c and m representing

its capacity and number of agents, $\langle a_1, a_2, \dots, a_m \rangle$ being the sequence of its agents, $w_i(l)$ and $\delta_i(l)$ denoting the scaled workload and workload change of a residing agent a_i at phase l .

For server s , its overall workload at a phase is the summation of all its agents' workloads, if we do not consider their overlapping. We define *expected relative workload* of server s at phase l as

$$r(l) = E\left[\frac{\sum_{i=1}^m w_i(l)}{m}\right], \quad (4.1)$$

Expected relative workload represents a server's overall workload w.r.t. its capacity. We use the scaled workload of each agent in order to ensure workload of different servers in the VM are comparable with each other.

Most of previous work designs migration strategy to minimize the degree of overload. Due to the non-negligible runtime overhead in migration, certain degree of overload must be tolerated. We reformulate the decision problem as "Given a bound of overload, find the minimal migration frequency". Therefore, the optimal migration timing is obtained by finding the maximal phase l^* for a given overload bound R , $0 < R < 1$, *i.e.*

$$r(l) \leq R. \quad (4.2)$$

R reflects the degree to which a server is considered to be overloaded. To find the maximal phase in (4.2), we need to calculate the expected value of scaled workload $w_i(l)$ for each residing agent. But the heterogeneity in terms of different server capacities and alteration of an agent's scaled workload due to service migration, add complexity to its calculation. Next, we try to derive the expression of scaled workload $w_i(l)$ with reference to the value $w_i^0(l)$ on the original server.

Consider server s makes a service migration at certain phase k . Let s' denotes the server after migration, in the *space domain*. We use (c', m') to represent its capacity and number of agents, respectively. For a residing agent, say a_i , its scaled workload and workload change after migration are denoted by $(w'_i(\cdot), \delta'_i(\cdot))$. The *workload conservation property* states the workload of an agent remains unchanged at the moment of migration. That is $w'_i(k)c'/m' = w_i(k)c/m$. In the *time domain*, with the additive scaled workload evolution model in Eq. (3.1), it is clear $w'_i(l) = w'_i(k) + \sum_{t=k}^{l-1} \delta'_i(t)$, for phases $l > k$. If there is no service migration at phase k , the scaled workload becomes $w_i(l) = w_i(k) + \sum_{t=k}^{l-1} \delta_i(t)$, for $l > k$. By Assumption 2, we have $\delta'_i(t)c'/m' = \delta_i(t)c/m$. Thus, we derive $w'_i(l) = w_i(l)cm'/(c'm)$. Together with (c^0, m^0) for the original server and $m = m^0$ due to no agent migration, the scaled workload of agent a_i on server s can be calculated as

$$w_i(l) = \frac{w_i^0(l)}{\tilde{c}}, \quad (4.3)$$

where $\tilde{c} = c/c^0$ is called the *relative capacity*. We can see that the scaled workload of an agent equals its initial scaled workload relative to the ratio of capacities between the current server and original one. This is caused by service migrations in the space domain.

Let μ_i be the means of scaled workload changes $\delta_i^0(\cdot)$, for $i = 1..m$. According to the additive workload evolution model, the scaled workload of agent a_i on its original server at phase l is $w_i^0(l) = w_0 + \sum_{t=0}^{l-1} \delta_i^0(t)$, where $w_0 = w_i^0(0)$ is scaled workload at the phase 0, *i.e.* the initial scaled workload of an agent on its original server. Since the scaled workload changes $\delta_i^0(\cdot)$ in different phases are i.i.d. random variables, we have the expected value of scaled workload $E[w_i(l)] = E[w_i^0(l)]/\tilde{c} = (w_0 + l\mu_i)/\tilde{c}$. Thus, the expected relative workload of server s becomes

$$r(l) = \frac{\sum_{i=1}^m (w_0 + l\mu_i)}{m\tilde{c}} = \frac{w_0 + l\bar{\mu}}{\tilde{c}}, \quad (4.4)$$

where $\bar{\mu} = (\sum_{i=1}^m \mu_i)/m$ is the average change rate of the scaled workload. If $\bar{\mu} \leq 0$, then the overall scaled workload of a server tends to remain unchanged or decrease, as the computational phase proceeds. It means there is no need to perform service migration any more. So, we only consider the cases with $\bar{\mu} > 0$ in our discussion. With $r(l) \leq R$, we have the following theorem to determine the optimal timing for service migrations.

Theorem 4.1 *For any server s with capacity c in a virtual machine, suppose the scaled workload changes of different agents on the original server, $\delta_i^0(\cdot)$, $i = 1..m$, have means μ_i and their average $\bar{\mu} > 0$. For a given overload bound R , $0 < R < 1$, of the expected relative workload, the optimal service migration timing l^* in phase, under Assumption 1 and Assumption 2, is*

$$l^* = \frac{\tilde{c}R - w_0}{\bar{\mu}}. \quad (4.5)$$

Intuitively, Theorem 4.1 states that the optimal service migration timing of a server equals to its available server capacity, apart from the initial workload, divided by the average rate of workload increase. Eq. (4.5) reflects the influence from both the time domain (by the average mean of scaled workload changes) and the space domain (by the relative server capacity \tilde{c}). It shows the timing is proportional to the relative server capacity. This is because both w_0 and each μ_i , are defined relative to capacity of the original server. To reduce the service migration frequency, we can select the destination server with the highest capacity. But this greedy strategy may allocate more resource to less heavily-loaded servers, which leads to resource imbalance in the virtual machine. We will discuss this tradeoff in Section 4.2.

4.2 Destination Server Selection

Servers in a virtual machine independently decide when to perform service migrations based on Theorem 4.1. At phases of the optimal service migration timing, their services and agents will be transmitted to the destination servers. As we have discussed in the previous subsection, the selection of a destination server is a tradeoff between the local and global performance optimizations. From an individual server's perspective, it is preferable to obtain an available server with the highest capacity as its destination so that its service migration frequency becomes minimal. But from the perspective of the entire virtual machine, it'd better allocate more powerful servers to more heavily-loaded servers. In this section, we introduce a migration gain function and try to find the minimal capacity, that a prospective destination server must have, for a given target gain value.

For server s in a VM, the execution time of its agent a_i at phase l , $t_i(l)$, is proportional to its workload relative to the capacity quota, *i.e.* its scaled workload, as $t_i(l) \propto w_i(l)$. The duration of a computational phase on a server is determined by its agent with the longest finish time, *i.e.* $t(l) = \max_{i=1..m} t_i(l)$. Consider server s makes a service migration to an available server, denoted by s' , with capacity c' at the optimal migration timing l^* according to Theorem 4.1. We require $c' > c$, otherwise this migration is meaningless w.r.t. the system performance improvement. Thus the execution time of each agent on server s' will be reduced. We define the *expected migration gain* for this service migration as

$$g(c') = E\left[\sum_{l=l^*}^{k-1} (t(l) - t'(l))\right] - E[O(w(l^*))], \quad (4.6)$$

where k is the prospective optimal migration timing for server s' with its supposed capacity c' according to Theorem 4.1. The execution time on server s' is $t'(l) \propto \max_{i=1..m} w'_i(l)$, and $O(\cdot)$ returns the migration overhead for a given scaled workload. The expected migration gain function describes the performance improvement, in terms of agents' execution time on the two servers until the next service migration, compared with the migration overhead. We do not use execution time of the entire virtual machine as a metric to express performance gain, due to the unpredictable behaviors of other servers in the VM and the independency of server execution. We expect the performance benefits in terms of a value G from a service migration. The target gain, G , is measured relative to the migration overhead. So, the selection of a destination server in a service migration is guided by finding the low bound of server capacity for a given non-negative target gain value G of the migration gain function, *i.e.*

$$g(c') \geq G. \quad (4.7)$$

A special case is to have $G = 0$, which means no benefit is yielded from a migration except to counteract the migration overhead. By increasing G , capacities of the destination servers will become higher accordingly. But to avoid resource allocation imbalance in the virtual machine, we need to choose the value of G appropriately. Since the state information of services on a server is usually maintained by a relatively limited number of variables, the main portion of the service migration overhead is caused by transferring codes and data of services. The size of data set is likely to change little since its initial assignment to a server and the code of services are immutable. So, the migration overhead of a server can be represented by a constant, denoted by H . However, our optimization approach is still applicable with the general overhead function, $O(\cdot)$. Since the execution time of an agent, $t_i(l)$, is solely determined by its scaled workload $w_i(l)$, we will use these two terms interchangeably. Next, we transform the migration gain function to find the minimal capacity of the destination server for a given target gain value.

With the expressions of $t(l)$, $t'(l)$ and $w_i(l)$, the expected migration gain (4.6) becomes

$$g(c') = E\left[\sum_{l=l^*}^{k-1} \left(\max_{i=1..m} \frac{w_i^0(l)}{\tilde{c}} - \max_{i=1..m} \frac{w_i^0(l)}{\tilde{c}'}\right)\right] - H = \left(\frac{1}{\tilde{c}} - \frac{1}{\tilde{c}'}\right) \sum_{l=l^*}^{k-1} E\left[\max_{i=1..m} w_i^0(l)\right] - H, \quad (4.8)$$

where $\tilde{c}' = c'/c^0$. Let μ_i and σ_i^2 be the mean and variance of the scaled workload change $\delta_i^0(\cdot)$. Since $w_i^0(l) = w_0 + \sum_{t=0}^{l-1} \delta_i^0(t)$ and the phase-wise scaled workload changes of each agent $\delta_i^0(\cdot)$ are i.i.d. random variables w.r.t. phases, the central limit theorem of statistics [19] guarantees that as l gets large, the distribution of $w_i^0(l)$ tends to become normally distributed with mean $w_0 + l\mu_i$ and variance $l\sigma_i^2$. Because $\delta_i^0(\cdot)$, for $i = 1..m$, are independent random variables, scaled workloads $w_i^0(\cdot)$ are also independently distributed.

So we have $Pr\{\max_{i=1..m} w_i^0(l) \leq x\} = \prod_{i=1}^m Pr\{w_i^0(l) \leq x\}$. Since each $w_i^0(\cdot) \geq 0$, it is clear

$$E\left[\max_{i=1..m} w_i^0(l)\right] = \int_0^\infty Pr\{\max_{i=1..m} w_i^0(l) > x\} dx = \int_0^\infty \left[1 - \prod_{i=1}^m \int_{-\infty}^x f_i(u) du\right] dx, \quad (4.9)$$

where $f_i(u)$ is the probability density function of normal distribution with parameters $w_0 + l\mu_i$ and $l\sigma_i^2$ for scaled workload $w_i^0(l)$. Therefore, we can derive the low bound of the destination server capacity, as shown in the following theorem.

Theorem 4.2 *For any server s with capacity c in a virtual machine, suppose the scaled workload changes of different agents on the original server, $\delta_i^0(\cdot)$, $i = 1..m$, have means μ_i and variances σ_i^2 . For a given*

non-negative target gain G of the expected migration gain function, the minimal capacity that a prospective destination server must have during a service migration at the optimal phase l^* , under Assumption 1 and Assumption 2, is

$$c^* = \max\{c^+, c + \frac{\bar{\mu}}{R}c^0\}, \quad (4.10)$$

where R is the overload bound of the expected relative workload in (4.2) and c^+ is the solution to equation:

$$\left(\frac{1}{\tilde{c}} - \frac{1}{\tilde{c}'}\right) \sum_{l=l^*}^{k-1} \left[\int_0^\infty \left(1 - \prod_{i=1}^m \int_{-\infty}^x f_i(u) du\right) dx \right] - H = G, \quad (4.11)$$

in which the optimal service migration phases $l^* = (\tilde{c}R - w_0)/\bar{\mu}$ and $k = (\tilde{c}'R - w_0)/\bar{\mu}$ by Theorem 4.1.

Proof: By applying $E[\max_{i=1..m} w_i^0(l)]$ in Eq. (4.9) to the migration gain function, we get the left-hand side part of Eq. (4.11). And it is clear the resulting gain function monotonously increases as its parameter c' becomes greater. So for a given target gain value G , the minimal capacity satisfying $g(c') \geq G$ is the solution to equation $g(c') = G$, denoted by c^+ .

At the same time, the prospective optimal migration phase k for the new server s' should be greater than phase l^* for server s . Otherwise, a new service migration will occur as soon as the previous one completes. We call this phenomena *idle migrations*. So, we have $k \geq l^* + 1$, i.e. $(R\tilde{c}' - w_0)/\bar{\mu} \geq (R\tilde{c} - w_0)/\bar{\mu} + 1$ and obtain $c' \geq c + \bar{\mu}c^0/R$. Therefore, the minimal capacity of a destination server takes the form in the theorem. \square

The expected gain function is hard to calculate by Eq. (4.9). A special case is when the scaled workload changes $\delta_i^0(\cdot)$, for $i = 1..m$ are i.i.d. random variables with the same mean and variance. This is reasonable for the SPMD applications, in which agents execute the same program on their own data sets and they have similar workload change behaviors. For these applications, the minimal destination server capacity can be determined by the following corollary.

Corollary 4.1 *For any server s with capacity c in a virtual machine, suppose the scaled workload changes of different agents on the original server, $\delta_i^0(\cdot)$, $i = 1..m$, have the same mean μ and variance σ^2 . For a given non-negative target gain G of the expected migration gain function, the minimal capacity that a prospective destination server must have during a service migration at the optimal phase l^* , under Assumption 1 and Assumption 2, is determined by Eq. (4.10), where c^+ is the solution to equation:*

$$\left(\frac{1}{\tilde{c}} - \frac{1}{\tilde{c}'}\right) \sum_{l=l^*}^{k-1} [w_0 + l\mu + \alpha(m)\sigma\sqrt{l}] - H = G, \quad (4.12)$$

in which $\alpha(m) = (2 \ln m)^{1/2} - (\ln \ln m + \ln 4\pi)/[2(2 \ln m)^{1/2}] + \gamma/(2 \ln m)^{1/2}$ and γ is Euler's constant (0.5772...).

Proof: Since the scaled workload change $\delta_i^0(\cdot)$, $i = 1..m$, have the same mean μ and variance σ^2 , the scaled workload $w_i^0(l)$, $i = 1..m$ are i.i.d. normally distributed random variables with the same mean $w_0 + l\mu$ and variance $l\sigma^2$. According to the extreme value theory [1], we have

$$E\left[\max_{i=1..m} w_i^0(l)\right] \approx w_0 + l\mu + \alpha(m)\sigma\sqrt{l}, \quad (4.13)$$

Therefore, the expected migration gain equals to the left-hand side of Eq. (4.12). Thus the minimal destination server capacity is determined according to Theorem 4.2. \square

We can see the minimal destination server capacity ensures the performance improvement due to a service migration equals to a given target gain in addition to the migration overhead. Since servers make their migration decisions independently, it is possible that two or more servers may decide to migrate to the same destination server according to Theorem 4.2. This phenomena is called *destination confliction* in service migration. To solve the confliction and reduce migration frequency, we allow the server with the largest prospective next optimal migration timing, determined by Theorem 4.1, to take the destination server. Other servers will try to choose those from the remaining available servers with the least sufficient capacities. After selecting a destination server and completing service migration to it, we update the server sequence of the VM and the capacity distribution sequence C . Since then, the renewal process begins a new round of virtual machine computation and each server determines its next migration timing.

5 Hybrid Migration Decision

The properties of autonomy and multi-hop mobility make mobile agent as a suitable technique to achieve load balancing in distributed computing. Like service migration, there is also a decision problem for agent migration. But the dynamic environment with server reconfiguration and interaction among different agents greatly add to the difficulty of its decision. In this section, we model the agent migration decision problem by dynamic programming and extend the service migration decision to incorporate hybrid mobility.

5.1 Agent Migration Decision

To adapt to the change of server capacities, agents may need to migrate. The decision question is when to migrate and which server the agent should travel to. A locally optimal strategy can be derived if we assume individual agents make their migration decisions independently. But unlike service migration, different agents may migrate to the same server. The interaction among them may cause this migration strategy useless. So, certain global information about the agent or workload distribution is needed. If we consider the global optimization of agent migration, the decision problem is equivalent to the task scheduling in distributed computing. In general, it is a NP-problem to find the globally optimal task distribution. However, in the bulk synchronous computation model for SPMD, agents proceed their computation in phases and their carry out the same task at each phase. It is possible to find a feasible optimal solution to agent migration with some global information. We relate the coordinated decision problem of agent migration to the remapping problem in parallel computing and use dynamic programming to derive a globally optimal solution.

We define the *state* of a VM as the workload distribution of different servers at certain phase. Assume the choice of agent migration while in state v incurs a cost $C(v, u)$, where u is a binary decision: migrate or not. $C(v, u)$ may be random. The decision process then passes into another state. The probability $p_{vq}(u)$ of passing into state q from v is dependent on the action chosen in state v . The expected total cost of a decision policy is the expected sum of the costs incurred at each decision step. An optimal decision policy minimizes the expected total cost.

Let $\mathcal{J}(v)$ be the expected total cost of the VM which starts in state v , and which is governed by the control

decisions. Then

$$\mathcal{J}(v) = \min_u \{C(v, u) + \sum_{q \in I} p_{vq} \mathcal{J}(q)\}, \quad (5.1)$$

where I is the set of states that the VM may enter from state v due to agent migration. For any decision process state v , let $next(v)$ denote the set of states reachable from v in one phase and $travel(v)$ be the entry state of an agent migration. Each decision state $r = \langle r_1, \dots, r_N, l+1 \rangle \in next(\langle v_1, \dots, v_N, l \rangle)$ has a transition probability

$$Pr\{r|v\} = \prod_{i=1}^N Pr\{r_i|v_i\},$$

where $Pr\{r_i|v_i\}$ is the probability of chain i passing from state v into state r_i in one phase. The execution time of the VM at state v is determined by the most heavily-loaded server, *i.e.* $t(v) = \max_i \{v_i\}$. Each v_i equals to the scaled workload of the agent with the longest finish time. That is $v_i = \max_k \{w_{i,k}(l)\}$. Eq. (5.1) can be re-written as

$$\mathcal{J}(v) = \min_u \begin{cases} t(v) + \sum_{v' \in next(v)} Pr(v'|v) \mathcal{J}(v') \\ t(v) + C(v, u) + \sum_{v' \in next(travel(v))} Pr(v'|v) \mathcal{J}(v'), \end{cases} \quad (5.2)$$

where the bottom equation on the right-hand side is the cost function associated with agent migration, and the top equation is associated with no migration.

According to the theory of Markov decision processes, the optimal decision to make from state v is the decision which minimizes the right-hand side of Eq. (5.2). If the number of phases the agent takes is some random variable with finite mean, then the system of equations given by (5.2) can be solved. The iteration algorithm proposed in [30] can be applied to solve the equation.

5.2 Interplay between Service and Agent Migration

With the introduction of agent migration, the set of residing agents on each server of a virtual machine will not be immutable any more. An agent may land on or leave from a server at any phase. So the service migration decision problem is quite different from the discussion in Section 4.

An agent may decide to migrate to a server that will perform a service migration to another server at the same phase. So the information, *e.g.* server capacity, for the agent to make decision may be incorrect at the moment of its migration. To avoid this problem, we require service migrations are decided and performed first, and after their completion, agents decide whether to migrate or not with the new system information. Now the agent distribution sequence $M = \langle m_1, m_2, \dots, m_N \rangle$ of a VM is no longer a constant. Instead, we treat it as a sequence of random variables whose distributions are determined by the dynamic agent migration strategy discussed in Section 5.1. At certain computational phase, the residing agents on server s are denoted by $\langle a_1, a_2, \dots, a_m \rangle$. They may come from different servers via agent migration. These agents are assigned to their original servers, represented by $\langle s_{(1)}, s_{(2)}, \dots, s_{(m)} \rangle$, at phase $l = 0$. Entries in this sequence may be identical, which reflects some agents come from the same server. The optimal service migration timing for a server in Theorem 4.1 will be extended by considering the dynamic composition of its agent set. As we can see, it is quite complicated to make optimal decisions with both service and agent migrations for general applications. Here we consider a simplified case where agents on their original servers have the same capacity quota. It can be realized by the task scheduler at the system initiation. The optimal service migration timing with agent migration is described in the following theorem.

Theorem 5.1 For any server s with capacity c in a virtual machine, its m agents may migrate from different servers. Suppose these residing agents are, at the initial phase, have equal initial capacity quota q_0 . Their scaled workload changes, $\delta_i^0(\cdot)$, $i = 1..m$, have the same mean μ . For a given overload bound R , $0 < R < 1$, of the expected relative workload, the optimal service migration timing l^* in phase, under Assumption 1 and Assumption 2, is

$$l^* = \frac{1}{\mu} \left(\frac{cR}{q_0 E[m]} - w_0 \right). \quad (5.3)$$

Proof: For a residing agent a_i on server s , suppose it originally resides on server $s_{(i)}$, whose capacity and initial number of agents are $c_{(i)}^0$ and $m_{(i)}^0$, respectively. Its initial capacity quota $q_0 = c_{(i)}^0/m_{(i)}^0$. According to the previous discussion, its scaled workload at phase l is $w_i(l) = w_i^0(l)c_{(i)}^0 m/(cm_{(i)}^0) = w_i^0(l)q_0 m/c$. Since $w_i^0(l)$ tends to have a normal distribution with mean $w_0 + l\mu$, we have $E[w_i^0(l)q_0] = (w_0 + l\mu)q_0$. Similarly, the initial scaled workloads of other agents on server s are also normally distributed. The m random variables $w_1^0(l)q_0, \dots, w_m^0(l)q_0$ are i.i.d. with the same mean $(w_0 + l\mu)q_0$. Since m is also a random variable, we have

$$\begin{aligned} r(l) &= \frac{1}{c} \cdot E \left[\sum_{i=1}^m w_i^0(l)q_0 \right] = \frac{1}{c} \cdot E \left[E \left[\sum_{i=1}^m w_i^0(l)q_0 \mid m \right] \right] = \frac{1}{c} \cdot E \left[m E \left[w_i^0(l)q_0 \right] \right] \\ &= \frac{1}{c} \cdot E[m] E \left[w_i^0(l)q_0 \right] = \frac{1}{c} q_0 (w_0 + l\mu) E[m]. \end{aligned}$$

By applying the above equation to $r(l) \leq R$, the optimal service migration timing l^* of server s can be derived as in the theorem. \square

According to Theorem 5.1, we can see the optimal service migration timing is reversely proportional to the expected value of agent number m . This can be explained that as the number of agents on a server increases, its overall workload tends to become greater and this leads to a reduced migration interval. Compared with Theorem 4.1, Eq. (5.3) is additionally related to the initial capacity quota of each agent due to agent and service migration. The destination server selection policy in Section 4.2 can also be extended when incorporating agent migration.

6 Simulation Results

To verify and analyze results of the migration decision according to our optimal model, we performed experiments to simulate the random processes of agent execution in a reconfigurable VM with 30 agents on each server. Simulations were conducted using Matlab and each data point of the simulation results was an average of 200 runs. The 95% confidence interval for each simulation result is also presented to demonstrate the robustness of the estimates. The distribution functions of agent workload evolution characterize the workload dynamics in three of various applications.

Case 1. In the first experiment, we were to find and verify the optimal computational phase before a service migration according to Theorem 4.1 with the objective of minimizing the migration frequency. The agents on a server were divided into three groups with 10 agents in each one. They might be run in different phases and their initial scaled workload changes obeyed the following distribution functions respectively.

$$\delta_1^0(l) = \begin{cases} 0.1 & \text{w.p. 0.50,} \\ 0 & \text{w.p. 0.50.} \end{cases} \quad \delta_2^0(l) = \begin{cases} 0.1 & \text{w.p. 0.50,} \\ 0 & \text{w.p. 0.25,} \\ -0.1 & \text{w.p. 0.25.} \end{cases} \quad \delta_3^0(l) = \begin{cases} 0.1 & \text{w.p. 0.25,} \\ 0 & \text{w.p. 0.50,} \\ -0.1 & \text{w.p. 0.25,} \end{cases}$$

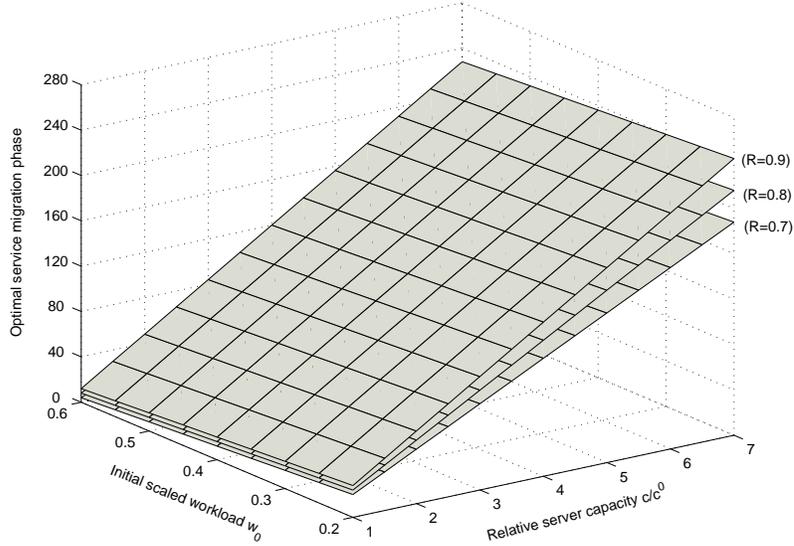


Figure 2: Optimal service migration timing (numerical results).

where w.p. means "with probability". The distribution of scaled workload change $\delta_1^0(l)$ implies the total workload of an agent keeps increasing. This is valid in the early phase of some search algorithms, such as the branch-and-bound method [12]. The branching procedure recursively expands a (sub)problem into subproblems such that the size of subproblem set is on the steady increase. The second distribution $\delta_2^0(l)$ indicates the bounding procedure computes the expanded subproblems, which completes part of agent's workload. But the agent tends to generate new subproblems with higher probability as the computation proceeds. In the third phase $\delta_3^0(l)$, the rates of subproblem generation and completion tend to be equal so that the agent's expected workload remains unchanged for a certain period.

We varied the initial scaled workload w_0 from 0.2 to 0.6 and the capacity of each server relative to the original server of its residing agents, \tilde{c} , from 1 to 7. The corresponding optimal migration timing changed from the 4th to 244th phase as shown in Fig. 2. The bound of overload, R , indicates the degree to which certain overload must be tolerated by a server to avoid frequent migration. Its value should be set according to the runtime overhead of service migration. By varying the overload bound R from 0.7 to 0.9, we can see the migration timing increases. Even with $w_0 = 0.5$, $R = 0.7$ and $\tilde{c} = 4$, the optimal migration phase is the 100th, at which much computation has been performed. So service migration is a feasible approach to realize virtual machine reconfiguration. Fig. 3 shows the migration phases simulated by random processes when the initial scaled workload w_0 is 0.5. This does not affect the experiment results. We can see our migration decision provides an accurate low bound in determining the service migration timing. The difference between simulated and predicted results tends to decrease as the overload bound R increases. The small confidence interval indicates the robustness of the estimates.

Case 2. The purpose of the second experiment is to verify the accuracy of destination server selection by Theorem 4.2 for a service migration. The agent composition and distribution functions of scaled workload change were the same as those in Case 1. The initial scaled workload of agents and overload bound were set as $w_0 = 0.5$ and $R = 0.8$ respectively. The service migration overhead was $H = 50$, as 100 times greater than the initial scaled workload. It confirms to our measurements of the runtime overhead of service migration in M-DSA [8].

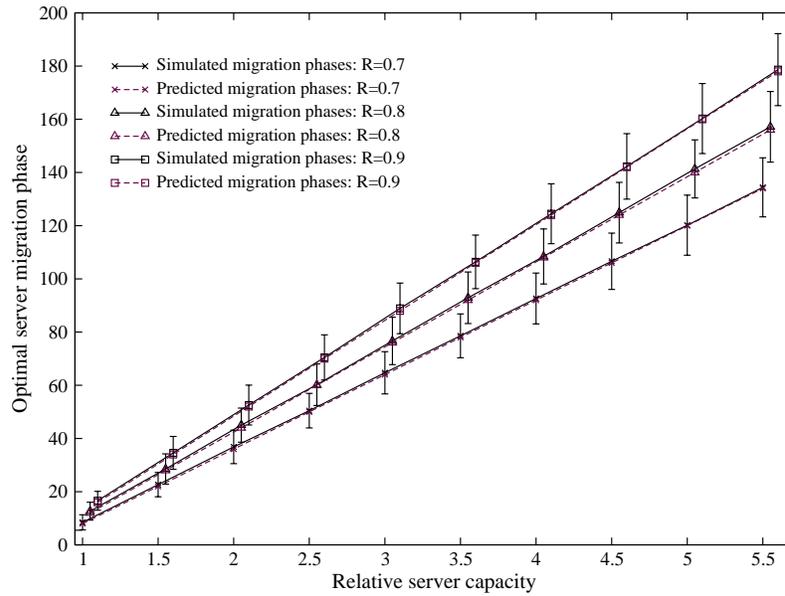


Figure 3: Optimal service migration timing with $w_0 = 0.5$ (simulation results) .

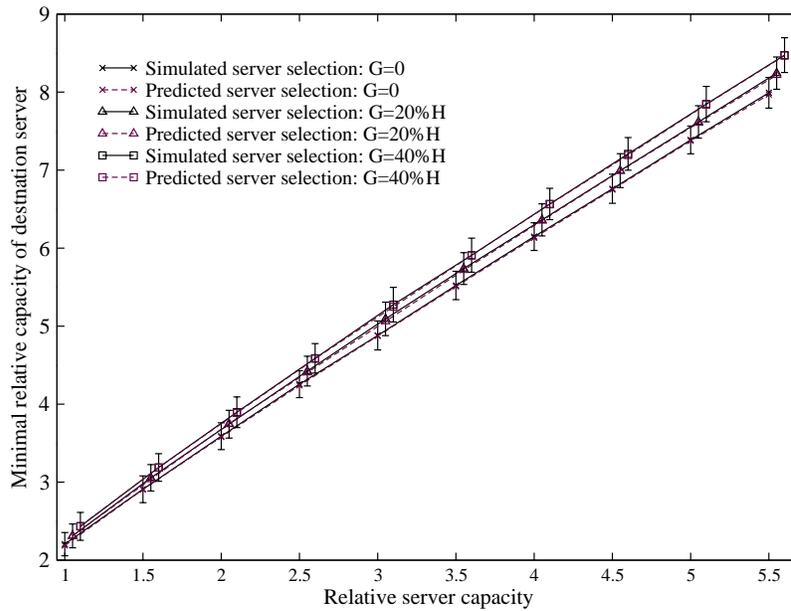


Figure 4: Destination server selection.

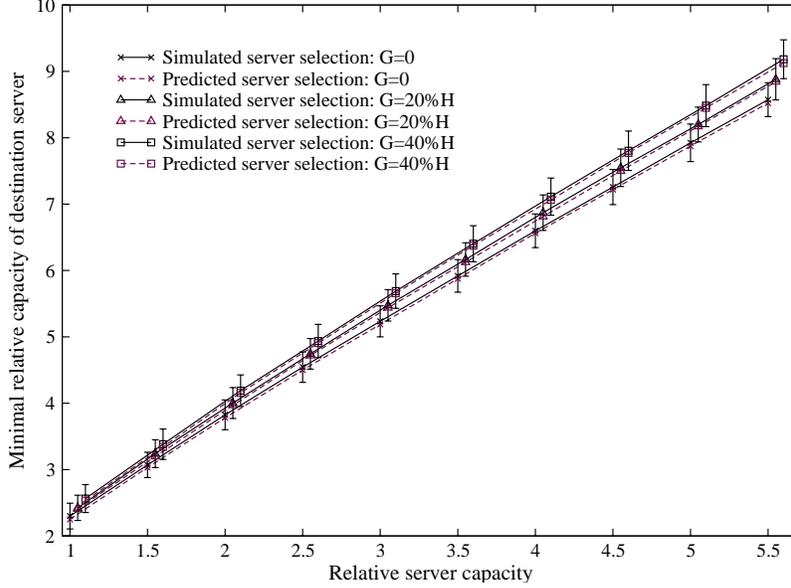


Figure 5: Destination server selection with extreme value theory.

Fig. 4 plots the minimal capacity that a prospective destination server must have. We varied the capacity of a server from 1 to 5.5 relative to the original server and measured the minimal relative capacity of destination server. We can see the predicted value is either equal to or slightly smaller than the simulated result. This is because our migration decision finds the low bound of destination server capacity. The difference is caused by approximating the scaled workload of an agent by a normal distribution according to the central limit theorem. From the figure, we can also see that the minimum value of destination capacity increases at a sub-linear rate. The target gain value G represents the expected benefits from service migration and it is measured relative to the migration overhead. As shown in the figure, if the target gain G is set to 20% of the migration overhead and the relative capacity of a server is 1, the minimal relative capacity of a destination server is about 2.3. Then during the next optimal migration timing, this service will be further migrated to an available server whose capacity should be at least 4.0 times greater than the original server of its residing agents. After migration, the process will be renewed and agent tasks will be executed until the next migration.

Case 3. In the third experiment, we repeated the second experiment, but assuming all agents of a server had the same distribution function of scaled workload change as

$$\delta^0(l) = \begin{cases} 0.1 & \text{w.p. } 0.50, \\ 0 & \text{w.p. } 0.25, \\ -0.1 & \text{w.p. } 0.25. \end{cases}$$

According to Corollary 4.1, the problem of finding the lower bound of destination server capacity can be approximated by using the extreme value theory.

Fig. 5 shows the accuracy of the prediction approximation. The results in solid line are due to simulation measurements. Compared with the predicted values in Case 2, Corollary 4.1 is a little less accurate in destination server selection. It is because the extreme value theory provides an upper bound approximation of the expected value for maximal jointly distributed random variables. But, in the figure, we can see the predicted results are still very close to the simulated ones. The advantage of applying Corollary 4.1 is to simplify calculation of the expected migration gain function.

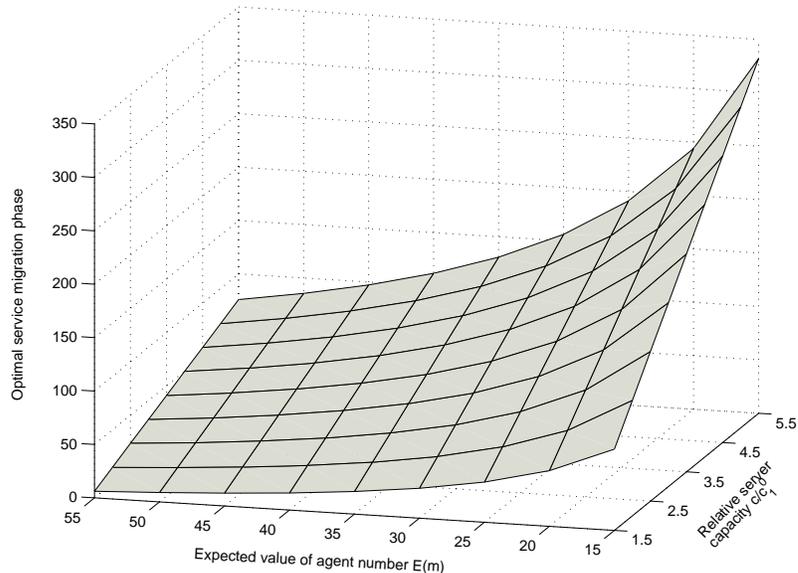


Figure 6: Optimal service migration timing for hybrid mobility (numerical results).

Case 4. The decision problem of agent migration has been modelled by dynamic programming in Section 5.1. It can be solved by an approach proposed in [30]. To analyze the interplay between agent and service migration, we conducted the fourth experiment. We applied Theorem 5.1 to extend the optimal service migration decisions in Case 1 and conducted simulations to verify the results. To meet the premise of the theorem, we assigned 30 agents to each server and ensured each agent was allocated the same initial capacity quota. All other settings were equivalent to those in Case 3.

Fig. 6 plots the optimal service migration timing by incorporating agent migration. We varied the capacity ratio between a server and the original one of the residing agents. For each ratio value, we measured the optimal migration timing corresponding to different expectations of the agent number, $E[m]$. This expected value represents the dynamics of agent migration. From the figure, we can see the migration timing decreases as the expectation of agent number increases. An extreme was reached when $E[m] = 55$ and server capacity ratio was 1.5. Its migration phase was about 6. This was resulted from many agents landing on this server and the overall workload rapidly overwhelming its capacity. As computation proceeded, agents might decide to migrate to other light-loaded servers. Thus, the server's next migration interval became expanded accordingly. The expected value of agent number for each server can be determined dynamically by the history information of agent membership during the computation. Fig. 7 presents the simulation results. We measured the migration timing for the capacity ratio equal to 1.5, 3.0 and 4.5. We can see the difference between simulated and predicted phases tends to be small around $E[m] = 30$, which is the initial number of agents on each server.

7 Conclusions

In this paper, we have studied the decision problem of service and agent migration for a reconfigurable virtual machine. The migration timing and destination server selection of service migration are modelled as stochastic optimization models. Thus, we have derived optimal migration policy in both time and space

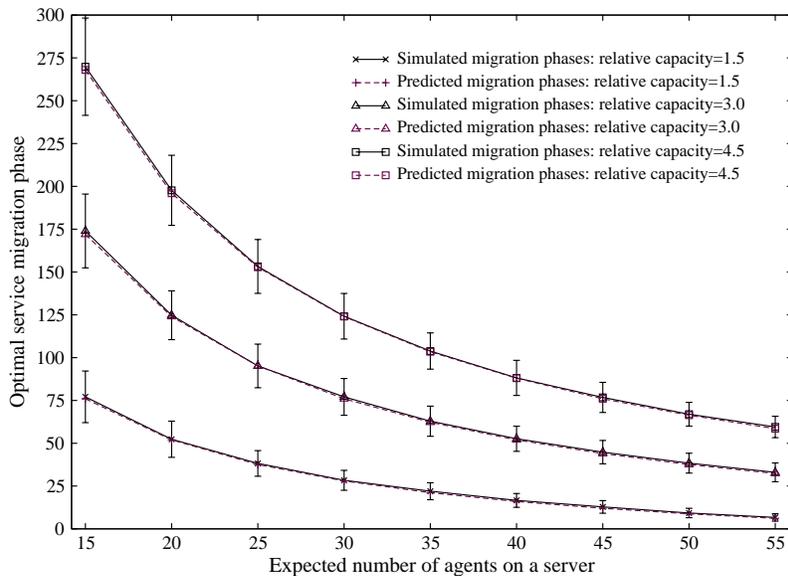


Figure 7: Optimal service migration timing for hybrid mobility with $w_0 = 0.5$ and $R = 0.8$.

domains for given overload bound and target gain value. We have formulated the agent migration decision problem by dynamic programming and analyzed the interplay of the two migration problems.

We note that although the assumption of scaled workload changes remaining the same distributions during migration was assumed in literature, determination of a distribution is non-trivial. Armstrong *et al.* experimented with the NAS benchmarks on different machines to determine the distribution types of the execution time [2]. They found that running the same job on different machines led to different execution time means and different distribution types. We have analyzed the impact of distribution alteration of scaled workload changes during service migration to the migration decision policy and derived the optimal migration timing in [9]. However, it becomes quite complicated when considering the destination server selection and interplays between service and agent migration. We will look for alternative approaches to find the optimal migration policy for them. The present solution to destination confictions, discussed in Section 4.2, is still primitive. It is difficult to make a globally optimal decision when several servers decide to migration to the same destination server according to the minimal capacity by our decision scheme. The solution is complicated by the dynamics of hybrid migration. New approaches will be integrated to our optimal model. In addition, we will investigate the support of fault tolerance by hybrid mobility, which is an important application of reconfigurable VM systems.

References

- [1] A. H.-S. Ang and W. H. Tang. *Probability Concepts in Engineering Plannig and Design*, volume II. Rainbow Bridge, 1984.
- [2] R. Armstrong, D. Hensgen, and T. Kidd. The relative performance of various mappng algorithms in independent of sizable variances in run-time predications. In *Proc. of the 7th Heterogeneous Computing Workshop*, 1998.

- [3] R. Basharahil, B. Wims, and C.-Z. Xu. Distributed shared array: An integration of message passing and multithreading on smp clusters. *Journal of Supercomputing*. (to appear), 2003.
- [4] J. Cao, D. P. Spooner, S. A. Jarvis, S. Saini, and G. R. Nudd. Agent-based grid load balancing using performance-driven task scheduling. In *Proc. of the Parallel and Distributed Processing Symposium (IPDPS)*, Nice, France, April 2003.
- [5] F. Dabek, M. F. Kaashoek, D. Karger, R. Morris, and I. Stoica. Wide-area cooperative storage with CFS. In *Proc. of the 18th ACM Symposium on Operating Systems Principles (SOSP)*, Chateau Lake Louise, Banff, Canada, October 2001.
- [6] R. Figueiredo, P. Dinda, and J. Fortes. A case for grid computing on virtual machines. In *Proc. of the 23rd International Conference on Distributed Computing Systems (ICDCS)*, May 2003.
- [7] N.-T. Fong, C.-Z. Xu, and L. Wang. Optimal periodic remapping of bulk synchronous computations on multiprogrammed distributed systems. In *Proc. of the 14th IEEE Int'l Parallel and Distributed Processing Symposium (IPDPS)*, pages 103–108, April 2000. To appear in JPDC.
- [8] S. Fu and C.-Z. Xu. M-DSA: A mobile distributed shared array system. Technical report, Department of Electrical and Computer Engineering, Wayne State University, 2003.
- [9] S. Fu and C.-Z. Xu. Migration decision for hybrid mobility in reconfigurable virtual machines. Technical Report CIC-09-03, Department of Electrical and Computer Engineering, Wayne State University, September 2003.
- [10] T. Garfinkel, B. Pfaff, J. Chow, M. Rosenblum, and D. Boneh. Terra: A virtual machine-based platform for trusted computing. In *Proc. of the 19th ACM Symposium on Operating Systems Principles (SOSP)*, October 2003.
- [11] H. Jiang and V. Chaudhary. Compile/run-time support for thread migration. In *Proc. of the 16th Int'l Parallel and Distributed Processing Symposium (IPDPS)*, April 2002.
- [12] E. L. Lawler and D. E. Wood. Branch and bound methods, a survey. *Operations Research*, 14:699–719, 1966.
- [13] R. A. Meyer and L. H. Seawright. A virtual machine time sharing system. *IBM System Journal*, 9(3):199–218, 1970.
- [14] D. Milojevic, F. Douglass, Y. Panedeine, R. Wheeler, and S. Zhou. Process migration. *ACM Computing Surveys*, 32(3):241–299, 2000.
- [15] K. Moizumi and G. Cybenko. The travelling agent problem. *Mathematics of Control, Signals and Systems*, 14(3):213–232, 2001.
- [16] D. M. Nicol and J. H. Saltz. Dynamic remapping of parallel computations with varying resource demands. *IEEE Transactions on Computers*, 37(9):1073–1087, September 1988.
- [17] W. Obeloer, C. Grewe, and H. Pals. Load management with mobile agents. In *Proc. of the 24th EUROMICRO Conference*, pages 1005–1012, Sweden, August 1998.
- [18] A. Rao, K. Lakshminarayanan, S. Surana, R. Karp, and I. Stoica. Load balancing in structured p2p systems. In *Proc. of the 2nd International Workshop on Peer-to-Peer Systems (IPTPS)*, February 2003.
- [19] S. M. Ross. *Introduction to Probability Models*. Academic Press, 6 edition, 1997.

- [20] C. P. Sapuntzakis, R. Chandra, B. Pfaff, J. Chow, M. S. Lam, and M. Rosenblum. Optimizing the migration of virtual computers. In *Proc. of the 5th Symposium on Operating Systems Design and Implementation (OSDI)*, December 2002.
- [21] E. G. Sirer, R. Grimm, A. J. Gregory, and B. N. Bershad. Design and implementation of a distributed virtual machine for networked computers. In *17th ACM Symposium on Operating Systems Principles (SOSP)*, pages 202–216, December 1999.
- [22] J. E. Smith. *An Overview of Virtual Machine Architectures*, October 2001. <http://www.ece.wisc.edu/jes/papers/vms.pdf>.
- [23] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proc. of the ACM SIGCOMM '01 Conference*, San Diego, California, August 2001.
- [24] C. A. Waldspurger and W. E. Weihl. Lottery scheduling: Flexible proportional-share resource management. In *Proc. of the 1th Symposium on Operating Systems Design and Implementation (OSDI)*, pages 1–11, November 1994.
- [25] D. Wang, N. Paciorek, and D. Moore. Java-based mobile agents. *CACM*, 42(3):92–102, March 1999.
- [26] A. Whitaker, M. Shaw, and S. Gribble. Denali: Lightweight virtual machines for distributed and networked applications. In *Proc. of the USENIX Technical Conference*, Monterey, CA, June 2002.
- [27] C.-Z. Xu and F. C. Lau. *Load Balancing in Parallel Computers: Theory and Practice*. Kluwer Academic Publishers, 1997.
- [28] C.-Z. Xu, L. Wang, and N.-T. Fong. Stochastic predication of execution time for dynamic bulk synchronous computation. *Journal of Supercomputing*, 21(1):91–103, January 2002.
- [29] C.-Z. Xu and B. Wims. A mobile agent based push methodology for global parallel computing. *Concurrency: Practice and Experience*, 14(8):705–726, July 2000.
- [30] G. Yin, C.-Z. Xu, and L. Wang. Optimal remapping in dynamic bulk synchronous computations via a stochastic control approach. *IEEE Transactions on Parallel and Distributed Systems*, 14(1):51–62, January 2003.
- [31] W. Zhu, C.-L. Wang, and F. C. M. Lau. JESSICA2: A distributed java virtual machine with transparent thread migration support. In *Proc. of the 4th Int'l Conference on Cluster Computing (CLUSTER 2002)*, September 2002.
- [32] X. Zhuang and S. Pande. Compiler scheduling of mobile agents for minimizing overheads. In *Proc. of the 23rd International Conference on Distributed Computing Systems (ICDCS)*, May 2003.