

## Chapter 2

# VM Placement via Resource Brokers in a Cloud Datacenter

### 2.1 Introduction

Resource management in cloud datacenters is one of the most important issues for cloud service providers because it directly affects their profit. Energy and performance guarantee are two major concern of it. In energy aspect, the total estimated energy bill of datacenters is \$11.5 billion and their energy bills double every five years [1, 2]. Also, in performance guarantee aspect, many researches insist that performance metrics such as throughput and response time should be considered as well as availability in IaaS SLA [3, 4]. If the IaaS SLA with the performance metrics are applied in public CSPs, performance management should be much more delicate to avoid the SLA penalty cost. Especially in VM placement in the cloud, an application quality of service (QoS)-based approach to allocate workload fairly in physical machines (PMs) and a power-based approach to consolidate VMs maximally are two basic goals respectively [5]. To reduce energy consumption in a cloud datacenter, dynamic right sizing (DRS) is a promising technology to dynamically adjust the number of active servers (i.e., servers whose power is switched on) in proportion to the measured user demands [6]. In DRS, energy saving can be achieved by enabling idle compute nodes that do not have any running VM instances to go into low-power mode (i.e., sleep or shut down). In order to maximize the energy efficiency via DRS, one of the primary adaptive resource management strategies is VM consolidation in which running VM instances can be dynamically integrated into the minimal number of compute nodes based on their resource utilization collected by a hypervisor monitoring module [7]. However, it is difficult to efficiently manage cloud resources because cloud users often have heterogeneous resource demands underlying multiple service applications which experience highly variable workloads. Therefore, inconsiderate VM consolidation might lead to undesirable performance degradation due to workload overloading. Although A virtualization technology allows CSPs to get many benefits such as getting high PM resource utilization via server consolidation and

elasticity in their resource usage. In virtualized environments, to guarantee performance in VMs sharing a PM, a hypervisor should provide isolation between VMs. However, while security isolation, fault isolation, and environment isolation are well guaranteed, the current virtualization technology does not provide effective performance isolation. Therefore, the VMs make an effect with each other and it causes performance degradation in the VMs. The phenomenon is called as performance interference [8].

Therefore, addressing the conflict between VM consolidation and dispersion is essential for effective cloud resource management. To achieve it, Patal and Shah [9] argued the need of cost modeling of a datacenter, and constructed the cost model. In the model, the total cost in a cloud datacenter includes the space, the power and the cooling recurring, and other cost. The other cost consists of maintenance and amortization of power and cooling system, personnel, software licenses, compute equipment depreciation, and so on.

In addition, we consider IaaS service level agreement (SLA) cost to the model. Although almost all cloud service providers only consider availability of VMs as IaaS SLA today, necessity of considering performance in IaaS SLA is raising as in [3, 4]. To achieve it, we define the SLA penalty cost ( $PC$ ) of compute node  $j$ ,  $PC_j(t)$  as depicted in Eq. (2.1) where  $UPC$  is the unit  $PC$ ,  $v_j(t)$  is resource usage of VMs in timeslot  $t$ , and  $\bar{d}_j(t)$  is the average performance degradation in timeslot  $t$ .

$$PC_j(t) = UPC \cdot v_j(t) \cdot \bar{d}_j(t). \quad (2.1)$$

Finally, we describe the partially total cost ( $PTC$ ) which is composed of costs only depending on VM placement in the total cost as depicted in Eq. (2.5). We note that the  $OC$  is an abbreviation of the PM operating cost including the power and the cooling recurring cost. In Eq. (2.2),  $J$  is the set of PMs in a cloud datacenter,  $UOC$  is the unit  $OC$ ,  $V_j(t)$  is a vector of  $v_j(t)$ , and  $p_j(V_j(t))$  is power consumption of node  $j$  when resource usage in timeslot  $t$  is  $V_j(t)$ .

$$\begin{aligned} PTC(t) &= \sum_{j \in J} OC_j(t) + PC_j(t) \\ &= \sum_{j \in J} UOC \cdot p_j(V_j(t)) + UPC \cdot v_j(t) \cdot \bar{d}_j(t). \end{aligned} \quad (2.2)$$

Based on the cost model, we formulate an optimization problem as depicted in Eqs. (2.3) and (2.4). The objective is to minimize the PM operating cost while keeping performance degradation less than the threshold.

$$\text{minimize } \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_t \sum_{j \in J} OC_j(t) \quad (2.3)$$

$$\text{subject to } \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_t \sum_{j \in J} PC_j(t) < \delta. \quad (2.4)$$

In this chapter, we handle VM placement schemes to solve the optimization problem. First, we present schemes for computing-aware initial VM placement. The computing-aware initial VM placement is operated to select the appropriate compute nodes to build VMs for executing applications. The placement algorithm is based on execution time prediction of target applications, and the prediction is achieved using methods for application and computing resource profiling and similarity analysis (Sect. 2.2). Second, we present schemes for VM reallocation based on resource utilization-aware two interactive actions: VM consolidation and dispersion. In the schemes, resource utilization of each compute node is carefully predicted by self-adjusting workload prediction. Based on the prediction, the actions are operated to balance VM consolidation and dispersion (Sect. 2.3).

## 2.2 Computing-Aware Initial VM Placement

### 2.2.1 Overview

In this chapter, we handle VM placement schemes to solve the optimization problem. First, we present a computing-aware initial VM placement algorithm. In the algorithm, each computing resource is ranked in order of execution performance of each application using profiled information, and the ranking is used for the selection of the appropriate compute nodes to build VMs for executing the applications. We note that that some part of this section is composed based on [10].

### 2.2.2 Computing-Aware Initial VM Placement Algorithm

Computing-aware initial VM placement is operated in four steps: (1) application profiling, (2) computing resource profiling, (3) similarity analysis, (4) initial VM placement. Given applications and computing resources, application profiling extracts each application's characteristics, and computing resource profiling records execution results of numerous benchmarks. Measuring similarity between each application and the numerous benchmarks, we can estimate execution performance of each application in each computing resource. Finally, initial VM placement is operated based on the estimated performance information.

**Application profiling.** Application profiling is achieved by extract micro-architecture independent characteristics (MICs) of each application. The MICs provide hardware-independent information of applications and distinguishes each application across various microarchitectures [11]. In Ref. [11], the authors seven categories of MICs as follows.

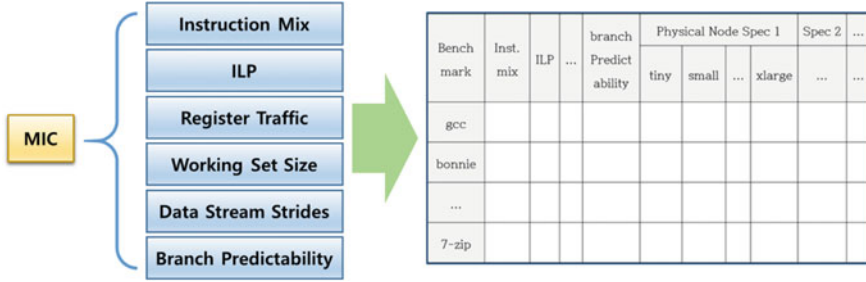


Fig. 2.1 MICs extraction

**itypes:** represent the percentage of load instruction, store instruction, branch arithmetic operation instruction.

**Reg:** characterize the register. Measure the utilization of all operation to write to or read from register and register dependency distance.

**PPM:** measure the accuracy of branch prediction on theoretical prediction-by-partial-matching (PPM).

**ILP:** represent the rate of the operation which is applicable to simultaneous processing inherent in application.

**Memreusedist:** represent the cache behavior of application and measure the re-usage distribution on memory.

**memfootprint:** measure the instruction and working set size of data stream.

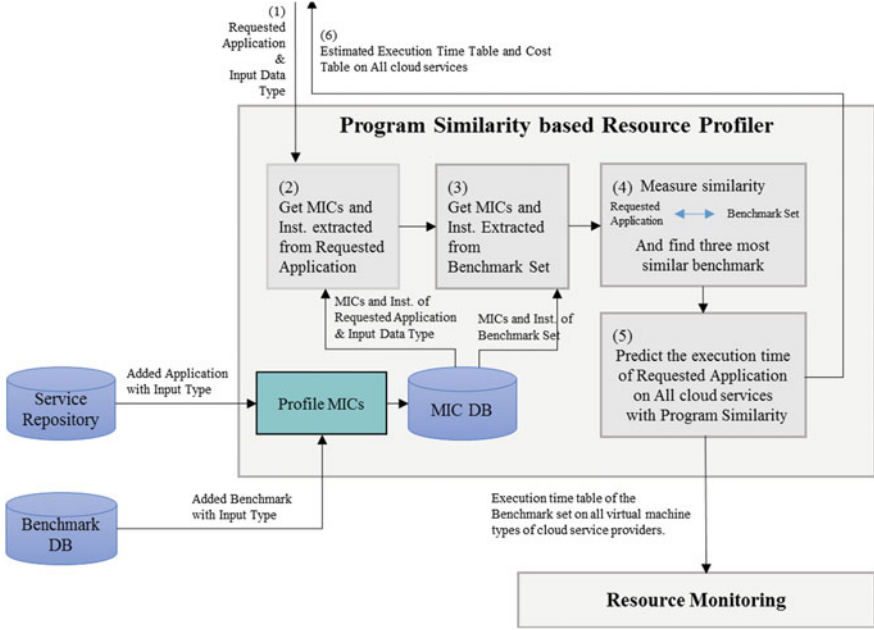
**Stride:** measure the difference of memory access interval when successive memory access occurs (Fig. 2.1).

**Computing resource profiling.** Computing resource profiling is achieved by numerous executions of benchmarks for each compute node. The execution results are normalized by dividing the execution times into the number of instructions (Fig. 2.2).

**Similarity analysis.** Figure 2.3 shows a procedure of program similarity-based execution time prediction. In the figure, execution time of each application is predicted by measuring similarity between the application and the benchmark sets used for computing resource profiling. Similarity  $S_{t,b_i}$  is defined as the reciprocal of Euclidean distance between MIC vectors of the target application  $t$  and benchmark  $b_i$ . The MIC vectors are defined as  $MIC_a = (C_1^a, C_2^a, \dots, C_n^a)$  for application  $a$  where

CSP	Eucalyptus			
exec/instruct	c1.medium(*10 <sup>^</sup> -10)	M1.large(*10 <sup>^</sup> -10)	M1.xlarge(*10 <sup>^</sup> -10)	C1.xlarge(*10 <sup>^</sup> -10)
bzip2-image(2.6MB)	1.42	1.43	1.416	1.41
imagemagick	1.59	1.61	1.59	1.58
gzip-image(2.6MB)	1.93	1.98	1.98	1.98
bzip2-video(453MB)	2.05	2.05	1.64	1.47
Expected exec/instruct	1.724078163	1.743629365	1.641754766	1.599016576

Fig. 2.2 Computing resource profiling example of memcoder [12] in Eucalyptus [13]



**Fig. 2.3** A procedure of program similarity-based execution time prediction

each element is the categories of MICs. After measuring the similarities, the resource profiler selects the three benchmarks  $(b_x, b_y, \dots, b_z)$  which have the highest similarity values, and predicts the execution time of the target application  $t$  in computing resource  $r_j$  using the measured similarity. Equation (2.5) shows the predicted execution time of the target application  $t$  in computing resource  $r_i$  where  $T_{a,r_j}$  is the execution time of application  $a$  in computing resource  $r_i$ ,  $inst_a$  is the number of instructions of application  $a$ . The formula is based on [14]

$$\hat{T}_{t,r_i} = \frac{inst_t}{S_{t,b_x} + S_{t,b_y} + S_{t,b_z}} \cdot \left( \frac{S_{t,b_x} \cdot T_{b_x}}{inst_{b_x}} + \frac{S_{t,b_y} \cdot T_{b_y}}{inst_{b_y}} + \frac{S_{t,b_z} \cdot T_{b_z}}{inst_{b_z}} \right) \quad (2.5)$$

**Computing-aware VM placement.** Algorithm 1 shows computing-aware placement algorithm. In the algorithm, available compute nodes are sorted based on the rank of estimation time prediction results of the target application  $t$  in computing resource  $r_j$  in ascending order. Then, the algorithm selects an available compute node who has the highest rank.

Algorithm 1. Computing-aware VM placement [10]

<b>Input</b>	$t, r_i$ ( $t$ : the target application, $r_i$ : computing resource type)
1:	get the sorted compute node list based on execute time prediction results of the target application $t$ in
	each computing resource $r_j$
2:	for each compute node $p_i$
3:	if compute node $p_i$ is available for computing resource type $r_i$
4:	$createVM = createNewVM(p_i, r_i)$
5:	Return $createVM$
6:	end if
7:	end for

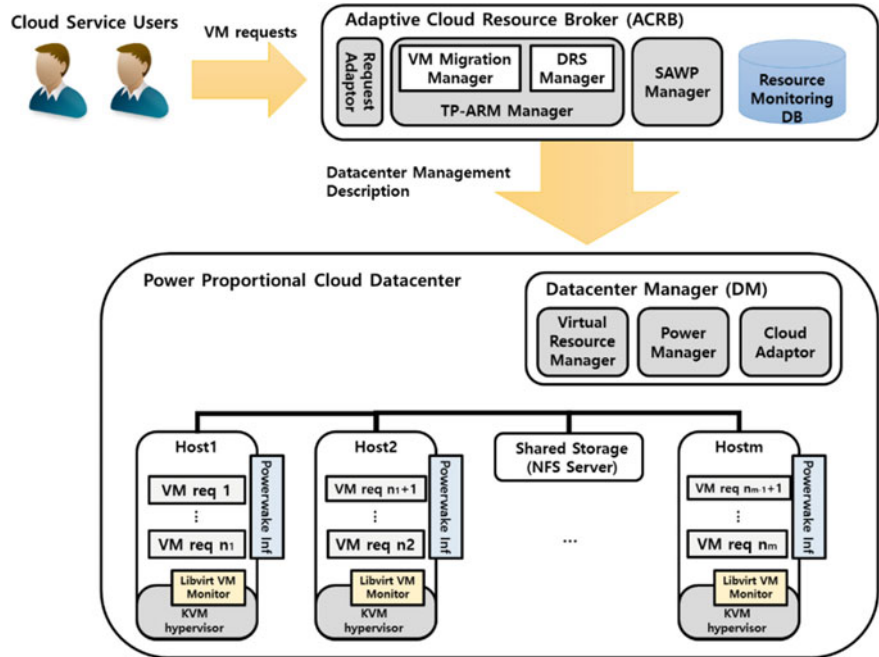
## 2.3 VM Reallocation Based on Resource Utilization-Aware VM Consolidation and Dispersion

### 2.3.1 Overview

In this section, we propose a Two Phase based Adaptive Resource Management (TP-ARM) scheme to address the above challenges for cloud datacenters. The energy consumption model based on TP-ARM was formulated with a performance cost (reputation loss) caused by an increased delay from downsizing active servers, by an energy cost from keeping particular servers active, and by a cost incurred from switching off servers on. Subsequently, we designed an automated cloud resource management system called the Adaptive Cloud Resource Broker (ACRB) system with the TP-ARM scheme. Moreover, we introduced our novel prediction method called Self-Adjusting Workload Prediction (SAWP) to increase the prediction accuracy of users' future demands even under unstatic and irregular workload patterns. The proposed SAWP method adaptively scales the history window size up or down according to the extracted workload's autocorrelations and sample entropies which measure the periodicity and burstiness of the workloads [15]. To investigate the performance characteristics of the proposed approaches, we conducted various experiments to evaluate energy consumption, resource utilization and completion time delay by live migration and DRS execution on a real testbed based on Openstack which is a well-known cloud platform using KVM hypervisor [16]. Through meaningful experimental results, we found that our proposed TP-ARM scheme and SAWP method provide significant energy savings while guaranteeing acceptable performance required by users in practice. We note that that some part of this section is composed based on [17].

### 2.3.2 System Architecture

In this subsection, we discuss the architecture of the automated ACRB system. Our considered cloud environment including ACRB, which supports deploying a



**Fig. 2.4** Adaptive Cloud Resource Brokering (ACRB) system including TP-ARM module and SAWP manager for green cloud datacenters [17]

resource management scheme in order to migrate VM requests and adjusts the number of active servers according to the workload level, is depicted in Fig. 2.4. There are  $m$  physical hosts and  $n$  VM requests in the cloud datacenter. In the cloud datacenter, the information on resource utilization is collected through KVM hypervisor based monitoring modules into the Resource Monitoring DB module, and reported to the ACRB which is responsible for solving the migration of allocated VM requests and sizing the datacenter. The ACRB has two modules: the TP-ARM module and the SAWP Manager. The TP-ARM module includes the VM Migration Manager and the DRS Manager. In the first phase, the VM Migration Manager is responsible for selecting the appropriate VM requests to be migrated based on the measured resource utilization level. We describe the metrics for determining the VM request migration in detail in Sect. 2.2. In the second phase, the DRS Manager is responsible for finding the optimal number of active servers in the cloud datacenter. The DRS plan derived by the DRS Manager based on the amount of submitted VM requests and the measured resource utilization from each VM request is delivered to the Datacenter Manager, and the determined percentage of idle servers are powered off. The SAWP Manager is responsible for adjusting the window size of historical data adaptively in order to predict the future demand of VM requests. The SAWP Manager is able to achieve the exact prediction of future demands even under varied workload levels by considering the periodicity and the

fluctuation of the historical data. The owner of the cloud datacenter has to minimize the costs for resource operations while boosting the benefits which can increase based on the good reputation of the observed QoS of the cloud services. In this paper, our ACRB tries to find a resource management solution to minimize the total cost of resource operations including two sub cost models: the energy consumption cost and the performance reputation cost.

From the perspective of the energy consumption cost, the VM Migration Manager tries to maximize the resource utilization of each physical host by consolidating running VM requests based on the whole offered load measured through the Libvirt VM Monitor module attached to each host. The Libvirt VM Monitor module gauges the utilization of resources such as CPU, memory, and I/O bandwidth in order to check whether whole hosts are overloaded [18, 19]. VM requests on overloaded hosts are preferably migrated to other hosts which have enough extra resource capacity to accommodate newly allocated VM requests. The DRS Manager sends the shutdown messages to servers which have to be powered off, while it sends magic packets to the ones which have to be powered on again. Obviously, the transition of servers from sleeping mode to active mode (aWake transition) requires additional energy consumption (note that the transition overhead from active to sleep (aSleep transition) can be negligible because it requires only a short time to be carried out compared to the aWake transition). Therefore, it is clear that frequent aWake transitions have to be discouraged in order to minimize unnecessary energy consumption. To simplify our model, we assume that the aSleep transition causes no additional energy consumption (in fact, it requires non-zero energy consumption). In terms of the performance reputation cost, the VM Migration Manager tries to decentralize running VM requests over multiple physical hosts in order to avoid QoS deterioration caused by VM interference. In cloud datacenters, the VM co-location interference is the key factor that makes servers undergo severe performance degradation [20, 21]. VM co-location interference is caused by resource contention which is reflected mainly by the number of co-located VM instances and the resource utilization of them. Briefly, VM co-location interference becomes larger as more VM instances are co-located on a common server, and subsequently, higher resource utilization occurs. Therefore, VM requests have to be scattered in order to avoid performance degradation by VM co-location interference as best as possible. Because of the complexity of the optimization for resource management in large-scale cloud datacenters, our TP-ARM scheme adopts a metaheuristic based on GA to obtain a near optimal solution for VM migration to achieve energy savings and QoS assurance in a cloud datacenter.

### 2.3.3 Cost Optimization Model of TP-ARM in Clouds

In this subsection, we introduce an energy cost model and performance reputation cost model based on our proposed TP-ARM scheme including VM live migration and DRS execution in the ACRM.



**Workload cost model for phase 1: VM migration.**

In general, VM requests have heterogeneous workloads in cloud datacenters. There are three types of VM request workloads: CPU, block I/O, and network I/O intensive workloads. The CPU intensive workloads such as scientific applications, high-definition video compression or big data analysis should be processed within an acceptable completion time determined by the cloud service users. The block I/O intensive workloads such as huge engineering simulations for critical areas include astrophysics, climate, and high energy physics which are highly data intensive [22]. These applications contain a large number of I/O accesses where large amounts of data are stored to and retrieved from disks. Network I/O intensive workloads such as Internet web services and multimedia streaming services have to be processed within a desirable response time according to their application type [18, 23]. Each VM request requires different resource utilizations according to their running applications. We categorized the resource utilization of running VM requests on a physical host in a cloud datacenter into two parts in this paper, the Flavor Utilization (FU) and the Virtual Utilization (VU). FU represents the ratio of the resource flavor (i.e., the specification of the resource requirement) of a VM request to the resource capacity of the physical host. VU represents the resource utilization of an assigned virtual resource. The Flavor Utilization  $FU_{j,i}^k$  of a VM request  $i$  on a physical host  $j$  in the resource component  $k$  and the Virtual Utilization  $VU_{j,i}^k$  of the VM request  $i$  on the physical host  $j$  in the resource component  $k$  are given by

$$FU_{j,i}^k = \frac{flv_i^k}{rcp_j^k} \quad (2.6)$$

$$RU_{j,i}^k = FU_{j,i}^k \cdot VU_{j,i}^k \quad (2.7)$$

where  $flv_i^k$  is the flavor of the VM request  $i$  in the resource component  $k$ ;  $rcp_j^k$  is the resource capacity of the physical host  $j$  in the resource component  $k$ , and  $RU_{j,i}^k$  is the actual resource utilization of the VM request  $i$  on the physical host  $j$  in the resource component  $k$ ; therefore, it describes the practical workload in the resource component  $k$ . Note that the FU of the VM request can be determined through its attached service description to the ACRM in advance, while the VU can only be measured by the internal monitoring module in the cloud server during the running duration of the VM request. In the period  $t$ , the VM migration plan is designed according to the FU of the submitted VM requests at period  $t$  and the measured RU of each host in the past history  $t - 1$ . The VM migration plan should satisfy the following constraints:

$$flv_i^k \leq S_k^t(i), \quad \forall i, k, t \quad (2.8)$$

$$flv_i^k \geq 0, \quad \forall i, k \quad (2.9)$$

$$S_{state}^t(i) = 1, \quad \forall i, t \quad (2.10)$$

where  $S_k^t(i)$  denotes the remain capacity of resource  $k$  in the physical server  $S$  which is assigned to the VM request  $i$  at time period  $t$ , and  $S_{state}^t(i)$  represents the state of the physical server  $S$  which is assigned to the VM request  $i$  at time period  $t$ . If the server  $S$  is in the sleep mode, then  $S_{state}^t(i) = 0$ , otherwise,  $S_{state}^t(i) = 1$  (i.e., it is in the active mode). Equation (2.8) represents that the total requirements of the resource capacity of the newly allocated VM requests and migrated VM requests at time period  $t$  cannot exceed the resource capacity provided by their assigned physical server  $S$ . Next, we consider a VM performance reputation which is determined based on the RU of each physical server. The VM co-location interference implies that the virtualization of the cloud supports resource isolation explicitly when multiple VM requests are running simultaneously on a common PM however, it does not mean the assurance of performance isolation between VM requests internally. There is a strong relationship between VM co-location interference and both of the number of co-located VMs and their resource utilization in the PM. As the number of co-located VM requests increase and the RU of the VM requests becomes larger, VM co-location interference becomes more severe. The VM Migration Manager selects server candidates to be migrated based on the amount of RU for each physical server. To achieve this, the RU size of each server is measured through an internal monitoring module, and the VM Migration Manager checks whether they exceed the predetermined RU threshold value  $RU_{thr}$ . In servers which have an RU size over the  $RU_{thr}$ , they are considered as candidates for migration servers during the next period. After the migration servers are chosen, the VM requests to be migrated and their destination servers are determined based on the performance reputation cost model. We propose the objective function of the performance reputation cost  $C_{repu}^t$  at time  $C_{repu}^t$  given by

$$C_{repu}^t = \rho_{inf} \sum_{\forall j} \sum_{\forall k} \omega_k |\mathcal{PM}_j^t| \cdot \left( \frac{\sum_{i=1}^{|\mathcal{PM}_j^t|} RU_{idx_{ji}^k}}{RU_{thr}^k} - 1 \right)^+ + \rho_{mig} T_{mig} \sum_{\forall j} \left| |\mathcal{PM}_j^{t-1}| - |\mathcal{PM}_j^t| \right| \quad (2.11)$$

where  $(x)^+ = \max(x, 0)$ , and  $\rho_{inf}$  and  $\rho_{mig}$  are the price constants of the VM interference cost and VM migration cost, respectively. We use  $T_{mig}$  to denote the processing time for the VM migration. The first term in the right hand of Eq. (2.11) represents the following: as the number of concurrent running VM requests on a physical server increases, the number of users experiencing undesirable performance degradation also increases. The second term represents the following: as the number

of migrated VM requests increases, the migration overhead is also increases. Therefore, a migration plan needs to be found that satisfies both avoiding unnecessary migration overhead and minimizing performance reputation degradation.

### Energy Consumption Cost Model for Phase 2: DRS Procedure.

To achieve a power-proportional cloud datacenter which consumes power only in proportion to the workload, we considered a DRS procedure which adjusts the number of active servers by turning them on or off dynamically [6]. Obviously, there is no need to turn all the servers in a cloud datacenter on when the total workload is low. In the DRS procedure, the state of servers which have no running applications can be transitioned to the power saving mode (e.g., sleep or hibernation) in order to avoid wasting energy. At each time  $t$ , the number of active servers in the cloud datacenter is determined by a DRS procedure plan according to the workload of the allocated VM requests. In order to successfully deploy the DRS procedure onto our system, we considered the switching overhead for adjusting the number of active servers (i.e., for turning servers in sleep mode on again). The switching overhead includes the following: (1) additional energy consumption from the transition from a sleep to active state (i.e., awaken transition); (2) wear-and-tear cost of the server; (3) fault occurrence by turning servers in sleep mode [6]. We considered the energy consumption as the overhead from DRS execution. Therefore, we define the constant  $P_{aWake}$  to denote the amount of energy consumption for the aWake transition of the servers. Then, the total energy consumption  $C_{energy}^t$  of the cloud datacenter at time  $t$  is given by

$$C_{energy}^t = \rho_p P_{active} \sum_{\forall j} \left( |\mathcal{PM}_j^t| \right)^- + \rho_p P_{aWake} T_{switch} \left( \sum_{\forall j} \left( |\mathcal{PM}_j^t| \right)^- - \sum_{\forall j} \left( |\mathcal{PM}_j^t| \right)^- \right)^+ \quad (2.12)$$

where  $(x)^- = \min(x, 1)$ . We use  $\rho_p$  to denote the constant of the power consumption price, and  $P_{active}$  and  $P_{aWake}$  are the amount of power consumption for the active mode and the aWake transition of the server, respectively.  $T_{switch}$  is the time requirement for the aWake transition. The first term on the right side of Eq. (2.12) represents the energy consumption for using servers to serve VM requests allocated to all the physical servers in the cloud datacenter at time  $t$ , and the second term represents the energy consumption for the awaken transition of sleeping servers. Especially, the second term implies that a frequent change in the number of active servers could increase an undesirable waste of energy. Note that the overhead by the transition from the active to the sleep state (i.e., asleep transition) is ignored in our model because the time required for the asleep transition is relatively short compared to the one for the awaken transition.

### Algorithm 2. Phase 1: VM migration in TP-ARM

---

Input : the VM requests allocation of each server  $\mathcal{PM}^{t-1}, \forall j$   
Output : the VM migration plan  $\mathcal{PM}^t, \forall j$

---

```

00: for each  $\mathcal{PM}_j^{t-1} \in \mathcal{PM}^{t-1}, \forall j$ 
01:   for each  $VM_{idx_{j,i}^{t-1}}, \forall i$ 
02:     for each resource component  $k$ 
03:       measure resource utilization  $RU_{idx_{j,i}^{t-1}}^k$  of the VM request with  $idx_{j,i}^{t-1}$ 
04:     end for
05:   end for
06: end for
07: calculate the average resource utilization  $\overline{RU}^k$  at all resource components  $k$ 
08: if  $\overline{RU}^k < RU_{thr^{low}}^k, \forall k$  then
09:   derive the VM migration plan  $\mathcal{PM}^t$  based on Eq.(8) through Algorithm 3.
10: else if  $\overline{RU}^k > RU_{thr^{high}}^k$  then
11:   go to Algorithm 2.
12: end if

```

---

The purpose of our algorithm is to achieve a desirable VM migration and DRS procedure plan that are energy efficient as well as a QoS aware resource management at each period iteratively. At period  $t - 1$ , the proposed TP-ARM approach aims to minimize the cost function in Eq. (2.8) by finding the solution  $\mathcal{PM}^t$  as follows,

$$\underset{\mathcal{PM}^t}{\text{minimize}} \quad C_{total}^t = C_{repu}^t + C_{energy}^t \quad (2.13)$$

Subject to Eqs. (2.8–2.10)

To solve the objective cost function in Eq. (2.13), we prefer a well-known evolutionary metaheuristic called the Genetic Algorithm in order to approximate the optimal plan  $\mathcal{PM}^t$  at each period because Eqs. (2.11) and (2.12) have non-linear characteristics. In next section, our proposed TP-ARM with the VM migration and DRS procedure is introduced in detail.

### 2.3.4 Heuristic Algorithms for the Proposed TP-ARM Scheme

In this subsection, we describe heuristic algorithms for the proposed TP-ARM scheme discussed in Algorithms 2, 3, and 4. First, the process of VM migration in the TP-ARM approach includes the following four steps: (1) monitor and collect the resource utilization data on VM requests for each physical server through the attached Libvirt based monitoring tools; (2) go step 3 if the average utilization of all the active servers are significantly low (i.e., below the predefined threshold), otherwise go to step 4; (3) choose active servers which are supposed to be turned

off, migrate all the VM instances on them to other servers and trigger DRS execution; (4) determine the number of sleeping servers which are supposed to be turned on and send magic packets to them for wake-up if the average utilization of all the active servers are significantly high (i.e., above the predefined threshold), otherwise maintain the current number of active servers. Algorithm 2 shows the procedure for Phase 1: VM migration in the TP-ARM scheme. From line 00 to 06, the resource utilization RU of all the VM requests on the physical servers in the cloud datacenter is measured by the Libvirt API monitoring module. The average resource utilization  $\overline{RU}^k$  at all resource components k is calculated in line 07. If the  $\overline{RU}^k$  is below the predetermined  $\overline{RU}^k_{thr^{low}}$  at all the resource components, then the optimal VM migration plan  $\mathcal{PM}^t$  is derived, and the algorithm is finished. Otherwise, if the  $\overline{RU}^k$  exceeds the predetermined  $\overline{RU}^k_{thr^{high}}$ , then Algorithm 2 for the DRS procedure is triggered.

### Algorithm 3. Phase 2: DRS procedure in TP-ARM

---

Input : the average resource utilization  $\overline{RU}^k, \forall k$   
the VM requests allocation of each server  $\mathcal{PM}^{t-1}, \forall j$   
the historical data of resource utilization  $\overline{RU}^k$  in  $t-1, t-2, \dots, t-\varpi$

---

Output : the set of powered-off servers to be turned on at time  $t$ .

---

00: determine the boundary size of the history window,  $\varpi' (\leq \varpi)$  through Algorithm 4.  
01: estimate the sign  $\varphi$  and angle  $\xi$  of an slope of the historical resource utilization curve from  $t-1$  to  $t-\varpi'$ .  
02: **if**  $\varphi \geq 0$  **then**  
03:      $\alpha = \alpha_{large} \cdot \xi$   
04: **end if**  
05: **if**  $\varphi < 0$  **then**  
06:      $\alpha = \alpha_{small} \cdot \xi$   
07: **end if**  
08: determine the number of sleeping servers supposed to be turned on according  
to  $\alpha \cdot \max_k \frac{\overline{RU}^k}{\overline{RU}^k_{thr^{high}}}$   
09: derive the set of powered-off servers to be turned on at time  $t$

---

Algorithm 3 shows the procedure for Phase 2: DRS in the TP-ARM scheme. In line 00, the history window size  $\varpi'$  is determined by the ACRM through Algorithm 4: the SAWP scheme that is described in next section. In line 02, we calculate the sign  $\varphi$  and the angle  $\xi$  of the slope of the historical resource utilization curve from the current time  $t-1$  to the time  $t-\varpi'$ . If  $\varphi \geq 0$  (i.e., the resource utilization is increased), and then, we get the coefficient  $\alpha$  based on  $\alpha_{large} \cdot \xi$  ( $\alpha_{large} > \alpha_{small}$ ) to adaptively react to the large workload level in the cloud datacenter. Otherwise, if  $\varphi < 0$  (i.e., the resource utilization is decreased), then we get  $\alpha$  based on  $\alpha_{small} \cdot \xi$  to maximize the energy saving of the cloud datacenter. In line 08, we determine the number of servers to be in active mode in the next period.

Algorithm 4 describes the GA for the TP-ARM scheme in detail.  $pop^h$  is a population with size P (even number) at the  $h$ th generation.  $h_{max}$  is the maximum of

the GA iteration count. In line 03,  $f(\cdot)$  is a fitness function of the total cost with two parameters, candidate solution  $\mathcal{PM}^{h,i}$  and the previous solution  $\mathcal{PM}^{t-1}$  at time  $t - 1$ . From line 04 to 06, the two candidate solutions are iteratively chosen randomly from pop to generate offsprings by crossover until there are no remaining unselected solutions in pop. From line 07 to 08, the fitness function values of each offspring are calculated similar to line 03. In line 09, all the parent solutions in pop and generated offsprings are sorted in ascending order for their corresponding fitness function values. In line 10, the next population including only P solutions that achieve good performance from the union of the original pop and derived offsprings is generated to improve the quality of the final solution. From line 11 to 12, to reduce the time complexity of the GA procedures, when we encounter the first solution that has a fitness function value below the predetermined fitness threshold value  $fv_{thr}$ , it counts as a final solution for the next period, and the algorithm is finished. In line 14, if we cannot find a solution that satisfies  $fv_{thr}$  when the iteration count reaches  $h_{max}$ , then we select a solution that has a minimum fitness function value in the population which satisfies the conditions in Eqs. (2.8–2.10) as a final solution for the next period. If there are no solutions to satisfy all the constraints, then we just preserve the current resource allocation vector shown from line 15 to 16. In the population of a GA, mutations are often applied in order to include new characteristics from the offsprings that are not inherited traits from the parents [24]. We did not consider mutations in our GA in this paper; however, it can be used to improve the quality of the GA for the TP-DRM in future work.

#### Algorithm 4. GA for searching VM migration plan in TP-ARM

---

Input : the set of VM requests allocation of whole servers  $\mathcal{PM}^{t-1}$  at time  $t - 1$   
Output : the set of VM requests migration plan,  $\mathcal{PM}^t$  at time  $t$

---

```

00: initialize  $\mathbf{pop}^h = \{\mathcal{PM}^{h,1}, \mathcal{PM}^{h,2}, \dots, \mathcal{PM}^{h,P}\}$ ,  $P = \text{size of population}$ ,
    and even number
01: while  $h \leq h_{max}$ 
02:   for each  $\mathcal{PM}^{h,i} \in \mathbf{pop}^h$ 
03:      $fv^{h,i} = f_{c_{total}}(\mathcal{PM}^{h,i}, \mathcal{PM}^{t-1})$ 
04:   while  $(\forall \mathcal{PM}^{h,i} \in \mathbf{pop}^h \text{ are selected as parents})$ 
05:      $(\mathcal{PM}^{h,i}, \mathcal{PM}^{h,j}) \leftarrow \text{select parents randomly from } \mathbf{pop}^h$ 
06:      $\mathbf{offspring}^h = \cup \text{crossover}(\mathcal{PM}^{h,i}, \mathcal{PM}^{h,j})$ 
07:   for each  $\mathbf{off\_PM}^{h,i} \in \mathbf{offspring}^h$ 
08:      $\mathbf{off\_fv}^{h,i} = f_{c_{total}}(\mathbf{off\_PM}^{h,i}, \mathcal{PM}^{t-1})$ 
09:   sort  $\mathbf{pop}^{h\sim} = \{\mathcal{PM}^{h,1}, \dots, \mathcal{PM}^{h,P}, \mathbf{off\_PM}^{h,1}, \dots, \mathbf{off\_PM}^{h,\frac{P}{2}}\}$  in
    ascending order of solutions' corresponding  $fv^{h,i}$  and  $\mathbf{off\_fv}^{h,i}$ 
10:    $\mathbf{pop}^{h+1} = \{\mathbf{pop}_1^{h\sim}, \dots, \mathbf{pop}_{\frac{P}{2}}^{h\sim}\}$ 
11:   if  $f_{c_{total}}(\mathbf{pop}_1^{h+1}, \mathcal{PM}^{t-1}) \leq fv_{thr}$  then
12:      $\mathcal{PM}^t = \mathbf{pop}_1^{h+1}$  and exit
13:    $h++$ 
14:    $\mathcal{PM}^t = \text{argmin}_{\mathbf{pop}_i^{h_{max}} \in \mathbf{pop}^{h_{max}}} (f_{c_{total}}(\mathbf{pop}_i^{h_{max}}, \mathcal{PM}^{t-1}))$  s.t. Eq (3), (4), (5)
15: if  $\mathcal{PM}^t = \emptyset$  then
16:    $\mathcal{PM}^t = \mathcal{PM}^{t-1}$ 

```

---

**Self-Adjusting Workload Prediction Scheme.** Algorithm 5 describes the procedure of the proposed SAWP algorithm in ACRB. The irregularity function  $g(\varepsilon, \delta)$  represents a level of unpredictability for future resource utilization where  $\varepsilon$  is its fluctuation value (i.e., levels of instability and aperiodicity), and  $\delta$  is its burstiness value (i.e., levels of sudden surge and decline), and both values are calculated based on [15]. According to the predetermined threshold values  $g_{thr}^{high}$  and  $g_{thr}^{low}$  with  $g(\varepsilon, \delta)$ , the history window size  $\varpi'$  is adaptively updated at each prediction process. A relatively long history window size is not suitable to react to recent changes in the workload but is tolerant to varied workload patterns over a short time while a short history window size is favorable to efficiently respond to the latest workload patterns but is not good for widely varying workloads. Consequently, the SAWP algorithm generally outperforms traditional prediction schemes during drastic utilization changes from various cloud applications because it is able to cope with temporary resource utilizations (i.e., does not reflect overall trends) by adjusting the history window size  $\varpi'$ .

#### Algorithm 5. Self-Adjusting Workload Prediction scheme

---

Input : the historical data of resource utilization  $\overline{RU}^k$  in  $t-1, t-2, \dots, t-\varpi$   
Output : the boundary size of history window  $\varpi'$

---

```

00:  $\varpi' = \varpi$ 
01: analyze the historical data of resource utilization  $\overline{RU}^k$  from  $t-1$  to  $t-\varpi$ 
02: extract the fluctuation size  $\varepsilon$  and the burstness value  $\delta$  based on [6]
03: if  $g(\varepsilon, \delta) > g_{thr}^{high}$  then
04:   increase  $\varpi'$  based on  $\frac{g(\varepsilon, \delta)}{g_{thr}^{high}}$ 
05: end if
06: if  $g(\varepsilon, \delta) < g_{thr}^{low}$  then
07:   decrease  $\varpi'$  based on  $\frac{g_{thr}^{low}}{g(\varepsilon, \delta)}$ 
08: end if
09: predict the future demand based on  $\overline{RU}^k$  from  $t-1$  to  $t-\varpi'$ 

```

---

### 2.3.5 Evaluation

In our experiments, we measured various metrics which affect the parameter decision for our proposed algorithms. To this end, we established five cluster servers as a cloud platform, one server for ACRB with a Mysql DB system, a power measuring device from Yocto-Watt [25] and a laptop machine called the VirtualHub for collecting and reporting information on the measured power consumption shown in Fig. 2.5. The hardware specification of each server for the cloud compute host is as follows: an Intel i7-3770 (8-cores, 3.4 Ghz), 16 GB RAM memory, and two 1 Gbps NICs (Network Interface Cards). In order to measure efficiently the power consumption of the cloud cluster server, we used a power measuring device model called YWATTKM1 by Yocto-Watt. This model has a measurement unit of 0.2 W for AC power with an error ratio of 3% and 0.002 W for the DC power with an error ratio of 1.5%. The VirtualHub collects information

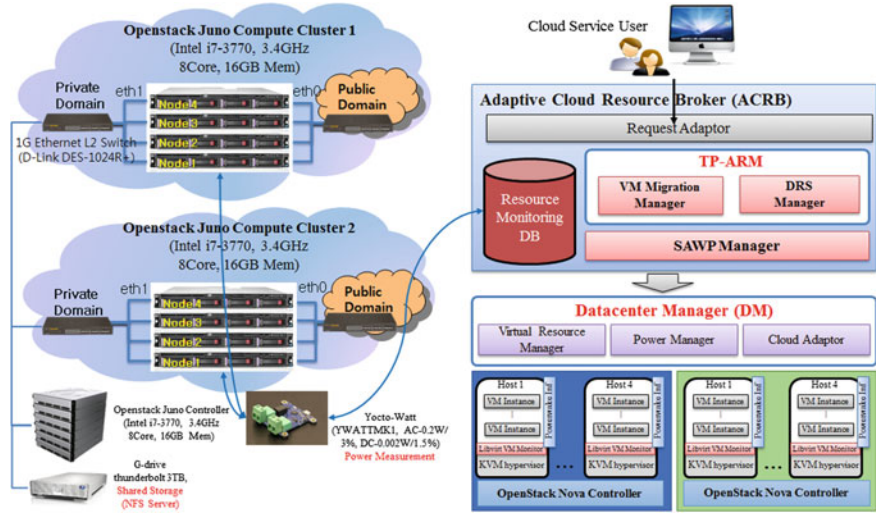


Fig. 2.5 Experimental environment

on power consumption from YWATTMK1 through the Yocto-Watt Java API and reports it to the power monitoring table of the Mysql DB system in the ACRB periodically.

The dynamic resource utilizations by each VM instance are measured via our developed VM monitoring modules based on the Libvirt API and are sent to the resource monitoring table of the Mysql DB system periodically. In addition, a SATA 3 TB hard disk called the G-drive was deployed as a NFS server in our testbed for live migration [26]. We adopted Openstack kilo version which is a well-known open source solution based on KVM Hypervisor as a cloud platform in our testbed. Finally, we used PowerWake package [27] to turn remotely off servers on via Wake on Lan (WOL) technology for DRS execution.

In Table 2.1, we show the average power consumption and resource utilization of two running applications: Montage m106-1.7 projection and ftp transfer. Montage project is an open source based scientific application, and it has been invoked by NASA/IPAC Infrared Science Archive as a toolkit for assembling Flexible Image Transport System (FITS) images into custom mosaics [10, 28]. The m106-1.7 projection in Montage is a CPU-intensive application while the ftp transfer is a network-intensive one. Therefore, the running of m106-1.7 causes a power consumption of about 75 Wh and a CPU utilization of about 15% whereas the power consumption by ftp transfer for a 1.5 GB test.avi file is about 60 Wh, and the network bandwidth usage is about 3.7 Mbps. That is, the CPU usage is the main part that affects the power consumption of server. In terms of DRS execution, the power consumption by an off server is about 2.5 Wh (note that this value is not zero because the NIC and some its peripheral components are still powered on to maintain the standby mode to receive the magic packets from the Powerwake




**Table 2.1** The average power consumption and resource utilization of two running applications: Montage m106-1.7 projection and ftp transfer

Host state		Power consumption (Wh)	CPU utilization	Mem utilization	Net bandwidth
Idle		55 Wh	1.5%	3%	20 Kbps (bytes)
Active	Montage m 106-1.7 (projection)	75 Wh	15%	3.7%	20 Kbps (bytes)
	test.avi downloading (ftp, 1.5 GB)	60 Wh	2%	3.7%	3.7 Mbps (bytes)
aSleep (to Power Off)		70–80 Wh (5–7 s)	1.8% (5–7 s)	3% (5–7 s)	20 Kbps (bytes)
Power off (hibernating)		2.5 Wh			
aWake (from power off)		78 Wh (50 s–1 min)			

controller) while the asleep and awaken transition procedures, which cause the switching overhead for DRS, require power consumption of about 80 Wh to turn the active servers off or to turn off servers on, respectively. The asleep transition procedure is trivial because it requires a short time (i.e., 5–7 s) to complete even though its power consumption is considerable whereas the awaken transition procedure requires a relatively long execution time (i.e., more than 1 min) and should be considered carefully. The overhead for the awaken transition would be a more serious problem in practice because its required execution time is generally far longer (i.e., above 10 min) for multiple servers of racks in a datacenter. Therefore, it is essential to consider the switching overhead for awaken transition to reduce efficiently the resource usage cost of the datacenter.

Figure 2.6 shows the additional energy consumption overhead of a single physical server from the DRS execution with respect to the asleep transition (i.e., from the active mode to the sleep mode) and the awaken transition (i.e., from the sleep mode to the active mode). Note that the asleep transition requires a relatively short processing time of about 7 or 8 s and causes an additional energy consumption of about 20% compared to the idle state of the server, while the awaken transition needs 60 or 90 s to be carried out. Although not included in this paper, we also verified that some of the HP physical servers in our laboratory required more than 10 min to be turned on. We expect that physical servers in a real cloud datacenter require tens of minutes or hundreds of minutes to get back from the sleep mode to the active mode. Moreover, the aWake transition causes an additional energy consumption of about 40% on average compared to the idle state of a server. These results imply that the frequent DRS execution might cause the degradation of the energy saving performance in a cloud datacenter. Our proposed TP-ARM scheme is able to avoid the unnecessary energy consumption overhead by the




**Fig. 2.6** Power consumption overhead from DRS execution of the test server with respect to the aSleep and aWake transition

awaken transition through the cost model by considering the DRS transition overhead shown in Eq. (2.12).


Figure 2.7 shows the resource utilization and the power consumption measured by the Libvirt API monitoring module and Yocto-Watt power measuring device. VM request instance-10 runs m101-1.0 mProj, instance-0f runs m108-1.7, and instance-0c runs the streaming server with the 180 MB movie file. Figure 2.7a–c shows that their resource utilization is consistent with the resource intensive characteristic of their running workload. Figure 2.7d shows the different energy consumptions according to each workload type. The energy consumption of the physical server is 60 Wh in the idle state; it is about 82 Wh when running m101-1.0 and 95 Wh when running both m101-1.0 and m108-1.7, while the streaming server just causes an additional energy consumption of about 2 Wh. As mentioned in Sect. 2.3.3, the main part of the resource components affecting the power consumption in a physical server is the CPU resource, and the effects of other components such as the memory, storage, and network interface cards are negligible in general. These results are consistent with the ones in Table 2.1. Therefore, these results imply that the energy consumption can be different according to which resource component has high utilization. Our proposed workload cost model of the TP-ARM scheme considers different weight values for each resource component in order to derive the practical cost of the resource management.

Figure 2.8 shows the performance of the CPU utilization for the VM live migration between the source machine (kdkCluster2) and the destination machine (kdkCluster1). There are VM request instances, such as running instance-33 and 34, which execute the compression of video files 5 and 4 GB in size, respectively.

Instance-33 executed the process of m101-1.6mProj at 12:40:14 local time and started migration to the kdkCluster1 at 12:44:00 local time. The migration of instance-33 was completed by 13:24:30 local time, and its execution of m101-1.6 mProj ended at 13:47:16. Namely, the completion time of m101-1.6 mProj at instance-33 was 67 min in total, and its migration times were over 30 min. If we




CPU utilization of VM instance 0000000c (Streaming server), VM instance 0000000f (Montage m108-1.7) and VM instance 00000010 (Montage m101-1.0)




Network transmission bytes of VM instance 0000000c (Streaming server), VM instance 0000000f (Montage m108-1.7) and VM instance 00000010 (Montage m101-1.0)

**Fig. 2.7** Libvirt API monitoring module shows the results of resource utilization of running VM requests in the kdkCluster1 in CPU utilization (a) network transmission bytes (b)

consider the time to complete m101-1.6 mProj in the case of no VM live migration, it is forecasted to be around 10 min. We can see that there exist quite big overheads in VM live migration.

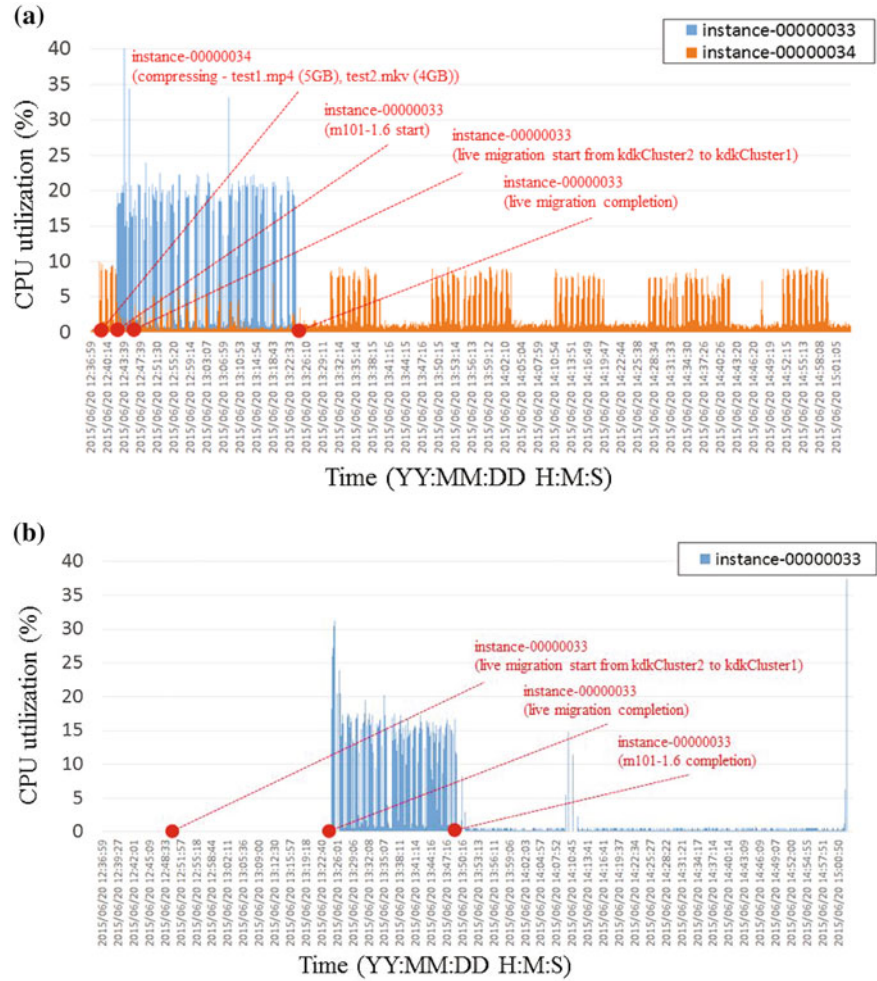


Disk block read bytes of VM instance 0000000c (Streaming server), VM instance 0000000f (Montage m108-1.7) and VM instance 00000010 (Montage m101-1.0)



Power consumption of VM instance 0000000c (Streaming server), VM instance 0000000f (Montage m108-1.7) and VM instance 00000010 (Montage m101-1.0)

**Fig. 2.7** (continued). Libvirt API monitoring module shows the results of resource utilization of running VM requests in the kdkCluster1 in Disk block read bytes (c) and Yocto-Watt module [8] was used to measure the power consumption of the kdkCluster1 as shown in (d)

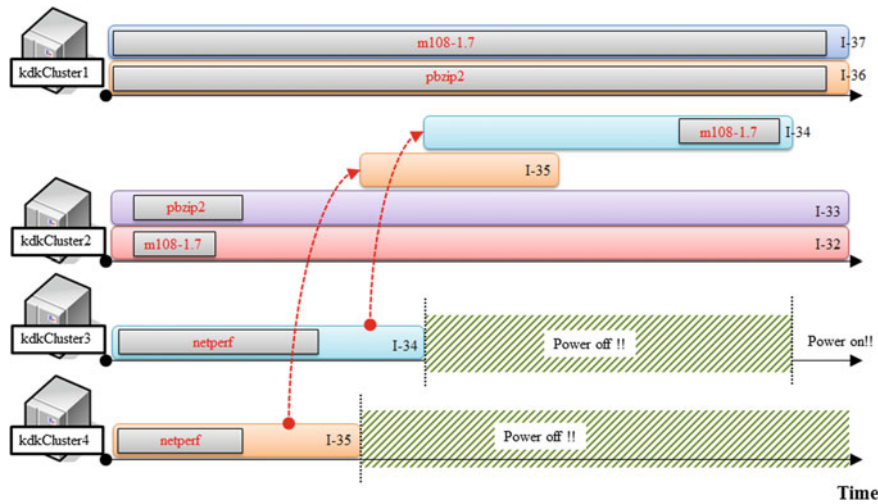


**Fig. 2.8** Results of CPU utilization of the migrated VM request from **a** a source server: kdkCluster2 to **b** a destination server:kdkCluster1

In addition, VM migration shows some problems in power consumption. We considered cost models to identify an efficient migration scheme which was defined in Eq. (2.11). Figure 2.9 shows the test environments and the operation of the VM migration when the proposed TP-ARM schemes are applied to the test cloud systems with heterogeneous applications and test programs including m108-1.7, pbzip2, and netperf. From the experiments, we obtained interesting results in the comparison of the utilization performance, which examined the CPU utilization monitoring results and the measured power consumption, respectively, shown in Fig. 2.10. From the results, we found rapidly changing instants of utilization when the cloud brokering system considered the power consumptions for each running

cloud clusters	Experiment Schedule			
kdkCluster1	• instance-00000037 montage m108-1.7			
	• instance-00000038 pbzip2 video.avi (2GB)			
kdkCluster2	• instance-00000032 montage m108-1.7 (04:58 ~ 05:12)	• instance-00000035 migration ( 4->2 ) (06:24 ~ 06:26)	• instance-00000034 migration ( 3->2 ) (08:13 ~ 08:15)	• instance-00000034 montage (09:20 ~ ) => workload
	• instance-00000033 pbzip2 video.avi (2GB) (05:07 ~ 05:24)			
kdkCluster3	• instance-00000034 netperf (04:58 ~ 05:58)	• instance-00000034 netperf (06:25 ~ 07:13)	• instance-00000034 migration ( 3->2 ) (08:13 ~ 08:15) • power off (08:16 ~ 09:25)	• power on (09:25 ~ )
kdkCluster4	• instance-00000035 netperf (04:58 ~ 05:13)	• instance-00000035 migration ( 4->2 ) (06:24 ~ 06:26) • power off (06:30 ~ 09:30)		


(a) The schedule of the use case




(b) The illustration of the use case

**Fig. 2.9** Test environments and operation of the VM migration when the proposed TP-ARM schemes are applied for test cloud servers (e.g. kdkCluster1-5 in lab test environments) with heterogeneous applications and test programs, such as m108-1.7, pbzip2, and netperf. **a** An example of test schedule under cloud testbed environments. **b** Illustration of the VM migration using two-phase power consumption cost models






CPU utilization of kdkCluster1 with VM instance-37 (m108-1.7) and VM instance-36 (pbzip2)




CPU utilization of kdkCluster2 with VM instance-35 (netperf, kdkCluster4 -> 2), VM instance-34 (m108-1.7, kdkCluster3 -> 2), VM instance-33 (pbzip2) and VM instance-32 (m108-1.7)

**Fig. 2.10** Performance comparison of CPU utilization using 4 test systems in lab environments. **a** and **b** show the results of the measured utilization during the VM live migration in test use case as shown in Fig. 2.9.




CPU utilization of kdkCluster3 with VM instance-34 (m108-1.7, kdkCluster3 -> 2)



CPU utilization of kdkCluster4 with VM instance-35 (netperf, kdkCluster4 -> 2)

**Fig. 2.10** (continued). Performance comparison of CPU utilization using 4 test systems in lab environments. **c** and **d** show the results of the measured utilization during the VM live migration in test use case as shown in Fig. 2.9





**Fig. 2.10** (continued). Performance comparison of CPU utilization using 4 test systems in lab environments. **e** Measured power consumption of test cloud systems (e.g. kdkCluster1–4) in lab cloud test environments

application under cloud data center environments. Additionally, the test set ‘net-perf’, a network intensive workload of a test cloud system (e.g., kdkCluster4 with instance-35) in Fig. 2.10d, showed very low performance in CPU utilization. It was enough to satisfy the threshold value to run the VM migration efficiently. And then, the proposed TP-ARM algorithm adjusts the VM migration procedure to reduce the power consumption sufficiently. Basically, the TP-ARM is triggered to migrate instance-35 to another system kdkCluster2; thereafter, it changes the kdkCluster’s sleeping mode in advance.

Next, because instance-34 of the kdkCluster system was running netperf, which generated a very low performance of utilization, the TP-ARM also did a migration of the instance-34 to the kdkCluster2 and transited to the sleeping mode as well. Therefore, we could expect efficient power consumption through the change in the sleeping mode of kdkCluster3 and kdkCluster4, respectively, as well as keeping active modes for both kdkCluster1 and kdkCluster2. Figure 2.10 shows the performance comparison of the power consumption for the VM migration and DRS process based on the TP-ARM algorithm, e.g., seen in Fig. 2.9. We can see a performance improvement of 60 Wh each in power consumption after the change in the sleeping mode in the case of the VM migration requests. Through the experiments, we verified that the proposed TP-ARM scheme, which carries out adaptive migration and asleep transition through real-time monitoring of resource utilization, has good performance in reducing the power consumption effectively. This study provides good results for achieving an energy efficient cloud service for users by not increasing QoS degradation as much.

### 2.3.6 Conclusion

In this paper, we introduced ACRB with Two Phases based the TP-ARM scheme for energy efficient resource management by real time based VM monitoring in a cloud data center. Our proposed approach is able to reduce efficiently the energy consumption of the servers without a significant performance degradation by live migration and DRS execution through a considerate model that considers switching overheads. The various experimental results based on the Openstack platform suggest that our proposed algorithms can be deployed to prevalent cloud data centers. The novel prediction method called SAWP is proposed in order to improve the accuracy of forecasting future demands even under drastic workload changes.

Especially, we evaluated the performance of our proposed TP-ARM scheme through various applications such as Montage, pbzip2, netperf, and the streaming server which have heterogeneous workload demands. Our TP-ARM scheme could maximize the energy saving performance of the DRS procedure by Phase 1: VM migration achieves consolidated resource allocation under a low workload level. Moreover, it could ensure the QoS of cloud service users by Phase 2: the DRS procedure increases adaptively the number of active servers under a high workload level. Through experiments based on a practical use case, our proposed scheme is not only feasible from a theoretical point of view but also practical in a real cloud environment. In future work, we will demonstrate that our proposed algorithm outperforms existing approaches for energy efficient resource management through various experiments based on the implemented system in practice.

## References

1. Prediction. Available: <http://searchstorage.techtarget.com.au/articles/28102-Predictions-2-9-Symantec-s-Craig-Scroggie>
2. R. Buyya, A. Beloglazov, J. Abawajy, Energy-efficient management of data center resources for cloud computing: a vision, architectural elements, and open challenges (2010)
3. S.A. Baset, Cloud SLAs: present and future. *SIGOPS Oper. Syst. Rev.* **46**(2), 57–66 (2012)
4. J. M. Myerson, *Best Practices to Develop Slas for Cloud Computing*. (IBM Corporation, New York, 2013) p. 9
5. A. Shankar U. Bellur, *Virtual Machine Placement in Computing Clouds CoRR*, abs/1011.5064 (2010)
6. M. Lin, A. Wierman, L.L.H. Andrew, E. Thereska, Dynamic right-sizing for power-proportional data centers. *IEEE/ACM Trans. Netw.* **21**(5), 1378–1391 (2013)
7. Z. Xiao, W. Song, Q. Chen, Dynamic resource allocation using virtual machines for cloud computing environment. *IEEE Trans. Parallel Distrib. Syst.* **24**(6), 1107–1117 (2013)
8. Y. Koh, R. Knauerhase, P. Brett, M. Bowman, Z. Wen, C. Pu, An analysis of performance interference effects in virtual environments. *IEEE Int. Symp. Perform Anal. Syst. Softw.* 200–209 (2007)
9. C.D. Patel, A.J. Shah, Cost model for planning, development and operation of a data center. *Development* **107**, 1–36 (2005)

10. W.-J. Kim, D.-K. Kang, S.-H. Kim, C.-H. Youn, Cost adaptive vm management for scientific workflow application in mobile cloud. *Mob. Netw. Appl.* **20**(3), 328–336 (2015)
11. K. Hoste, A. Phansalkar, L. Eeckhout, A. Georges, L. K. John, K. De Bosschere, *Performance prediction based on inherent program similarity PACT*, vol 9 (Seattle, washinton 2006), p. 114
12. Memcoder. Available: <https://linux.die.net/man/1/mencoder>
13. Eucalyptus. Available: <http://www.eucalyptus.com>
14. K. Hoste, L. Eeckhout, Microarchitecture-independent workload characterization. *IEEE Micro* **27**(3), 63–72 (2007)
15. A. Ali-Eldin, J. Tordsson, E. Elmroth, M. Kihl, *Workload Classification for Efficient Auto-Scaling of Cloud Resources*. (2005)
16. OpenStack. Available: <http://www.openstack.org/>
17. D.-K. Kang, F. Al-Hazemi, S.-H. Kim, M. Chen, L. Peng, C.-H. Youn, Adaptive VM management with two phase power consumption cost models in cloud datacenter. *Mob. Netw. Appl.* **21**(5), 793–805 (2016)
18. M. Chen, Y. Zhang, L. Hu, T. Taleb, Z. Sheng, Cloud-based wireless network: virtualized, reconfigurable, smart wireless network to enable 5G technologies. *Mob. Netw. Appl.* **20**(6), 704–712 (2015)
19. M. Chen, H. Jin, Y. Wen, V. Leung, Enabling technologies for future data center networking: a primer. *IEEE Netw.* **27**(4), 8–15 (2013)
20. F. Xu, F. Liu, L. Liu, H. Jin, B.B. Li, B.B. Li, iAware: making live migration of virtual machines interference-aware in the cloud. *IEEE Trans. Comput.* **63**(12), 3012–3025 (2014)
21. D. Gupta, L. Cherkasova, R. Gardner, A. Vahdat, Enforcing performance isolation across virtual machines in xen, *Proceedings 7th ACM/IFIP/USENIX international conference middleware*, pp. 342–362, (2006)
22. A. Nisar, W.K. Liao, A. Choudhary, Scaling parallel I/O performance through I/O delegate and caching system, *2008 SC—International conference for high performance computing (Storage and Analysis, SC, Networking, 2008)*
23. M. Chen, Y. Zhang, Y. Li, S. Mao, V.C.M. Leung, EMC: Emotion-aware mobile cloud computing in 5G. *IEEE Netw.* **29**(2), 32–38 (2015)
24. K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evol. Comput.* **6**(2), 182–197 (2002)
25. YOCTO-WATT. Available: <http://www.yoctopuce.com/EN/products/usb-electrical-sensors/yocto-watt>
26. G-Technology. Available: <http://www.g-technology.com/products/g-drive>
27. PowerWake. Available: <http://manpages.ubuntu.com/manpages/utopic/man1/powerwake.1.html>
28. Montage. Available: <http://montage.ipac.caltech.edu/>

Cloud Broker and Cloudlet for Workflow Scheduling

Youn, C.-H.; Chen, M.; Dazzi, P.

2017, IX, 212 p. 92 illus., 48 illus. in color., Hardcover

ISBN: 978-981-10-5070-1