OpenStack Neat: a framework for dynamic and energy-efficient consolidation of virtual machines in OpenStack clouds

Anton Beloglazov*,† and Rajkumar Buyya

Cloud Computing and Distributed Systems (CLOUDS) Laboratory, Department of Computing and Information Systems, The University of Melbourne, Melbourne, Australia

SUMMARY

Dynamic consolidation of virtual machines (VMs) is an efficient approach for improving the utilization of physical resources and reducing energy consumption in cloud data centers. Despite the large volume of research published on this topic, there are very few open-source software systems implementing dynamic VM consolidation. In this paper, we propose an architecture and open-source implementation of OpenStack Neat, a framework for dynamic VM consolidation in OpenStack clouds. OpenStack Neat can be configured to use custom VM consolidation algorithms and transparently integrates with existing OpenStack deployments without the necessity of modifying their configuration. In addition, to foster and encourage further research efforts in the area of dynamic VM consolidation, we propose a benchmark suite for evaluating and comparing dynamic VM consolidation algorithms. The proposed benchmark suite comprises Open-Stack Neat as the base software framework, a set of real-world workload traces, performance metrics and evaluation methodology. As an application of the proposed benchmark suite, we conduct an experimental evaluation of OpenStack Neat and several dynamic VM consolidation algorithms on a five-node testbed, which shows significant benefits of dynamic VM consolidation resulting in up to 33% energy savings. Copyright © 2014 John Wiley & Sons, Ltd.

1. INTRODUCTION

Cloud computing has revolutionized the information and communication technology (ICT) industry by enabling on-demand provisioning of elastic computing resources on a pay-as-you-go basis. However, cloud data centers consume huge amounts of electrical energy resulting in high operating costs and carbon dioxide (CO_2) emissions to the environment. It is estimated that energy consumption by data centers worldwide has risen by 56% from 2005 to 2010 and was accounted to be between 1.1% and 1.5% of the global electricity use in 2010 [1]. Furthermore, carbon dioxide emissions of the ICT industry were estimated to be 2% of the global emissions, which is equivalent to the emissions of the aviation industry [2].

To address the problem of high energy use, it is necessary to eliminate inefficiencies and waste in the way electricity is delivered to computing resources, and in the way these resources are utilized to serve application workloads. This can be carried out by improving the physical infrastructure of data centers, as well as resource allocation and management algorithms. Recent advancement in the data center design resulted in a significant increase of infrastructure efficiency. As reported by the Open Compute project, Facebook's Oregon data center achieved a power usage effectiveness

---

*Correspondence to: Anton Beloglazov, CLOUDS Lab, Department of Computing and Information Systems, The University of Melbourne, Melbourne, Australia.
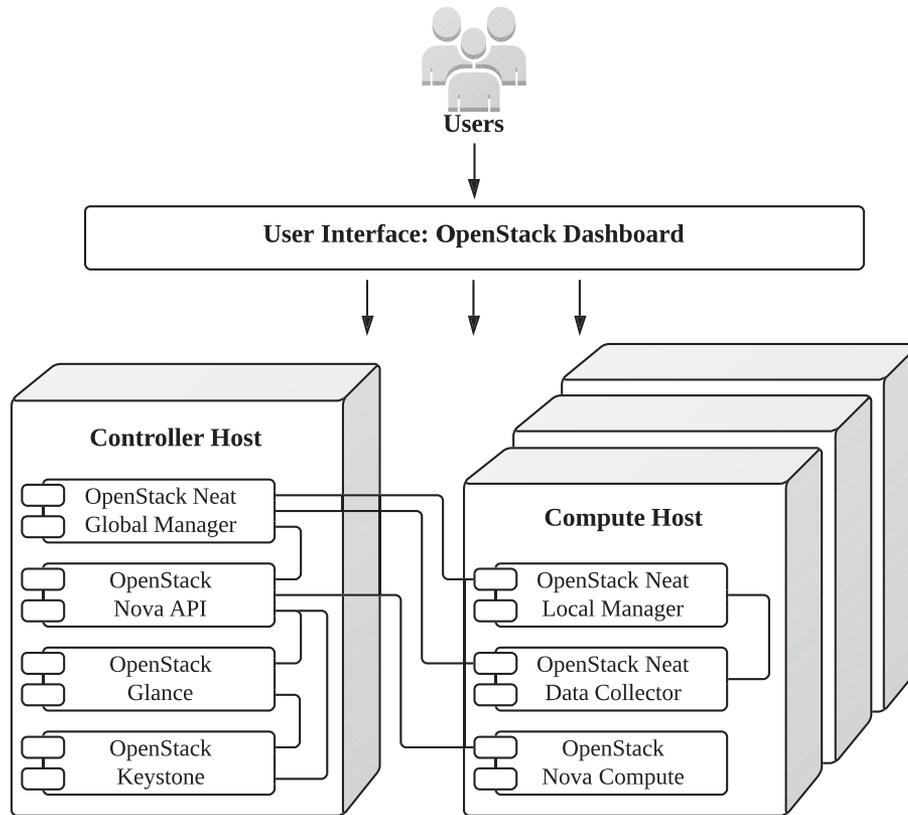†E-mail: a.beloglazov@student.unimelb.edu.au

Figure 1. The combined deployment of OpenStack and OpenStack Neat.

(PUE) of 1.08 [3], which means that approximately 91% of the data center's energy consumption is consumed by the computing resources. Therefore, now it is important to focus on optimizing the way the resources are allocated and utilized to serve application workloads.

One method to improve the utilization of resources and reduce energy consumption is dynamic consolidation of virtual machines (VMs) [4–26] enabled by *live migration*, the capability of transferring a VM between physical servers (referred to as hosts or nodes) with a close-to-zero downtime. Dynamic VM consolidation consists of two basic processes: migrating VMs from underutilized hosts to minimize the number of active hosts and offloading VMs from hosts when those become overloaded to avoid performance degradation experienced by the VMs, which could lead to a violation of the QoS requirements. Idle hosts are automatically switched to a low-power mode to eliminate the static power and reduce the overall energy consumption. When required, hosts are reactivated to accommodate new VMs or VMs being migrated from other hosts. Even though a large volume of research has been published on the topic of dynamic VM consolidation, there are very few open-source software implementations.

In this work, we introduce an architecture and implementation of *OpenStack Neat*[‡]: an open-source software framework for distributed dynamic VM consolidation in cloud data centers based on the OpenStack platform[§]. Figure 1 depicts a typical deployment of the core OpenStack services, OpenStack Neat services, and their interaction, which will be discussed in detail in the following sections. The deployment may include multiple instances of compute and controller hosts. The OpenStack Neat framework is designed and implemented as a transparent add-on to OpenStack, which means that the OpenStack installation need not be modified or specifically configured to benefit from OpenStack Neat. The framework acts independently of the base OpenStack platform

---

[‡]The OpenStack Neat framework, http://openstack-neat.org/.

[§]The OpenStack cloud platform, http://openstack.org/.

and applies VM consolidation processes by invoking public APIs of OpenStack. The purpose of the OpenStack Neat framework is twofold: (1) providing a fully operational open-source software for dynamic VM consolidation that can be applied to existing OpenStack clouds and (2) providing an extensible software framework for conducting research on dynamic VM consolidation.

OpenStack Neat is designed and implemented following the distributed approach to dynamic VM consolidation introduced and evaluated in our previous works [27–29]. The target environment is an infrastructure as a service (IaaS), for example, Amazon EC2, where the provider is unaware of applications and workloads served by the VMs and can only observe them from outside. We refer to this property of IaaS environments as being application agnostic. The proposed approach to distributed dynamic VM consolidation consists in splitting the problem into four sub-problems [29]:

1. Deciding if a host is considered to be underloaded, so that all VMs should be migrated from it, and the host should be switched to a low-power mode.
2. Deciding if a host is considered to be overloaded, so that some VMs should be migrated from it to other active or reactivated hosts to avoid violating the QoS requirements.
3. Selecting VMs to migrate from an overloaded host.
4. Placing VMs selected for migration on other active or reactivated hosts.

This approach has two major advantages compared with traditional fully centralized VM consolidation algorithms: (1) splitting the problem simplifies its analytical treatment by allowing the consideration of the sub-problems independently; and (2) the approach can be implemented in a partially distributed manner by executing the underload/overload detection and VM selection algorithms on compute hosts and the VM placement algorithm on the controller host, which can optionally be replicated. Distributed VM consolidation algorithms enable the natural scaling of the system to thousands of compute nodes, which is essential for large-scale cloud providers. For instance, Rackspace, a well-known IaaS provider, currently manages tens of thousands of servers. Moreover, the number of servers continuously grows: Rackspace has increased the total server count in the second quarter of 2012 to 84,978 up from 82,438 servers at the end of the first quarter [30].

In addition, to facilitate research efforts and future advancements in the area of dynamic VM consolidation, we propose a benchmark suite for evaluating and comparing dynamic VM consolidation algorithms comprising OpenStack Neat as the base software framework, real-world workload traces from PlanetLab, performance metrics, and evaluation methodology (Section 5).

The *key contributions* of this work are the following:

- An architecture of an extensible software framework for dynamic VM consolidation designed to transparently integrate with OpenStack installations and allowing configuration-based substitution of multiple implementations of algorithms for each of the four defined sub-problems of dynamic VM consolidation.
- An open-source software implementation of the framework in Python released under the Apache 2.0 license and publicly available online.
- An implementation of several algorithms for dynamic VM consolidation proposed and evaluated by simulations in our previous works [27–29].
- An initial version of a benchmark suite comprising the software framework, workload traces, performance metrics, and methodology for evaluating and comparing dynamic VM consolidation solutions following the distributed model.
- Experimental evaluation of the framework on a five-node OpenStack deployment using real-world application workload traces collected from more than a thousand PlanetLab VMs hosted on servers located in more than 500 places around the world [31]. According to the estimates of potential energy savings, the evaluated algorithms reduce energy consumption by up to 33% with a limited performance impact.

The remainder of the paper is organized as follows. In the next section, we discuss the related work, followed by the overall design and details of each component of the OpenStack Neat framework in Section 3. In Section 4, we describe the implemented VM consolidation algorithms. In Section 5, we propose a benchmark suite for evaluating distributed dynamic VM consolidation

algorithms. We experimentally evaluate the framework and analyze the results in Section 6. We conclude the paper with a discussion of scalability and future research directions in Section 7 and conclusions in Section 8.

2. RELATED WORK

Research work related to this paper can be divided into two categories: (1) practically implemented and publicly available open-source software systems and (2) theoretical work on various approaches to dynamic VM consolidation. Despite the large volume of research published on the topic of dynamic VM consolidation, there are very few software implementations publicly available online. To the best of our knowledge, the earliest open-source implementation of a VM consolidation manager is the Entropy project.¶ Entropy is an open-source VM consolidation manager for homogeneous clusters developed by Hermenier *et al.* [4] and released under the Lesser General Public License.

Entropy is built on top of Xen and focused on two objectives: (1) maintaining a configuration of the cluster, where all VMs are allocated sufficient resources, and (2) minimizing the number of active hosts. To optimize the VM placement, Entropy periodically applies a two-phase approach. First, a constraint programming problem is solved to find an optimal VM placement, which minimizes the number of active hosts. Then, another optimization problem is solved to find a target cluster configuration with the minimal number of active hosts that also minimizes the total cost of reconfiguration, which is proportional to the cost of VM migrations. Instead of optimizing the VM placement periodically as Entropy, OpenStack Neat detects host underload and overload conditions and dynamically resolves them, which allows the system to have a more fine-grained control over the host states. Although Entropy may find a more optimal VM placement by computing a globally optimal solution, all aspects of VM placement optimization must be computed by a central controller, thus limiting the scalability of the system.

Feller *et al.* [5] proposed and implemented a framework for distributed management of VMs for private clouds called Snooze,‖ which is open source and released under the General Public License version 2. In addition to the functionality provided by the existing cloud management platforms, such as OpenStack, Eucalyptus, and OpenNebula, Snooze implements dynamic VM consolidation as one of its base features. Another difference is that Snooze implements hierarchical distributed resource management. The management hierarchy is composed of three layers: local controllers on each physical node, group managers managing a set of local controllers, and a group leader dynamically selected from the set of group managers and performing global management tasks. The distributed structure enables fault tolerance and self-healing by avoiding single points of failure and automatically selecting a new group leader if the current one fails. Snooze also integrates monitoring of the resource usage by VMs and hosts, which can be leveraged by VM consolidation policies. These policies are intended to be implemented at the level of group managers and therefore can only be applied to subsets of hosts. This approach partially solves the problem of scalability of VM consolidation by the cost of losing the ability of optimizing the VM placement across all the nodes of the data center. OpenStack Neat enables scalability by distributed underload/overload detection and VM selection, and potentially replicating the VM placement controllers. In contrast to Snooze, it is able to apply global VM placement algorithms for the selected migration VMs by taking into account the full set of hosts. Another difference is that OpenStack Neat transparently integrates with OpenStack, a mature open-source cloud platform widely adopted and supported by the industry, thus ensuring long-term development of the platform.

The second category of related work includes various theoretical approaches to the problem of dynamic VM consolidation. These can be further divided into application-specific and application-agnostic approaches. In application-specific approaches, the authors make assumptions about the

---

¶The Entropy VM manager, http://entropy.gforge.inria.fr/.
‖The Snooze cloud manager, http://snooze.inria.fr/.

execution environment or applications. For example, there are a number of works where the authors focus on enterprise environments by developing static and semi-static consolidation techniques [6, 7], assuming the knowledge of application priorities [8], or applying off-line profiling of applications [9]. There has been extensive research on VM consolidation tailored to specific application types, such as focusing on Web [10–15] and high-performance computing (HPC) applications [16, 17] or managing both Web and HPC applications [18, 19]. In contrast, our approach is application agnostic, meaning that there are no application-specific assumptions. This makes the system easily applicable to IaaS environments, where various types of user applications share computing resources.

One of the first works in which dynamic VM consolidation was applied to minimize energy consumption in an application-agnostic way has been performed by Nathuji and Schwan [20]. The authors have proposed an architecture of a system where the resource management is divided into local and global policies. At the local level, the system leverages power management strategies of the guest operating system (OS). The global manager applies its policy to decide whether the VM placement needs to be adapted. Zhu *et al.* [21] studied a similar problem of automated resource allocation and capacity planning. They proposed three individual controllers each operating at a different timescale, which place compatible workloads onto groups of servers, react to changing conditions by reallocating VMs, and allocate resources to VMs within the servers to satisfy the service-level agreements (SLAs).

Bobroff *et al.* [22] proposed an algorithm for dynamic VM consolidation under SLA constraints based on time-series analysis and forecasting of the resource demand. Nandi *et al.* [23] formulated the problem of VM placement as a stochastic optimization problem with a constraint on the probability of a host overload, taking into account three dimensions of resource constraints. The authors introduced two versions of the optimization problem and corresponding greedy heuristics depending on the distribution model of the VM resource usage: for known distributions and for unknown distributions with known mean and variance. A drawback of all the described application-agnostic approaches is that they are centralized: a single algorithm running on the master node leverages the global view of the system to optimize the VM allocation. Such a centralized approach limits the scalability of the system, when the number of physical machines grows over thousands and tens of thousands, like in Rackspace [30].

In this work, we follow a distributed approach proposed in our previous works [27–29], where every compute host locally solves the problems of underload/overload detection and VM selection. Then, it sends a request to a global manager to place only the selected migration VMs on other hosts. A similar approach was followed by Wood *et al.* [24] in their system called Sandpiper aimed at load balancing in virtualized data centers using VM live migration. The main objective of the system is to avoid host overloads referred to as hot spots by detecting them and migrating overloaded VMs to less-loaded hosts. The authors applied an application-agnostic approach, referred to as a black-box approach, in which VMs are observed from outside, without any knowledge of applications resident in the VMs. A hot spot is detected when the aggregate usage of a host's resources exceeds the specified threshold for k out of n last measurements, as well as for the next predicted value. Another proposed approach is gray-box, when certain application-specific data are allowed to be collected. The VM placement is computed heuristically by placing the most loaded VM to the least loaded host. The difference from our approach is that VMs are not consolidated; therefore, the number of active hosts is not reduced to save energy.

Feller *et al.* [25] proposed a set of algorithms for dynamic VM consolidation and evaluated them using the Snooze cloud manager discussed earlier. The approach and algorithms proposed in the paper can be seen as an extension of our previous work [27]. The authors proposed threshold-based heuristics for underload and overload detection using averaged resource usage data collected from the CPU, random access memory (RAM), and network interface of the hosts. In contrast to our previous works [27, 28], the VM placement algorithm is applied only to subsets of hosts, therefore disallowing the optimization of the VM placement across all the hosts of the data center.

3. SYSTEM DESIGN

The aim of the OpenStack Neat project is to provide an extensible framework for dynamic consolidation of VMs based on the OpenStack platform. Extensibility in this context means the ability to implement new VM consolidation algorithms and apply them in OpenStack Neat without the necessity to modify the source code of the framework itself. Different implementations of the algorithms can be plugged into the framework by modifying the appropriate options in the configuration file.

OpenStack Neat provides an infrastructure required for monitoring VMs and hypervisors, collecting resource usage data, transmitting messages and commands between the system components, and invoking VM live migrations. The infrastructure is agnostic of VM consolidation algorithms in use and allows implementing custom decision-making algorithms for each of the four sub-problems of dynamic VM consolidation: host underload/overload detection, VM selection, and VM placement. The implementation of the framework includes the algorithms proposed in our previous works [27–29]. In the following sections, we discuss the requirements and assumptions, integration of the proposed framework with OpenStack, and each of the framework's components.

3.1. Requirements and assumptions

The components of the framework are implemented in the form of OS services running on the compute and controller hosts of the data center in addition to the core OpenStack services. The current implementation of OpenStack Neat assumes a single instance of the controller responsible for finding a new placement of the VMs selected for migration. However, because of distributed underload/overload detection and VM selection algorithms, the overall scalability is significantly improved compared with existing completely centralized VM consolidation solutions.

OpenStack Neat relies on live migration to dynamically relocate VMs across physical machines. To enable live migration, it is required to set up a *shared storage* and correspondingly configure OpenStack Nova (i.e., the OpenStack Compute service) to use this storage for storing VM instance data. For instance, a shared storage can be provided using the Network File System (NFS) or the GlusterFS distributed file system [32].

OpenStack Neat uses a database (which can be distributed) for storing information about VMs and hosts, as well as resource usage data. It is possible to use the same database server used by the core OpenStack services. In this case, it is only required to create a new database and user for OpenStack Neat. The required database tables are automatically created by OpenStack Neat on the first launch of its services.

Another requirement is that all the compute hosts must have a user, which is enabled to switch the host into a low-power mode, such as *Suspend to RAM*. This user account is used by the global manager to connect to the compute hosts via the Secure Shell protocol and switch them into the sleep mode when necessary. More information on deactivating and reactivating physical nodes is given in Section 3.4.1.

3.2. Integration with OpenStack

OpenStack Neat services are installed independently of the core OpenStack services. Moreover, the activity of the OpenStack Neat services is transparent to the core OpenStack services. This means that OpenStack does not need to be configured in a special way to be able to take advantage of dynamic VM consolidation implemented by OpenStack Neat. It also means that OpenStack Neat can be added to an existing OpenStack installation without the need to modify its configuration.

The transparency is achieved by the independent resource monitoring implemented by OpenStack Neat, and the interaction with the core OpenStack services using their public APIs. The OpenStack APIs are used for obtaining information about the current state of the system and performing VM migrations. In particular, the APIs are used to obtain the current mapping of VMs to hosts, hardware characteristics of hosts, parameters of VM flavors (i.e., instance types), and VM states and to invoke VM live migrations. Although OpenStack Neat performs actions affecting the current state of the

system by relocating VMs across hosts, it is transparently handled by the core OpenStack services because VM migrations are invoked via the public OpenStack APIs, which is equivalent to invoking VM migrations manually by the system administrator.

One of the implications of this integration approach is that the VM provisioning and destruction processes are handled by the core OpenStack services, while OpenStack Neat discovers new VM instances through OpenStack APIs. In the following sections, we refer to hosts running the Nova Compute service, that is, hosting VM instances, as compute hosts and a host running the other OpenStack management services but not hosting VM instances as the controller host.

3.3. System components

OpenStack Neat is composed of a number of components and data stores, some of which are deployed on the compute hosts and some on the controller host, which can potentially have multiple replicas. As shown in Figure 2, the system is composed of three main components:

- *Global manager*—a component that is deployed on the controller host and makes global management decisions, such as mapping VM instances to hosts, and initiating VM live migrations
- *Local manager*—a component that is deployed on every compute host and makes local decisions, such as deciding that the host is underloaded or overloaded and selecting VMs to migrate to other hosts
- *Data collector*—a component that is deployed on every compute host and is responsible for collecting data on the resource usage by VM instances and hypervisors and then storing the data locally and submitting it to the central database, which can also be distributed
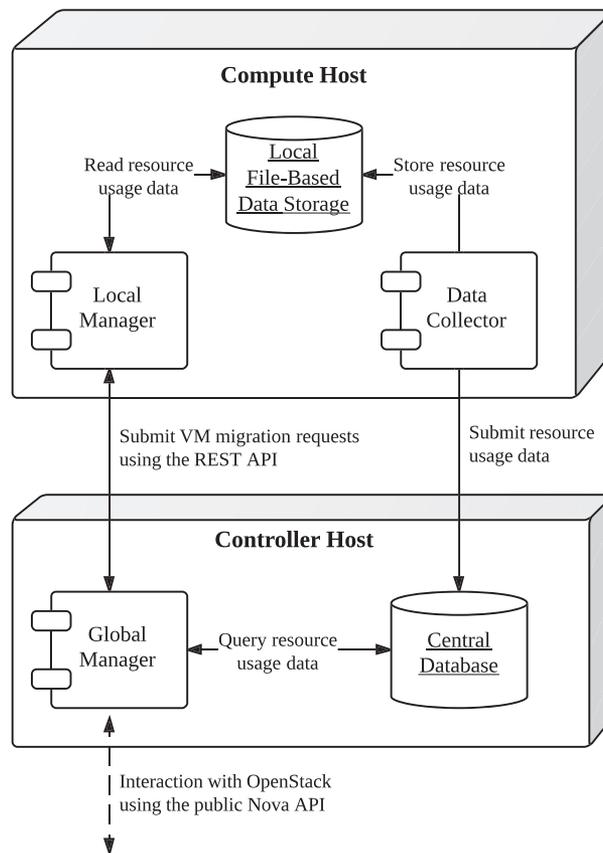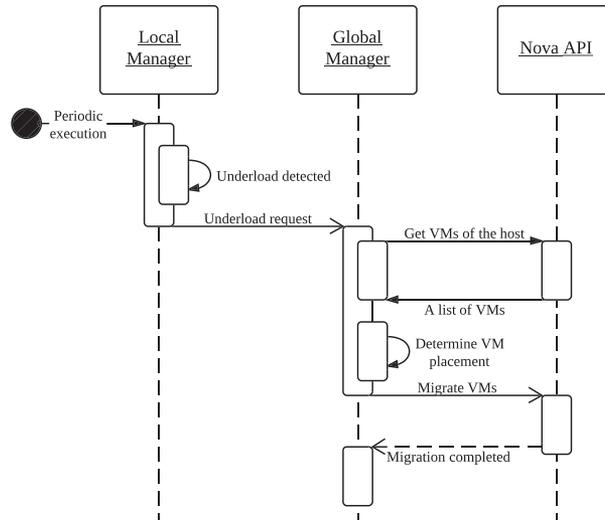

Figure 2. The deployment diagram.

Figure 3. The global manager: a sequence diagram of handling an underload request.

The deployment model may vary for each particular system depending on its requirements. For instance, the central database can be deployed on a separate physical node or be distributed across multiple physical nodes. The location and deployment of the database server are transparent to OpenStack Neat, which only requires a configuration parameter to be set to the network address of the database front-end server. For simplicity, in our experiment testbed, the database server is deployed on the same physical node hosting the global manager, as shown in Figure 2.

3.4. Global manager

The global manager is deployed on the controller host and is responsible for making VM placement decisions and initiating VM migrations. It exposes a representational state transfer (REST) Web service, which accepts requests from local managers. The global manager processes two types of requests: (1) relocating VMs from an underloaded host and (2) offloading a number of VMs from an overloaded host.

Figure 3 shows a sequence diagram of handling a host underload request by the global manager. First, a local manager detects an underload of the host using the specified-in-the-configuration underload detection algorithm. Then, it sends an underload request to the global manager including the name of the underloaded host. The global manager calls the OpenStack Nova API to obtain the list of VMs currently allocated to the underloaded host. This is required to obtain a view of the actual state of the system, as some VMs currently allocated to the host may still be in migration. Once the list of VMs is received, the global manager invokes the VM placement algorithm with the received list of VMs along with their resource usage and states of hosts fetched from the database as arguments. Then, according to the VM placement generated by the algorithm, the global manager submits the appropriate VM live migration requests to the OpenStack Nova API and monitors the VM migration process to determine when the migrations are completed. Upon the completion of the VM migrations, the global manager switches the now idle source host into the *Suspend to RAM* state. Because the *Suspend to RAM* state provide low-latency transitions, it is possible to quickly re-enable the host if required as discussed in Section 3.4.1.

As shown in Figure 4, handling overload requests is similar to underload requests. The difference is that instead of sending just the host name, the local manager also sends a list of Universally Unique Identifiers (UUID) of the VMs selected by the configured VM selection algorithm to be offloaded from the overloaded host. Once the request is received, the global manager invokes the specified configuration VM placement algorithm and passes as arguments the list of VMs received from the local manager to be placed on other hosts along with other system information. If some of the VMs are placed on hosts that are currently in the low-power mode, the global manager reactivates
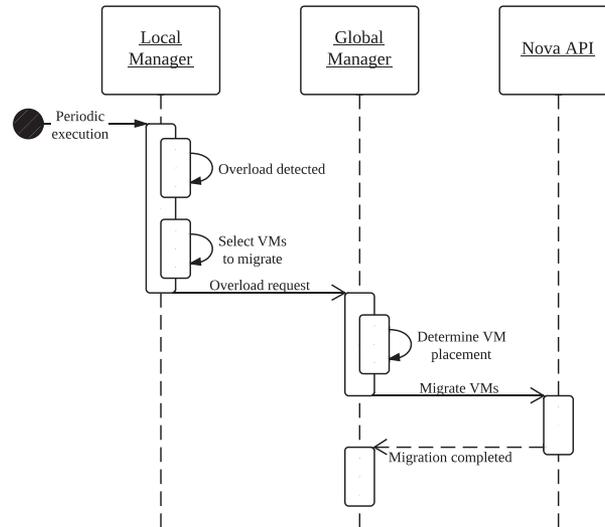
Figure 4. The global manager: a sequence diagram of handling an overload request.

them using the Wake-on-LAN technology, as described in Section 3.4.1. Then, similar to handling underload requests, the global manager submits VM live migration requests to the OpenStack Nova API.

3.4.1. Switching power states of hosts. One of the main features required to be supported by the hardware and OS in order to take advantage of dynamic VM consolidation to save energy is the Advanced Configuration and Power Interface (ACPI). The ACPI standard defines platform-independent interfaces for power management by the OS. The standard is supported by Linux, the target OS for the OpenStack platform. ACPI defines several sets of power states, the most relevant of which is the sleep state S3, referred to as *Suspend to RAM*. Meisner *et al.* [33] showed that power consumption of a typical blade server can be reduced from 450 W in the active state to just 10.4 W in the S3 state. The transition latency is currently mostly constrained by the power supply unit of the server, which leads to a total latency of approximately 300 ms. This latency is acceptable for the purposes of dynamic VM consolidation, as VM live migrations usually take tens of seconds.

The Linux OS provides an API to programmatically switch the physical machine into the sleep mode. The machine can be reactivated over the network using the Wake-on-LAN technology. This technology has been introduced in 1997 by the Advanced Manageability Alliance formed by Intel and IBM and is currently supported by most modern servers.

3.5. Local manager

The local manager component is deployed on every compute host as an OS service. The service periodically executes a function that determines whether it is necessary to reallocate VMs from the host. A high-level view of the workflow performed by the local manager is shown in Figure 5. At the beginning of each iteration, it reads from the local storage the historical data on the resource usage by the VMs and hypervisor stored by the data collector. Then, the local manager invokes the specified-in-the-configuration underload detection algorithm to determine whether the host is underloaded. If the host is underloaded, the local manager sends an underload request to the global manager's REST API to migrate all the VMs from the host and switch the host to a low-power mode.

If the host is not underloaded, the local manager proceeds to invoking the specified-in-the-configuration overload detection algorithm. If the host is overloaded, the local manager invokes the configured VM selection algorithm to select VMs to offload from the host. Once the VMs to migrate
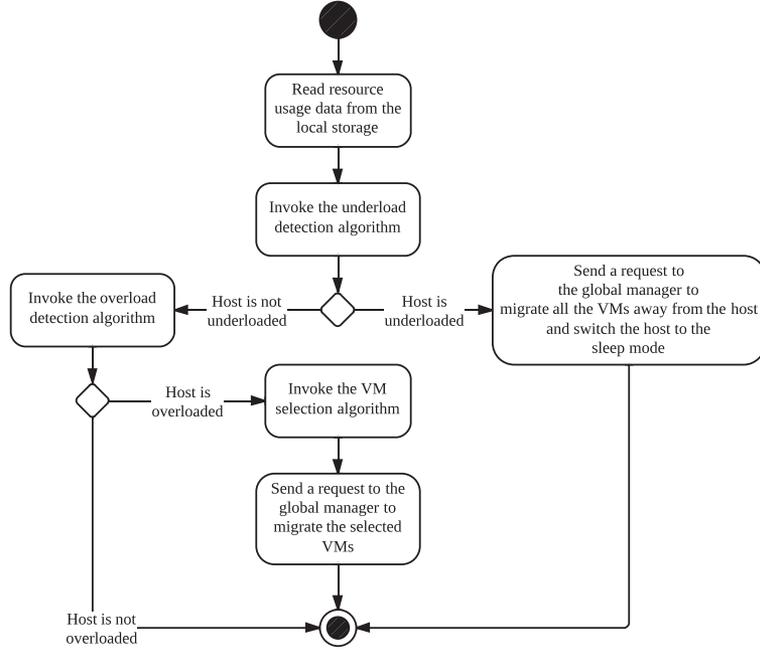
Figure 5. The local manager: an activity diagram.

from the host are selected, the local manager sends an overload request to the global manager's REST API to migrate the selected VMs. Similar to the global manager, the local manager can be configured to use custom underload detection, overload detection, and VM selection algorithms using the configuration file.

3.6. Data collector

The data collector is deployed on every compute host as an OS service and periodically collects the CPU utilization data for each VM running on the host, as well as data on the CPU utilization by the hypervisor. The collected data are stored in the local file-based data store and also submitted to the central database. The data are stored as the average megahertz consumed by a VM during the last measurement interval of length τ. In particular, the CPU usage $C_i^v(t_0, t_1)$ of a VM i, which is a function of the bounds of a measurement interval $[t_0, t_1]$, is calculated as shown in (1).

$$C_i^v(t_0, t_1) = \frac{n_i^v F\left(\tau_i^v(t_1) - \tau_i^v(t_0)\right)}{t_1 - t_0}, \tag{1}$$

where n_i^v is the number of virtual CPU cores allocated to VM i, F is the frequency of a single CPU core in megahertz, and $\tau_i^v(t)$ is the CPU time consumed by VM i up to time t. The CPU usage of the hypervisor $C_j^h(t_0, t_1)$ is calculated as a difference between the overall CPU usage and the CPU usage by the set of VMs allocated to the host, as shown in (2).

$$C_j^h(t_0, t_1) = \frac{n_j^h F\left(\tau_j^h(t_1) - \tau_j^h(t_0)\right)}{t_1 - t_0} - \sum_{i \in \mathcal{V}_j} C_i^v(t_0, t_1), \tag{2}$$

where n_j^h is the number of physical cores of the host j, $\tau_j^h(t)$ is the CPU time consumed by the host overall up to time t, and $\mathcal{V}_j$ is the set of VM allocated to host j. The CPU usage data are stored as

integers. This data format is used because of its portability: the stored values can be approximately converted to the CPU utilization percentages for any host or VM type. In general, to more accurately handle heterogeneous environments, this performance metric can be replaced with a more abstract measure, such as EC2 compute units in Amazon EC2.

The actual data are obtained using libvirt's API[**] in the form of the CPU time consumed by VMs and hosts overall to date. Using the CPU time collected at the previous time step, the CPU time for the last time interval is calculated. According to the CPU frequency of the host and the length of the time interval, the CPU time is converted into the required average megahertz consumed by the VM over the last time interval. Then, using the VMs' CPU utilization data, the CPU utilization by the hypervisor is calculated. The collected data are both stored locally and submitted to the database.

At the beginning of every iteration, the data collector obtains the set of VMs currently running on the host using the Nova API and compares them with the VMs running on the host at the previous time step. If new VMs have been found, the data collector fetches the historical data about them from the central database and stores the data in the local file-based data store. If some VMs have been removed, the data collector removes the data about these VMs from the local data store.

While OpenStack Neat oversubscribes the CPU of hosts by taking advantage of information on the real-time CPU utilization, it does not overcommit RAM. In other words, RAM is still a constraint in placing VMs on hosts; however, the constraint is the maximum amount of RAM that can be used by a VM statically defined by its instance type, rather than the real-time RAM consumption. The reason for that is that RAM is a more critical resource compared with the CPU, as an application may fail because of insufficient RAM, whereas insufficient CPU may just slow down its execution.

3.7. Data stores

As shown in Figure 2, the system contains two types of data stores:

- *Central database*—a database server, which can be deployed either on the controller host or on one or more dedicated hosts
- *Local file-based data storage*—a data store deployed on every compute host and used for temporarily caching the resource usage data to use by the local managers in order to avoid excessive database queries

The details about the data stores are given in the following subsections.

3.7.1. Central database.
The central database is used for storing historical data on the resource usage by VMs and hypervisors, as well as hardware characteristics of hosts. To ensure scalability, the database can be provided by any existing distributed database system, for example, the MySQL Cluster [34]. The database is populated by the data collectors deployed on compute hosts. There are two main use cases when the data are retrieved from the central database instead of the local storage of compute hosts. First, it is used by local managers to fetch the resource usage data after VM migrations. Once a VM migration is completed, the data collector deployed on the destination host fetches the required historical data from the database and stores them locally to be used by the local manager.

The second use case of the central database is when the global manager computes a new placement of VMs on hosts. VM placement algorithms require information on the resource consumption of all the hosts in order to make global allocation decisions. Therefore, every time there is a need to place VMs on hosts, the global manager queries the database to obtain the up-to-date data on the resource usage by hypervisors and VMs.

[**]The libvirt virtualization API, http://libvirt.org/.

3.7.2. Local file-based data store. A local manager at each iteration requires data on the resource usage by the VMs and hypervisor of the corresponding host in order to pass them to the underload/overload detection and VM placement algorithms. To reduce the number of queries to the database over the network, apart from submitting the data into the database, the data collector temporarily stores the data locally. This way, the local manager can just read the data from the local file storage and avoid having to retrieve the data from the central database.

4. VM CONSOLIDATION ALGORITHMS

As mentioned earlier, OpenStack Neat is based on the approach to the problem of dynamic VM consolidation, proposed in our previous work [29], which consists in dividing the problem into four sub-problems: (1) host underload detection; (2) host overload detection; (3) VM selection; and (4) VM placement. The implementation of the framework includes several algorithms proposed in our previous works [27–29] with slight modifications. The overview of the implemented algorithms is given for reference only. It is important to note that the presented algorithms are not the main focus of the current paper. The focus of the paper is the design of the framework for dynamic VM consolidation, which is capable of handling multiple implementations of consolidation algorithms and can be switched between the implementations through configuration.

4.1. Host underload detection

In our experiments, we use a simple heuristic for the problem of underload detection shown in Algorithm 1. The algorithm calculates the mean of the n latest CPU utilization measurements and compares it with the specified threshold. If the mean CPU utilization is lower than the threshold, the algorithm detects a host underload situation. The algorithm accepts three arguments: the CPU utilization threshold, the number of last CPU utilization values to average, and a list of CPU utilization measurements.

Algorithm 1: The averaging threshold-based underload detection algorithm

---
1 **Input:** threshold, n, utilization   **Output:** whether the host is underloaded
2 **if** *not empty* utilization **then**
3   utilization $\leftarrow$ last n values of utilization
4   mean_utilization $\leftarrow$ `sum`(utilization) / `len`(utilization)
5   **return** mean_utilization $\leq$ threshold
6 **return False**

4.2. Host overload detection

OpenStack Neat includes several overload detection algorithms, which can be enabled by modifying the configuration file. One of the simple included algorithms is the averaging threshold-based (THR) overload detection algorithm. The algorithm is similar to Algorithm 1, while the only difference is that it detects overload situations if the mean of the n last CPU utilization measurements is *higher* than the specified threshold.

Another overload detection algorithm included in the default implementation of OpenStack Neat is based on estimating the future CPU utilization using *local regression* (i.e., the locally weighted scatterplot smoothing (LOESS) method), referred to as the local regression robust (LRR) algorithm [28] shown in Algorithm 2. The algorithm calculates the LOESS parameter estimates and uses them to predict the future CPU utilization at the next time step, taking into account the VM migration time. In addition, the LRR algorithm accepts a safety parameter, which is used to scale the predicted CPU utilization to increase or decrease the sensitivity of the algorithm to potential overloads.

Algorithm 2: The Local Regression Robust (LRR) overload detection algorithm

---

1 Input: threshold, param, n, migration_time, utilization
2 Output: whether the host is overloaded
3 if `len`(utilization) $<$ n **then**
4   **return False**
5 estimates $\leftarrow$ `loess_parameter_estimates`(*last* n *values of* utilization)
6 prediction $\leftarrow$ estimates[0] + estimates[1] * (n + migration_time)
7 return param * prediction $\geq$ threshold

A more complex overload detection algorithm included in OpenStack Neat is the Markov overload detection (MHOD) algorithm, which enables the system administrator to explicitly specify a constraint on the overload time fraction (OTF) value as a parameter of the algorithm, while maximizing the time between VM migrations, thus improving the quality of VM consolidation [29].

4.3. VM selection

Once a host overload has been detected, it is necessary to determine what VMs are the best to be migrated from the host. This problem is solved by VM selection algorithms. An example of such an algorithm is simply randomly selecting a VM from the set of VMs allocated to the host. Another algorithm shown in Algorithm 3 is called minimum migration time maximum CPU utilization. This algorithm first selects VMs with the minimum amount of RAM to minimize the live migration time. Then, out of the selected subset of VMs, the algorithm selects the VM with the maximum CPU utilization averaged over the last n measurements to maximally reduce the overall CPU utilization of the host.

Algorithm 3: The Minimum Migration Time Maximum CPU utilization (MMTMC) algorithm

---

1 Input: n, vms_cpu_map, vms_ram_map   **Output:** a VM to migrate
2 min_ram $\leftarrow$ `min`(*values of* vms_ram_map)
3 max_cpu $\leftarrow$ 0
4 selected_vm $\leftarrow$ **None**
5 foreach vm*,* cpu *in* vms_cpu_map **do**
6   **if** vms_ram_map[vm] $>$ min_ram **then**
7   **continue**
8   vals $\leftarrow$ last n values of cpu
9   mean $\leftarrow$ `sum`(vals) / `len`(vals)
10   **if** max_cpu $<$ mean **then**
11   max_cpu $\leftarrow$ mean
12   selected_vm $\leftarrow$ vm
13 return selected_vm

4.4. VM placement

The VM placement problem can be seen as a bin packing problem with variable bin sizes, where bins represent hosts, bin sizes are the available CPU capacities of hosts, and items are VMs to be allocated with an extra constraint on the amount of RAM. As the bin packing problem is NP-hard, it is appropriate to apply a heuristic to solve it. OpenStack Neat implements a modification of the best fit decreasing (BFD) algorithm, which has been shown to use no more than $11/9 \cdot OPT + 1$ bins, where OPT is the number of bins of the optimal solution [35].

The implemented modification of the BFD algorithm shown in Algorithm 4 includes several extensions: the ability to handle extra constraints, namely, consideration of currently inactive hosts, and a constraint on the amount of RAM required by the VMs. An inactive host is only activated

when a VM cannot be placed on one of the already active hosts. The constraint on the amount of RAM is taken into account in the first fit manner; that is, if a host is selected for a VM as a best fit according to its CPU requirements, the host is confirmed if it just satisfies the RAM requirements. In addition, similar to the averaging underload and overload detection algorithms, the algorithm uses the mean values of the last n CPU utilization measurements as the CPU constraints. The worst-case complexity of the algorithm is $(n + m/2)m$, where n is the number of physical nodes and m is the number of VMs to be placed. The worst case occurs when every VM to be placed requires a new inactive host to be activated.

Algorithm 4: The Best Fit Decreasing (BFD) VM placement algorithm

---

1 **Input:** n, hosts_cpu, hosts_ram, inactive_hosts_cpu, inactive_hosts_ram, vms_cpu, vms_ram
2 **Output:** a map of VM UUIDs to host names
3 vm_tuples ← empty list
4 **foreach** vm, cpu *in* vms_cpu **do**
5   vals ← last n values of cpu
6   append a tuple of the mean of vals, vms_ram[vm], and vm to vm_tuples
7 vms ← `sort_decreasing`(vm_tuples)
8 host_tuples ← empty list
9 **foreach** host, cpu *in* hosts_cpu **do**
10   append a tuple of cpu, hosts_ram[host], host to host_tuples
11 hosts ← `sort_increasing`(host_tuples)
12 inactive_host_tuples ← empty list
13 **foreach** host, cpu *in* inactive_hosts_cpu **do**
14   append a tuple of cpu, inactive_hosts_ram[host], host to inactive_host_tuples
15 inactive_hosts ← `sort_increasing`(inactive_host_tuples)
16 mapping ← empty map
17 **foreach** vm_cpu, vm_ram, vm_uuid *in* vms **do**
18   mapped ← False
19   **while** *not* mapped **do**
20   allocated ← False
21   **foreach** _, _, host *in* hosts **do**
22   **if** hosts_cpu[host] ≥ vm_cpu *and* hosts_ram[host] ≥ vm_ram **then**
23   mapping[vm_uuid] ← host
24   hosts_cpu[host] ← hosts_cpu[host]- vm_cpu
25   hosts_ram[host] ← hosts_ram[host]- vm_ram
26   mapped ← True
27   allocated ← True
28   **break**
29   **if** *not* allocated **then**
30   **if** *not empty* inactive_hosts **then**
31   activated_host ← pop the first from inactive_hosts
32   append activated_host to hosts
33   hosts ← `sort_increasing`(hosts)
34   hosts_cpu[activated_host[2]] ← activated_host[0]
35   hosts_ram[activated_host[2]] ← activated_host[1]
36   **else**
37   **break**
38 **if** `len`(*vms*) == `len`(mapping) **then**
39   **return** mapping
40 **return** *empty map*

---

5. A BENCHMARK SUITE FOR EVALUATING DISTRIBUTED DYNAMIC VM CONSOLIDATION ALGORITHMS

Currently, research in the area of dynamic VM consolidation is limited by the lack of a standardized suite of benchmark software, workload traces, performance metrics, and evaluation methodology. Most of the time, researchers develop their own solutions for evaluating the proposed algorithms, which are not publicly available later on. This complicates further research efforts in the area because of the limited opportunities for comparing new results with prior solutions. Moreover, the necessity in implementing custom evaluation software leads to duplication of efforts. In this work, we propose an initial version of a benchmark suite for evaluating dynamic VM consolidation algorithms following the distributed approach of splitting the problem into four sub-problems discussed earlier. The proposed benchmark suite consists of four major components:

1. OpenStack Neat, a framework for distributed dynamic VM consolidation in OpenStack clouds providing a base system implementation and allowing configuration-based switching of different implementations of VM consolidation algorithms
2. A set of workload traces containing data on the CPU utilization collected every 5 min from more than a thousand PlanetLab VMs deployed on servers located in more than 500 places around the world [31]
3. A set of performance metrics capturing the following aspects: quality of VM consolidation, quality of service delivered by the system, overhead of VM consolidation in terms of the number of VM migration, and execution time of the consolidation algorithms
4. Evaluation methodology prescribing the approach of preparing experiments, deploying the system, generating workload using the PlanetLab traces, and processing and analyzing the results

We believe that the availability of such a benchmark suite will foster and facilitate research efforts and future advancements in the area of dynamic VM consolidation. In addition, researchers are encouraged to publicize and share the implemented consolidation algorithms to simplify performance comparisons with future solutions. One approach to sharing algorithm implementations is the addition of an extra package to the main branch of OpenStack Neat that will contain contributed algorithms. This would provide a central location, where anyone can find the up-to-date set of consolidation algorithms to use in their research. Therefore, processing and managing the inclusion of such submissions into the main public repository of OpenStack Neat will be performed as part of the project. The following sections provide more information on the workload traces, performance metrics, and evaluation methodology of the proposed benchmark suite. The performance evaluation discussed in Section 6 is an example of application of the benchmark suite.

5.1. Workload traces

To make experiments reproducible, it is important to rely on a set of input traces to reliably generate the workload, which would allow the experiments to be repeated as many times as necessary. It is also important to use workload traces collected from a real system rather than artificially generated, as this would help to reproduce a realistic scenario. This chapter uses workload trace data provided as a part of the CoMon project, a monitoring infrastructure of PlanetLab [31]. The traces include data on the CPU utilization collected every 5 min from more than a thousand VMs deployed on servers located in more 500 places around the world. Ten days of workload traces collected during March and April 2011 has been randomly chosen, which resulted in a total of 11,746 24-h-long traces. The full set of workload traces is publicly available online[††].

The workload from PlanetLab VMs is representative of an IaaS cloud environment, such as Amazon EC2, in the sense that the VMs are created and managed by multiple independent users and the infrastructure provider is not aware of what particular applications are executing in the VMs.

[††]The PlanetLab traces, http://github.com/beloglazov/planetlab-workload-traces.
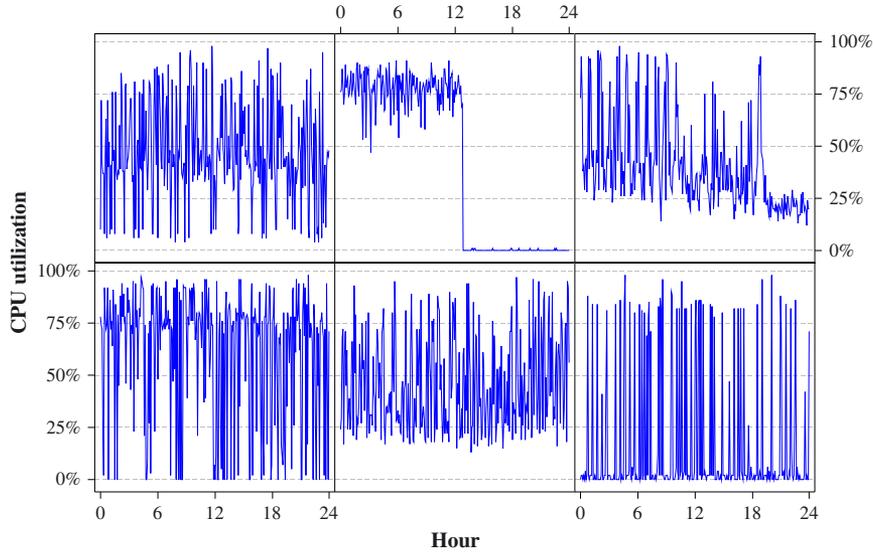
Figure 6. A sample of six selected CPU utilization traces over a 24-h period.

Furthermore, this implies that the overall system workload is composed of multiple heterogeneous applications, which also corresponds to a typical IaaS environment.

To stress the system in the experiments, the original workload traces have been filtered to leave only the ones that exhibit high variability. In particular, only the traces that satisfy the following two conditions have been selected: (1) at least 10% of time the CPU utilization is lower than 20% and (2) at least 10% of time the CPU utilization is higher than 80%. This significantly reduced the number of workload traces, resulting in only 33 out of 11,746 24-h traces left. The set of selected traces and filtering script is available online [36]. A sample of six traces is shown in Figure 6.

The resulting number of traces was sufficient for the experiments, whose scale was limited by the size of the testbed described in Section 6.1. If a larger number of traces is required to satisfy larger-scale experiments, one approach is to relax the conditions of filtering the original set of traces. Another approach is to randomly sample with replacement from the limited set of traces. If another set of suitable workload traces becomes publicly available, it can be included in the benchmark suite as an alternative.

5.2. Performance metrics

For effective performance evaluation and comparison of algorithms, it is essential to define performance metrics that capture the relevant characteristics of the algorithms. One of the objectives of dynamic VM consolidation is the minimization of energy consumption by the physical nodes, which can be a metric for performance evaluation and comparison. However, energy consumption is highly dependent on the particular model and configuration of the underlying hardware, efficiency of power supplies, implementation of the sleep mode, and so on. A metric that abstracts from the mentioned factors, but is directly proportional and can be used to estimate energy consumption, is the time of a host being idle, aggregated over the full set of hosts. By using this metric, the quality of VM consolidation can be represented by the increase in the aggregated idle time of hosts. However, this metric depends on the length of the overall evaluation period and the number of hosts. To eliminate this dependence, we propose a normalized metric referred to as the aggregated idle time fraction (AITF) defined as shown in (3).

$$AITF = \frac{\sum_{h \in \mathcal{H}} t_i(h)}{\sum_{h \in \mathcal{H}} t_a(h)}, \tag{3}$$

where $\mathcal{H}$ is a set of hosts, $t_i(h)$ is the idle time of the host h, and $t_a(h)$ is the total activity time of the host h. To quantify the overall QoS delivered by the system, we apply the aggregated overload time fraction (AOTF), which is based on the OTF metric [29] and defined as shown in (4).

$$AOTF(u_t) = \frac{\sum_{h \in \mathcal{H}} t_o(h, u_t)}{\sum_{h \in \mathcal{H}} t_b(h)}, \qquad (4)$$

where $t_o(h, u_t)$ is the overload time of the host h calculated according to the overload threshold u_t and $t_b(h)$ is the total busy (non-idle) time of the host h. We propose evaluating the overhead of dynamic VM consolidation in the system in terms of the number of VM migrations initiated as a part of dynamic consolidation. Apart from that, the execution time of various components of the system including the execution time of the VM consolidation algorithms is evaluated.

5.3. Performance evaluation methodology

One of the key points of the proposed performance evaluation methodology is the minimization of manual steps required to run an experiment through automation. Automation begins from scripted installation of the OS, OpenStack services, and their dependence on the testbed's nodes, as described in the OpenStack installation guide [32]. The next step is writing scripts for preparing the system for an experiment, which includes starting up the required services, booting VM instances, and preparing them for starting the workload generation.

While most of the mentioned steps are trivial, workload generation is complicated by the requirement of synchronizing the time of starting the workload generation on all the VMs. Another important aspect of workload generation is the way workload traces are assigned to VMs. Typically, the desired behavior is assigning a unique workload trace out of the full set of traces to each VM. Finally, it is necessary to create and maintain a specific level of CPU utilization for the whole interval between changes of the CPU utilization level defined by the workload trace for each VM.

This problem is addressed using a combination of a CPU load generation program[‡‡] and a workload distribution Web service and clients deployed on VMs [36]. When a VM boots from a pre-configured image, it automatically starts a script that polls the central workload distribution Web service to be assigned a workload trace. Initially, the workload distribution Web service drops requests from clients deployed on VMs to wait for the moment when all the required VM instances are booted up and ready for generating workload. When all clients are ready, the Web service receives a command to start the workload trace distribution. The Web service starts replying to clients by sending each of them a unique workload trace. Upon receiving a workload trace, every client initiates the CPU load generator and passes the received workload trace as an argument. The CPU load generator reads the provided workload trace file and starts generating CPU utilization levels corresponding to the values specified in the workload trace file for each time frame.

During an experiment, OpenStack Neat continuously logs various events into both the database and log files on each host. After the experiment, the logged data are used by special result processing scripts to extract the required information and compute the performance metrics discussed in Section 5.2, as well as the execution time of various system components. This process should be repeated for each combination of VM consolidation algorithms under consideration. Once the required set of experiments is completed, other scripts are executed to perform automated statistical tests and plotting graphs for comparing the algorithms.

The next section presents an example of application of the proposed benchmark suite and in particular applies the following: (1) OpenStack Neat as the dynamic VM consolidation framework; (2) the filtered PlanetLab workload traces discussed in Section 5.1; (3) the performance metrics defined in Section 5.2; and (4) the proposed evaluation methodology. The full set of scripts used in the experiments is available online [36].

[‡‡]The CPU load generator, http://github.com/beloglazov/cpu-load-generator.

6. PERFORMANCE EVALUATION

In this section, we apply the benchmark suite proposed in Section 5 to evaluate OpenStack Neat and several dynamic VM consolidation algorithm discussed in Section 4.

6.1. Experiment testbed

The testbed used for performance evaluation of the system consisted of the following hardware:

- 1 × Dell Optiplex 745

  - Intel(R) Core(TM) 2 CPU (two cores and two threads) 6600 at 2.40 GHz
  - 2-GB DDR2-667
  - Seagate Barracuda 80 GB, 7200 RPM SATA II (ST3808110AS)
  - Broadcom 5751 NetXtreme Gigabit Controller

- 4 × IBM System x3200 M3

  - Intel(R) Xeon(R) CPU (four cores and eight threads), X3460 at 2.80 GHz
  - 4-GB DDR3-1333
  - Western Digital 250 GB, 7200 RPM SATA II (WD2502ABYS-23B7A)
  - Dual Gigabit Ethernet (2 × Intel 82574L Ethernet Controller)

- 1 × Netgear ProSafe 16-Port 10/100 Desktop Switch FS116

The Dell Optiplex 745 machine was chosen to serve as the controller host running all the major OpenStack services and the global manager of OpenStack Neat. The four IBM System x3200 M3 servers were used as compute hosts, that is, running OpenStack Nova, and local managers and data collectors of OpenStack Neat. All the machines formed a local network connected via the Netgear FS116 network switch.

Unfortunately, there was a hardware problem preventing the system from taking advantage of dynamic VM consolidation to save energy. The problem was that the compute nodes of our testbed did not support the *Suspend to RAM* power state, which is the most suitable for the purpose of dynamic VM consolidation. This state potentially provides very low switching latency, in the order of 300 ms, while reducing the energy consumption to a negligible level [33]. Another approach would be to shut down inactive nodes; however, the latency is too high to quickly re-enable them. Therefore, rather than measuring the actual energy consumption by the servers, the AITF metric introduced in Section 5.2 was used to compare algorithms and estimate potential energy savings.

### 6.2. Experiment setup and algorithm parameters

From the point of view of experimenting with close-to-real-world conditions, it is interesting to allocate as many VMs on a compute host as possible. This would create a more dynamic workload and stress the system. At the same time, it is important to use full-fledged VM images representing realistic user requirements. Therefore, the Ubuntu 12.04 Cloud Image [37] was used in the experiments, which is one of the Ubuntu VM images available in Amazon EC2.

Because the compute hosts of the testbed contained limited amount of RAM, to maximize the number of VMs served by a single host, it was necessary to use a VM instance type with the minimum amount of RAM sufficient for Ubuntu 12.04. The minimum required amount of RAM was empirically determined to be 128 MB. This resulted in a maximum of 28 VMs being possible to instantiate on a single compute host. Therefore, to maximize potential benefits of dynamic VM consolidation on the testbed containing four compute nodes, the total number of VM instances was set to 28, so that in an ideal case, all of them can be placed on a single compute host, while the other three hosts are kept idle. Out of the 33 filtered PlanetLab workload traces discussed in Section 5.1, 28 traces were randomly selected, that is, one unique 24-h trace for each VM instance. The full set of selected traces is available online [36].

During the experiments, all the configuration parameters of OpenStack Neat were set to their default values except for the configuration of the overload detection algorithm. The overload detection algorithm was changed for each experiment by going through the following list of algorithms and their parameters:
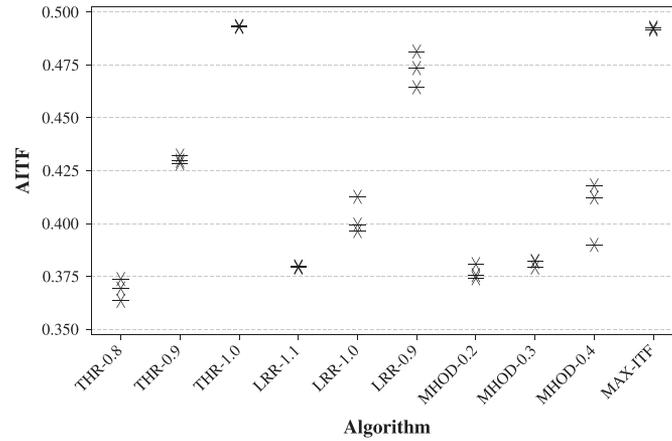
(a) The AITF metric
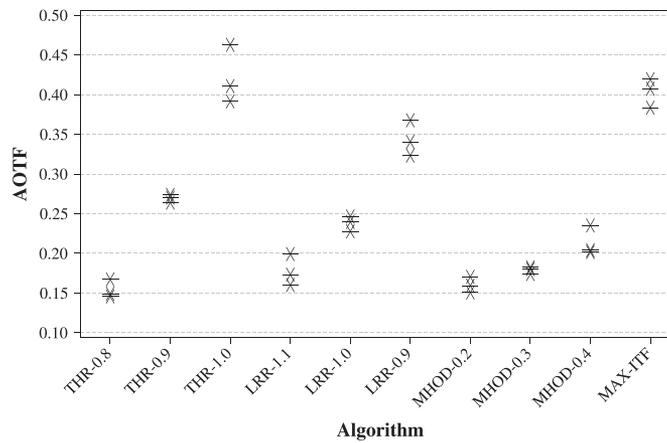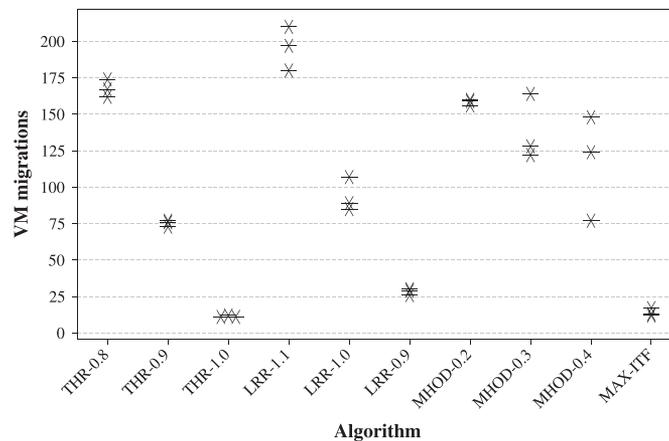
(b) The AOTF metric

(c) The number of VM migrations

Figure 7. The experiment results. (a) The AITF metric. (b) The AOTF metric. (c) The number of VM migrations.

1. MAX-ITF algorithm—a baseline algorithm that never detects host overloads, leading to the maximum ITF for the host, where the algorithm is used.
2. The THR algorithm with the *n* parameter set to 2, and the CPU utilization threshold set to 0.8, 0.9, and 1.0
3. The LRR algorithm with the safety parameter set to 0.9, 1.0, and 1.1
4. The MHOD algorithm with the OTF parameter set to 0.2, 0.3, and 0.4

Each experiment was run three times to handle the variability caused by random factors, such as the initial VM placement, workload trace assignment, and component communication latency. All the system initialization and result processing scripts, along with the experiment result packages are available online [36].

6.3. Experiment results and analysis

The results of the experiments are graphically depicted in Figure 7. The mean values of the obtained AITF and AOTF metrics and the number of VM migrations along with their 95% confidence intervals (CIs) are displayed in Table I. The results of MAX-ITF show that for the current experiment setup, it is possible to obtain high values of AITF of up to 50%, while incurring a high AOTF of more than 40%. All the THR, LRR, and MHOD allow tuning of the AITF values by adjusting the algorithm parameters. For the THR algorithm, the mean AITF increases from 36.9% to 49.2% with the corresponding decrease in the QoS level from 15.4% to 42.2% by varying the CPU utilization threshold from 0.8 to 1.0. The mean number of VM migrations decreases from 167.7 for the 80% threshold to 11.3 for the 100% threshold. The THR algorithm with the CPU utilization threshold set to 100% reaches the mean AITF shown by the MAX-ITF algorithm, which is expected as setting the threshold to 100% effectively disables host overload detection. Similarly, adjusting the safety parameter of the LRR algorithm from 1.1 to 0.9 leads to an increase of the mean AITF from 37.9% to 47.3% with a growth of the mean AOTF from 17.8% to 34.4% and decrease of the mean number of VM migrations from 195.7 to 28.3. THR-1.0 reaches the mean AITF of 49.2% with the mean AOTF of 42.2%, while LRR-0.9 reaches a close mean AITF of 47.3% with the mean AOTF of only 34.4%, which is a significant decrease compared with the AOTF of THR-1.0.

Varying the OTF parameter of the MHOD algorithm from 0.2 to 0.4 leads to an increase of the mean AITF from 37.7% to 40.7% with an increase of the mean AOTF from 16.0% to 21.4%. First of all, it is important to note that the algorithm meets the specified QoS constraint by keeping the value of the AOTF metric below the specified OTF parameters. However, the resulting mean AOTF is significantly lower than the specified OTF parameters: 17.9% for the 30% OTF and 21.4% for the 40% OTF. This can be explained by a combination of two factors: (1) the MHOD algorithm is parameterized by the per-host OTF, rather than AOTF, which means that it meets the OTF constraint for each host independently; (2) because of the small scale of the experiment testbed, a single underloaded host used for offloading VMs from overloaded hosts is able to significantly skew the AITF metric. The AITF metric is expected to be closer to the specified OTF parameter for large-scale OpenStack Neat deployments. A comparison of the results produced by LRR-1.1 and LRR-1.0 with

Table I. The experiment results (mean values with 95% CIs).

Algorithm	AITF	AOTF	VM migrations
THR-0.8	36.9% (35.6, 38.2)	15.4% (12.5, 18.3)	167.7 (152.7, 182.6)
THR-0.9	43.0% (42.6, 43.5)	27.0% (25.7, 28.1)	75.3 (70.2, 80.5)
THR-1.0	49.2% (49.2, 49.4)	42.2% (33.0, 51.3)	11.3 (9.9, 12.8)
LRR-1.1	37.9% (37.9, 38.0)	17.8% (12.8, 22.7)	195.7 (158.3, 233.0)
LRR-1.0	40.3% (38.1, 42.4)	23.8% (21.4, 26.1)	93.7 (64.6, 122.8)
LRR-0.9	47.3% (45.2, 49.4)	34.4% (28.8, 40.0)	28.3 (23.2, 33.5)
MHOD-0.2	37.7% (36.8, 38.5)	16.0% (13.5, 18.5)	158.3 (153.2, 163.5)
MHOD-0.3	38.1% (37.7, 38.5)	17.9% (16.8, 18.9)	138.0 (81.6, 194.4)
MHOD-0.4	40.7% (37.0, 44.4)	21.4% (16.7, 26.0)	116.3 (26.6, 206.0)
MAX-ITF	49.2% (49.1, 49.3)	40.4% (35.8, 44.9)	14.0 (7.4, 20.6)

Table II. Energy consumption estimates.

Algorithm	Energy (kWh)	Base energy (kWh)	Energy savings (%)
THR-0.8	25.99	34.65	24.99
THR-0.9	24.01	33.80	28.96
THR-1.0	22.09	32.93	32.91
LRR-1.1	25.66	34.50	25.63
LRR-1.0	24.96	34.18	26.97
LRR-0.9	22.60	33.20	31.93
MHOD-0.2	25.70	34.53	25.59
MHOD-0.3	25.59	34.48	25.76
MHOD-0.4	24.72	34.12	27.54
MAX-ITF	22.07	32.94	33.01

MHOD-0.2 and MHOD-0.4 reveals that the MHOD algorithm leads to lower values of the AOTF metric (higher level of QoS) for approximately equal values of AITF.

By using the obtained AITF and AOTF metrics for each algorithm and data on power consumption by servers, it is possible to compute estimates of potential energy savings relatively to a non-power-aware system assuming that hosts are switched to the sleep mode during every idle period. To obtain a lower bound on the estimated energy savings, it is assumed that when dynamic VM consolidation is applied, the CPU utilization of each host is 80% when it is active and non-overloaded, and 100% when it is overloaded. According to the data provided by Meisner *et al.* [33], power consumption of a typical blade server is 450 W in the fully utilized state, 270 W in the idle state, and 10.4 W in the sleep mode. Using the linear server power model proposed by Fan *et al.* [38] and the power consumption data provided by Meisner *et al.* [33], it is possible to calculate power consumption of a server at any utilization level.

To calculate the base energy consumption by a non-power-aware system, it is assumed that in such a system, all the compute hosts are always active with the load being distributed across them. Because the power model applied in this study is linear, it is does not matter how exactly the load is distributed across the servers. The estimated energy consumption levels for each overload detection algorithm, along with the corresponding base energy consumption by a non-power-aware system, and percentages of the estimated energy savings are presented in Table II.

According to the estimates, MAX-ITF leads to the highest energy savings over the base energy consumption of approximately 33% by the cost of substantial performance degradation (AOTF = 40.4%). The THR, LRR, and MHOD algorithms lead to energy savings from approximately 25% to 32% depending on the specified parameters. Similar to the earlier comparison of algorithms using the AITF metric, LRR-0.9 produces energy savings close to those of THR-1.0 (31.93% compared with 32.91%), while significantly reducing the mean AOTF from 42.2% to 34.4%. The MHOD algorithm produces approximately equal or higher energy savings than the LRR algorithm with lower mean AITF values, that is, higher levels of QoS, while also providing the advantage of specifying a QoS constraint as a parameter of the algorithm. The obtained experiment results confirm the hypothesis that dynamic VM consolidation is able to significantly reduce energy consumption in an IaaS cloud with a limited performance impact.

Table III lists the mean values of the execution time along with 95% CIs measured for each overload detection algorithm during the experiments for some of the system components: processing of underload and overload requests by the global manager (GM), overload detection algorithms executed by the local manager (LM), and iterations of the data collector (DC). Request processing by the global manager takes on average between 30 and 60 s, which is mostly determined by the time required to migrate VMs. The mean execution time of the MHOD algorithm is higher than those of THR and LRR, while still being under half a second, resulting in a negligible overhead considering that it is executed at most once in 5 min. The mean execution time of an iteration of the data collector is similarly under a second, which is also negligible considering that it is executed only once in 5 min.

Table III. The execution time of components in seconds (mean values with 95% CIs).

Algorithm	GM underload	GM overload	LM overload	DC
THR	33.5 (26.4, 40.5)	60.3 (54.0, 66.7)	0.003 (0.000, 0.006)	0.88 (0.84, 0.92)
LRR	34.4 (27.6, 41.1)	50.3 (47.8, 52.8)	0.006 (0.003, 0.008)	0.76 (0.73, 0.80)
MHOD	41.6 (27.1, 56.1)	53.7 (50.9, 56.6)	0.440 (0.429, 0.452)	0.92 (0.88, 0.96)
MAX-ITF	41.7 (9.6, 73.7)	—	0.001 (0.000, 0.001)	1.03 (0.96, 1.10)

7. SCALABILITY REMARKS AND FUTURE DIRECTIONS

Scalability and eliminating single points of failure are important benefits of designing a dynamic VM consolidation system in a distributed way. According to the approach adopted in the design of OpenStack Neat, the underload/overload detection and VM selection algorithms are able to inherently scale with the increased number of compute hosts. This is because they are executed independently on each compute host and do not rely on information about the global state of the system. In regard to the database setup, there exist distributed database solutions, for example, the MySQL Cluster [34].

On the other hand, in the current implementation of OpenStack Neat, there is assumed to be only one instance of the global manager deployed on a single controller host. This limits the scalability of VM placement decisions and creates a single point of failure. However, even with this limitation, the overall scalability of the system is significantly improved compared with existing completely centralized VM consolidation solutions. Compared with centralized solutions, the only functionality implemented in OpenStack Neat by the central controller is the placement of VMs selected for migration, which constitute only a fraction of the total number of VMs in the system. To address the problem of a single point of failure, it is possible to run a second instance of the global manager, which initially does not receive requests from the local managers and is automatically activated when the primary instance of the global manager fails. However, the problem of scalability is more complex because it is necessary to have multiple independent global managers concurrently serving requests from local managers.

Potentially, it is possible to implement replication of the global manager in line with OpenStack's approach to scalability by replication of its services. From the point of view of communication between the local and global managers, replication can be simply implemented by a load balancer that distributes requests from the local managers across the set of replicated global managers. A more complex problem is synchronizing the activities of the replicated global managers. It is necessary to avoid situations when two global managers place VMs on a single compute host simultaneously, because that would imply that they use an out-of-date view of the system state. One potential solution to this problem could be a continuous exchange of information between global managers during the process of execution of the VM placement algorithm, that is, if a host is selected by a global manager for a VM, it should notify the other global managers to exclude that host from their sets of available destination hosts.

There are several future directions for improvements of the current work. First of all, it is important to propose a synchronization model and implement replication of global managers to achieve a completely distributed and fault-tolerant dynamic VM consolidation system. Next, the data collector component should be extended to collect other types of data in addition to the CPU utilization that can be used by VM consolidation algorithms. Another direction is performance optimization and testing of OpenStack Neat on large-scale OpenStack deployments.

It is important to further refine the proposed initial version of the benchmark suite, for example, to provide a more concrete evaluation methodology along with tools for conducting automated experiments. It would also be beneficial to add alternative sets of workload traces, which would help to evaluate VM consolidation algorithms with a variety of workloads. As mentioned before, we accept contributions of implementations of VM consolidation algorithms that could be included in OpenStack Neat as a part of the proposed benchmark suite. Such a central repository of

VM consolidation algorithms would greatly benefit future research in the area by simplifying and facilitating performance comparisons of alternative solutions, and avoiding duplication of research efforts.

8. CONCLUSIONS

In this paper, we have proposed a design and implementation of an open-source framework for dynamic VM consolidation in OpenStack clouds, called OpenStack Neat. The framework follows a distributed model of dynamic VM consolidation, where the problem is divided into four sub-problems: host underload detection, host overload detection, VM selection, and VM placement. Through its configuration, OpenStack Neat can be customized to use various implementations of algorithms for each for the four sub-problems of dynamic VM consolidation. OpenStack Neat is transparent to the base OpenStack installation by interacting with it using the public APIs, and not requiring any modifications of OpenStack's configuration. We have also proposed a benchmark suite comprising OpenStack Neat as the base software framework, a set of PlanetLab workload traces, performance metrics, and methodology for evaluating and comparing dynamic VM consolidation algorithms following the distributed model.

The experiment results and estimates of energy consumption have shown that OpenStack Neat is able to reduce energy consumption by the compute nodes of a four-node testbed by 25–33%, while resulting in a limited application performance impact from approximately 15% to 40% AOTF. The MHOD algorithm has led to approximately equal or higher energy savings with lower mean AOTF values compared with the other evaluated algorithms, while also allowing the system administrator to explicitly specify a QoS constraint in terms of the OTF metric.

The performance overhead of the framework is nearly negligible, taking on average only a fraction of a second to execute iterations of the components. The request processing of the global manager takes on average between 30 and 60 s and is mostly determined by the time required to migrate VMs. The results have shown that dynamic VM consolidation brings significant energy savings with a limited impact on the application performance. The proposed framework can be applied in both further research on dynamic VM consolidation and real OpenStack cloud deployments to improve the utilization of resources and reduce energy consumption.

REFERENCES

1. Koomey J. *Growth in Data Center Electricity Use 2005 to 2010*. Analytics Press: Oakland, CA, 2011.
2. Gartner, Inc. *Gartner Estimates ICT Industry Accounts for 2 Percent of Global CO2 Emissions*. Gartner Press: Stamford, CT, USA, Release (April 2007).
3. The Open Compute project—energy efficiency. (Available from: http://opencompute.org/about/energy-efficiency/) [Accessed on 21 November 2012].
4. Hermenier F, Lorca X, Menaud J, Muller G, Lawall J. Entropy: a consolidation manager for clusters. *Proceedings of the 2009 ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments*, Washington, DC, USA, 2009; 41–50.
5. Feller E, Rilling L, Morin C. Snooze: a scalable and autonomic virtual machine management framework for private clouds. *Proceedings of the 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, Ottawa, Canada, 2012; 482–489.
6. Verma A, Dasgupta G, Nayak TK, De P, Kothari R. Server workload analysis for power minimization using consolidation. *Proceedings of the 2009 USENIX Annual Technical Conference*, San Diego, CA, USA, 2009; 28–41.
7. Speitkamp B, Bichler M. A mathematical programming approach for server consolidation problems in virtualized data centers. *IEEE Transactions on Services Computing (TSC)* 2010; **3**(4):266–278.
8. Cardosa M, Korupolu M, Singh A. Shares and utilities based power consolidation in virtualized server environments. *Proceedings of the 11th IFIP/IEEE Integrated Network Management (IM)*, Long Island, NY, USA, 2009; 327–334.
9. Kumar S, Talwar V, Kumar V, Ranganathan P, Schwan K. vManage: loosely coupled platform and virtualization management in data centers. *Proceedings of the 6th International Conference on Autonomic Computing (ICAC)*, Barcelona, Spain, 2009; 127–136.
10. Kusic D, Kephart JO, Hanson JE, Kandasamy N, Jiang G. Power and performance management of virtualized computing environments via lookahead control. *Cluster Computing* 2009; **12**(1):1–15.
11. Kusic D, Kandasamy N, Jiang G. Combined power and performance management of virtualized computing environments serving session-based workloads. *IEEE Transactions on Network and Service Management (TNSM)* 2011; **8**(3):245–258.

12. Jung G, Hiltunen MA, Joshi KR, Schlichting RD, Pu C. Mistral: dynamically managing power, performance, and adaptation cost in cloud infrastructures. *Proceedings of the 30th International Conference on Distributed Computing Systems (ICDCS)*, Genoa, Italy, 2010; 62–73.
13. Wang X, Wang Y. Coordinating power control and performance management for virtualized server clusters. *IEEE Transactions on Parallel and Distributed Systems (TPDS)* 2011; **22**(2):245–259.
14. Guenter B, Jain N, Williams C. Managing cost, performance, and reliability tradeoffs for energy-aware server provisioning. *Proceedings of the 30th Annual IEEE International Conference on Computer Communications (INFOCOM)*, Shanghai, China, 2011; 1332–1340.
15. Ardagna D, Panicucci B, Trubian M, Zhang L. Energy-aware autonomic resource allocation in multitier virtualized environments. *IEEE Transactions on Services Computing (TSC)* 2012; **5**(1):2–19.
16. Berral JL, Goiri Í, Nou R, Julià F, Guitart J, Gavaldà R, Torres J. Towards energy-aware scheduling in data centers using machine learning. *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*, Passau, Germany, 2010; 215–224.
17. Yang J, Zeng K, Hu H, Xi H. Dynamic cluster reconfiguration for energy conservation in computation intensive service. *IEEE Transactions on Computers* 2012; **61**(10):1401–1416.
18. Carrera D, Steinder M, Whalley I, Torres J, Ayguadé E. Autonomic placement of mixed batch and transactional workloads. *IEEE Transactions on Parallel and Distributed Systems (TPDS)* 2012; **23**(2):219–231.
19. Goiri Í, Berral J, Oriol Fitó J, Julià F, Nou R, Guitart J, Gavaldà R, Torres J. Energy-efficient and multifaceted resource management for profit-driven virtualized data centers. *Future Generation Computer Systems (FGCS)* 2012; **28**(5):718–731.
20. Nathuji R, Schwan K. VirtualPower: coordinated power management in virtualized enterprise systems. *ACM SIGOPS Operating Systems Review* 2007; **41**(6):265–278.
21. Zhu X, Young D, Watson BJ, Wang Z, Rolia J, Singhal S, McKee B, Hyser C, Gmach D., Gardner R., Christian T., Cherkasova L. 1000 Islands: integrated capacity and workload management for the next generation data center. *Proceedings of the 5th International Conference on Autonomic Computing (ICAC)*, Chicago, IL, USA, 2008; 172–181.
22. Bobroff N, Kochut A, Beaty K. Dynamic placement of virtual machines for managing SLA violations. *Proceedings of the 10th IFIP/IEEE International Symposium on Integrated Network Management (IM)*, Munich, Germany, 2007; 119–128.
23. Nandi B, Banerjee A, Ghosh S, Banerjee N. Stochastic VM multiplexing for datacenter consolidation. *Proceedings of the 9th IEEE International Conference on Services Computing (SCC)*, Honolulu, HI, USA, 2012; 114–121.
24. Wood T, Shenoy P, Venkataramani A, Yousif M. Black-box and gray-box strategies for virtual machine migration. *Proceedings of the 4th USENIX Symposium on Networked Systems Design & Implementation*, Cambridge, MA, USA, 2007; 229–242.
25. Feller E, Rohr C, Margery D, Morin C. Energy management in IaaS clouds: a holistic approach. *Proceedings of the 5th IEEE International Conference on Cloud Computing (IEEE CLOUD)*, Honolulu, HI, USA, 2012; 204–212.
26. Beloglazov A, Buyya R, Lee YC, Zomaya A. A taxonomy and survey of energy-efficient data centers and cloud computing systems. In *Advances in Computers*, Vol. 82, Zelkowitz M (ed.). Academic Press: Waltham, MA, USA, 2011; 47–111.
27. Beloglazov A, Abawajy J, Buyya R. Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing. *Future Generation Computer Systems (FGCS)* 2011; **28**(5):755–768.
28. Beloglazov A, Buyya R. Optimal online deterministic algorithms and adaptive heuristics for energy and performance efficient dynamic consolidation of virtual machines in cloud data centers. *Concurrency and Computation: Practice and Experience (CCPE)* 2012; **24**(13):1397–1420.
29. Beloglazov A, Buyya R. Managing overloaded hosts for dynamic consolidation of virtual machines in cloud data centers under quality of service constraints. *IEEE Transactions on Parallel and Distributed Systems (TPDS)* 2013; **24**(7):1366–1379. DOI: 10.1109/TPDS.2012.240.
30. Rackspace hosting reports second quarter 2012 results. (Available from: http://ir.rackspace.com/phoenix.zhtml?c=221673&p=irol-newsArticle&ID=1723357) [Accessed on 6 November 2012].
31. Park KS, Pai VS. CoMon: a mostly-scalable monitoring system for PlanetLab. *ACM SIGOPS Operating Systems Review* 2006; **40**(1):65–74.
32. Beloglazov A, Piraghaj SF, Alrokayan M, Buyya R. Deploying OpenStack on CentOS using the KVM hypervisor and GlusterFS distributed file system. *Technical Report*, CLOUDS-TR-2012-3, CLOUDS Laboratory, The University of Melbourne: Australia, 2012.
33. Meisner D, Gold B, Wenisch T. PowerNap: eliminating server idle power. *ACM SIGPLAN Notices* 2009; **44**(3): 205–216.
34. MySQL Cluster. (Available from: http://www.mysql.com/products/cluster/) [Accessed on 23 November 2012].
35. Yue M. A simple proof of the inequality FFD (L)< 11/9 OPT (L)+ 1, for all l for the FFD bin-packing algorithm. *Acta Mathematicae Applicatae Sinica (English Series)* 1991; **7**(4):321–331.
36. Beloglazov A. Scripts for setting up and analyzing results of experiments using OpenStack Neat. (Available from: http://github.com/beloglazov/ccpe-2014-experiments) [Accessed on 10 November 2013].
37. Ubuntu 12.04 (Precise Pangolin) Cloud images. (Available from: http://uec-images.ubuntu.com/precise/current/) [Accessed on 22 November 2012].
38. Fan X, Weber WD, Barroso LA. Power provisioning for a warehouse-sized computer. *Proceedings of the 34th Annual International Symposium on Computer Architecture (ISCA)*, San Diego, CA, USA, 2007; 13–23.