Allocation of Virtual Machines in Cloud Data Centers – A Survey of Problem Models and Optimization Algorithms*

Zoltán Ádám Mann
Budapest University of Technology and Economics

February 2015

Abstract

Data centers in public, private, and hybrid cloud settings make it possible to provision virtual machines (VMs) with unprecedented flexibility. However, purchasing, operating, and maintaining the underlying physical resources incurs significant monetary costs and also environmental impact. Therefore, cloud providers must optimize the usage of physical resources by a careful allocation of VMs to hosts, continuously balancing between the conflicting requirements on performance and operational costs. In recent years, several algorithms have been proposed for this important optimization problem. Unfortunately, the proposed approaches are hardly comparable because of subtle differences in the used problem models. This paper surveys the used problem formulations and optimization algorithms, highlighting their strengths and limitations, also pointing out the areas that need further research in the future.

Keywords: Cloud computing, data center, virtual machine, live migration, VM placement, VM consolidation, green computing

1 Introduction

In recent years, the increasing adoption of cloud computing has transformed the IT industry [20]. From a user's perspective, the practically unlimited scalability, the avoidance of up-front investments, and usage-based payment schemes make cloud computing a very attractive option. Beside globally available public cloud solutions, enterprises also take advantage of similar solutions in the form of private clouds and hybrid clouds.

Large, virtualized data centers are serving the ever growing demand for computation, storage, and networking. The efficient operation of data centers is increasingly important and complex [4]. Beside the traditional cost factors of equipment, staff, etc., energy consumption is playing an increasing role, because of both its costs and its environmental impact. According to a recent study, data center energy consumption is the fastest growing part of the energy consumption of the ICT ecosystem; moreover, the initial cost of purchasing the equipment for a data center is already outweighed by the cost of its ongoing electricity consumption [32].

Cloud data centers typically make extensive use of virtualization technology, in order to ensure isolation of applications while at the same time allowing a healthy utilization of physical resources. Virtual machines (VMs) are either provided directly to the customers in case of an Infrastructure-as-a-Service (IaaS) provider, or are used to wrap the provisioned applications in case of Software-as-a-Service (SaaS) or Platform-as-a-Service (PaaS) providers [108].

An attractive option for saving energy in data centers is to *consolidate* the virtual machines to the minimal number of physical hosts and switching the unused hosts off or at least to a less power-hungry mode of operation (e.g., sleep mode). However, too aggressive VM consolidation can lead to overloaded hosts with negative effects on the delivered quality of service (QoS), thus potentially violating the service level

*This is a preprint that has been submitted to a scientific journal for peer-review

agreements (SLA) with the customers. Hence, VM allocation must find the optimal balance between QoS and energy consumption [19, 89].

Good VM allocation also helps to serve as many customer requests as possible with the given set of resources, and thus amortizing the expenses related to purchasing, operations, and maintenance of the equipment (computing, network, and storage elements, as well as the physical data center infrastructure with cooling, redundant power supplies, etc.). In fact, achieving good utilization of server capacities was one of the key drivers behind the wide spread of virtualization technology. Today, virtualization and the live migration of VMs between hosts are key enablers of efficient resource allocation in data centers [8].

Beside using its own data center, a cloud provider can – in times of extremely high demand – use VMs from other providers as well, for example in a cloud federation or hybrid cloud setting [24]. This way, the cloud provider can serve its customers without restrictions. However, this further enlarges the search space for the best allocation.

In this paper, we focus on the VM allocation problem, i.e., determining the placement of VMs on physical hosts or using external providers, taking into account the QoS guarantees, the costs associated with using the hosts – with special emphasis on energy consumption – and the penalties resulting from VM migrations. Several algorithms have been proposed in the literature for this important and challenging optimization problem. However, these algorithms address slightly different versions of the problem, differing for example in the way the communication between hosts is modeled or how multi-core CPUs are handled. Lacking a generally accepted definition of the VM allocation problem, or some versions of the problem, many researchers came up with many different versions, and these differences can have substantial impact on algorithm runtime and/or on the applicability of the algorithm. This somewhat chaotic situation is even worsened by the fact that some authors failed to explicitly and precisely define the version of the problem that they are addressing, so that this must be figured out indirectly from the algorithms that they proposed or the way they evaluated their algorithms.

The primary aim of this paper is to "tidy up" the relevant problem formulations. Specifically, we start with a discussion of the context and the actors of the VM allocation problem (Section 2), followed by a description of the characteristics of the problem in Section 3. Section 4 presents a survey of the problem formulations existing in the literature, showing how those works fit into our general framework. Although our main focus is on problem formulations, we complete the survey of the literature with a brief description of the algorithms that have been proposed and how they were evaluated (Section 5). This is followed by a more detailed description of the most important algorithmic works of the field (Section 6), a discussion of the areas that we believe will need further research in the future (Section 7), and our concluding remarks (Section 8).

In this paper, we are mostly concerned with the details of *problem formulations* and their *algorithmic implications*. Technical details relating to infrastructure, architecture, and implementation issues are covered only as necessary for the aim of the paper.

2 Problem context

The VM allocation problem is one of the core challenges of using the cloud computing paradigm efficiently. Cloud computing encompasses several different setups, and depending on this, also the VM allocation problem has different flavors.

Usually, cloud computing scenarios are classified along two dimensions [108, 90]. One dimension concerns the nature of the offered service, differentiating between three categories: Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS). The other dimension refers to whether the service is provisioned in-house (private cloud), by a public provider (public cloud), or a combination of the two (hybrid cloud). The three possibilities along both dimensions give rise to 9 different possibilities.

Another classification focuses on service deployment scenarios [62]. Here it is assumed that a Service Provider (SP) would like to deploy a service on the infrastructure provided by one or more Infrastructure Providers (IPs) [94]. Depending on the relationship(s) between the SP and IP(s), the following scenarios are distinguished [62]:

- Public cloud: the SP makes use of the IP's infrastructure offering available to the general public.
- Private cloud: the SP uses its own resources, so that it also acts as IP.

- Bursted cloud: a hybrid of the above two, in which both in-house resources and resources rented from a public IP are used.
- Federated cloud: the SP contracts only one IP, but the IP collaborates with other IPs to share the load in a manner that is transparent to the SP.
- Multi-cloud: the SP uses multiple IPs to deploy (parts of) the service.
- Cloud broker: the SP contracts a single broker, which contracts multiple IPs but hides the complexity of the multi-cloud setup from the SP.

From our point of view, the crucial observation is that in each scenario, there is a need to optimize the allocation of VMs to physical resources, but this optimization may be performed by different actors and may have different characteristics, depending on the exact setup [84, 37]. Using the classification of Li et al., the VM allocation problem occurs in the respective scenarios as follows:

- Public cloud: the IP must optimize the utilization of its resources, in order to find the best balance between the conflicting requirements on profitability, performance, dependability, and environmental impact.
- Private cloud: the same kind of optimization problem occurs for the provider that acts as both SP and IP in this case[1].
- Bursted cloud: two slightly different optimization problems occur:
  - The IP must solve the same kind of optimization problem as above.
  - The SP must solve a similar problem for its own resources, extended by the possibility to off-load some VMs to an external IP.
- Federated cloud: the IPs must solve an optimization problem similar to the one the SP faces in the bursted cloud setup, i.e., optimization of own resources coupled with workload off-loading decisions.
- Multi-cloud: again, two different optimization problems occur:
  - The IPs must solve the same kind of optimization problem as in the public cloud setup.
  - The SP must solve an optimization problem in which the optimal allocation of parts of the service to the IPs is decided.
- Cloud broker: from an optimization point of view, this is the same as the multi-cloud scenario, with the broker taking the role of the SP.

In the following, we try to describe the VM allocation problem in a manner that is general enough to cover the above variants, and make the differences explicit only when this is necessary. We use the term *Cloud Provider (CP)* to refer to the entity who must carry out the VM allocation (which can be either the SP or the IP, depending on the setup). We assume that the CP must allocate VMs to a set of available resources. In general, there can be two kinds of resources: they can belong either directly to the CP, or the CP can also rent resources from *external CPs (eCPs)*. Depending on the specific setup, it is possible that the CP has only own resources and there are no eCPs, but it is also possible that the CP has no own resources, it can only select from eCPs. The case in which both internal resources and eCPs are available can be seen as the common generalization of all of the above scenarios.

3 Problem characteristics

Depending on the exact setup, there can be some differences in the most appropriate problem formulation, but the main characteristics of the VM allocation problem are in most cases the following:

[1] However, there can be subtle differences, e.g., SLAs tend to be less formal, VM sizes are more flexible etc.

- The CP accommodates VMs on the available physical machines (PMs) or by renting capacity from eCPs.
- The number of VMs changes over time as a result of upcoming requests to create additional VMs or to remove existing VMs.
- The resource requirements (e.g., computational power, memory, storage, network communication) of a VM can vary over time.
- The PMs have given capacity in terms of these resources.
- The usage of resources incurs monetary costs and consumes electric power. The magnitude of the costs and power consumption may depend on the type, state, and utilization of the resources.
- VMs can be migrated from one PM to another by means of live migration. This takes some time and creates additional load for the involved PMs and the network.
- PMs that are not used by any VM can be switched to a low-energy state.
- If the QoS requirements of the customer are not met, this may result in a penalty.

In the following, we investigate these aspects in more details.

3.1 VMs

A VM is usually characterized by:

- The number of CPU cores
- Required CPU capacity per core (e.g., in MIPS)
- Required RAM size (e.g., in GB)
- Required disk size (e.g., in GB)

Additionally, there can be requirements concerning the communication (bandwidth, latency) between pairs of VMs or a VM and the customer.

All of a VM's resource requirements can vary over time. Depending on the type of application(s) running on the VM, the VM's resource requirements can be relatively stable, changing periodically (e.g., in daily rhythm), or oscillating chaotically. In order to optimize resource usage, the CP must be well aware of the current resource requirements of the VMs and, even more importantly, the resource requirements expected for the near future [44].

In a public cloud setting, it is common that the CP offers standardized types of VMs. In a private cloud setting, customers usually have more freedom in specifying the parameters of a requested VM.

3.2 Resources

The resources available to the CP can be of two types:

- PMs, owned by the CP
- eCPs, from which VMs can be leased

These two resource types are significantly different. The CP's own PMs are *white-box* resources: the CP has detailed information about their state (e.g., power consumption characteristics, current workload, temperature) and it is the CP's responsibility to optimize the usage of these resources. On the other hand, eCPs represent *black-box* resource pools: the CP has no knowledge about the underlying physical infrastructure, it only knows the interface to request and manage VMs. Obviously, the CP has no direct influence on the underlying physical resources in this case.

Another important difference is that utilizing VMs from eCPs incurs direct costs that are normally higher than using the CP's own resources, since they also cover the eCP's profit. Therefore, a CP will usually first try to use its own resources, and use eCPs only as an extension in times of demand peaks. It is also possible that a CP has no resources on its own, and uses eCPs only [39].

Own PMs can reside in one or more Data Centers (DCs). If two PMs reside in different DCs, this usually leads to higher latencies in the communication between them, compared to the case when both PMs are in the same DC. Also live migration is usually done only within DC boundaries.

3.3 PM characteristics

The *utilization* or *load* of a PM measures to what extent its resources are utilized by the VMs residing on it. The most critical resource in terms of utilization is the CPU. On the one hand, it is the CP's interest to achieve high CPU utilization, in order to make the best use of the available resources. On the other hand, if CPU load is too high, this makes it likely that the VMs residing on the given PM do not receive the required capacity, which may lead to SLA violations and damage customer satisfaction. Too high CPU load may also lead to over-heating and it can accelerate aging of the hardware. For these reasons, many researchers concentrated on CPU load.

However, other resources like memory or disk space can also become a bottleneck [95]. Of particular interest is the cache, because current virtualization technologies do not ensure isolation of the cache usage of individual VMs accommodated by the same PM, leading to contention between them [58, 100]. Thus, it is important to model and predict the performance interference that can be expected when co-locating a pair of VMs [56].

Power consumption of a PM is a monotonously increasing function of the CPU load [55]. Determining the exact dependence of power consumption on CPU load is a non-trivial problem on its own and is even application-dependent [57]. Also, the load of other system components (e.g., disk) may play an important role. However, a linear approximation of power consumption as a function of CPU load works quite well across a wide range of applications and platforms [83]. Hence, several authors assumed linear dependence on CPU load [6, 51, 38, 60, 40, 93].

The amount of energy actually consumed by a PM does not only depend on power efficiency, but also on the duration. As shown by Srikantaiah et al., consolidating an increased amount of workload on a server improves energy consumption up to a certain point, when the usage of some resource of the server starts to saturate. Further increasing the load of the server leads to a slow-down of the execution of applications; since jobs take longer, the energy per job starts to increase [89].

Energy consumption of a server has a substantial static component that does not depend on the load of the server: even if a PM is "empty," i.e., it accommodates no VM, its energy consumption is non-negligible. In order to save more energy, it is therefore necessary to switch empty PMs to a *low-energy state*. In the simplest case, a PM has two states: On and Off. More sophisticated models include multiple states with different characteristics, e.g., On, Sleep, Hibernate, and Off. Realistically, switching between states takes some time, the amount of which depends on the source and target states. For instance, switching between On and Sleep is usually much quicker than switching between On and Hibernate; however, Hibernate will consume less energy than Sleep [44]. Nevertheless, most of the existing works use only a simplified two-state model.

In order to react to variations in the utilization, PMs usually offer – either directly, or through the virtualization platform – several possibilities. Dynamic voltage and frequency scaling (DVFS) is widely used to scale up or down the frequency of the CPU: in times of high load, the frequency is scaled up in order to increase performance at the cost of higher power consumption, whereas in times of low load, it is scaled down to decrease power consumption [60]. Using virtualization, it is possible to explicitly size the VMs by defining their share of the physical resources, and VMs can also be resized dynamically [29]. Scaling requests from the VMs can be used by the virtualization layer to determine the necessary physical scaling [74].

3.4 eCP characteristics

eCPs may offer VMs in two possible ways: either the eCP pre-defined some VM configurations from which customers can choose (example: Amazon EC2), or customers can define their own VM configuration by

specifying the needed amount from each resource (example: IC Cloud); this can make a difference in the achievable efficiency [45].

There can be considerable differences between eCPs concerning prices and pricing schemes, and even the same eCP may offer multiple pricing schemes [63]. For example, some providers offer discounted long-term rental rates and higher rates for the pay-as-you-go model [39]. The latter is often based on time quanta like hours. Further, there may be a fee proportional to the usage of some resources like network bandwidth [61]. In recent years, a further pricing scheme emerged: spot instances, the price of which depends on the current load of the provider. When the provider has a lot of free capacity, spot instances are cheap, but they become more expensive when the load of the provider is getting higher. Consumers can specify until what price they would like to keep the spot instance [30].

3.5 Communication and networking

VMs are used by customers to perform certain tasks, which are often parts of a bigger application, e.g., tiers of a multi-tier application [51]. This results in communication between the VMs. In some cases, this can mean the transfer of huge amounts of data, which may lead to an unacceptable increase in latency or response time as well as increased energy consumption in the affected hardware elements (PMs, routers, switches, etc.).

For the above reasons, it is beneficial to place VMs that communicate intensively with each other on the same PM, or at least within the same DC [8]. On the other hand, VMs that belong to the same application may exhibit correlation between their loads, increasing the probability that they will peak at the same time; this also has to be considered carefully in the VM allocation [101].

In some cases, the available network bandwidth can become a bottleneck. Some authors model network bandwidth the same way as any other resource of the PMs [11, 26, 75, 85, 105]. Others focus specifically on the communication among the VMs and try to minimize the resulting communication cost [70] or makespan [5]. Some works use a detailed network model with one or more layers of switches and communication links among switches and between switches and PMs, based on different topologies [50, 11]. Analogous problems arise also concerning the communication between multiple clouds [13].

A strongly related issue is the mapping of data on storage nodes. Some applications use huge amounts of data that are to be mapped on specialized storage nodes, leading to considerable network traffic between compute nodes and storage nodes. In such cases, the placement of VMs on compute nodes and the placement of application data on storage nodes are two interrelated problems that must be considered together in order to avoid unnecessarily high network loads [59].

Beside communication among VMs and between VMs and storage nodes, there is also communication with entities outside the cloud. An important example are the users. In several applications, the response time experienced by users is critical. The response time is the sum of the network round trip time and the processing time, and can thus be optimized by serving user requests from a data center offering a combination of low latency to the respective user and quick processing [53].

3.6 SLAs

By SLA, we mean any agreement between the CP and its customers on the expected service quality. The SLA defines Service Level Objectives: key measures to determine the appropriateness of the service (e.g., availability or response time). The SLA can be a formal document, specifying exactly for each SLO the performance indicators, the way they are measured, target values, as well as financial penalties for the case of non-fulfillment [92]. However, in many cases – notably in private cloud settings, where the provider and the customers belong to the same organization – SLAs can be less formal and less detailed. It is also possible that there is no written SLA at all. But even in such a case, customers do have expectations about service quality, and failure to fulfill those expectations damages the reputation of the CP, which will in the long run lead to customer churn and thus to profit loss.

Hence it is in all cases the CP's financial interest to pay attention to the – explicit or implicit – SLAs and try to avoid or at least minimize the number of SLA violations. This constrains the consolidation opportunities because too aggressive VM consolidation and overbooking of PMs would increase the probability of SLA violations [95].

We may differentiate hard and soft SLOs. A *hard* SLO must be fulfilled in any case. A *soft* SLO should be fulfilled as much as possible, but may be violated (usually at the price of a financial penalty). From a problem formalization point of view, hard SLOs must be modeled as constraints, whereas soft SLOs are no constraints but the number violations of a soft SLO must be minimized, and hence it will be part of the objective function.

Another distinction concerns the level of abstraction of the SLOs. Basically, we can differentiate between *user-level* SLOs describing quality metrics as observed by users (e.g., application response time, application throughput) and *system-level* SLOs defining the underlying technical objectives (e.g., system availability). Generally, user-level SLOs are more appropriate indicators of service performance; nevertheless, from a provider point of view, it is easier to control the system-level metrics, which will then indirectly determine the user-level metrics. For this reason, translating user-level objectives to system-level requirements is an important problem on its own [28].

An *SLA violation* occurs if one or more of the SLOs are not met. In many cases, this is the result of a situation in which a VM is not being allocated the required capacity, for instance because of too aggressive consolidation. But also other factors, e.g. inappropriate sizing of VMs or inadequate elasticity solutions can lead to an inability to serve requests within the boundaries stated in the SLA.

3.7 Live migration

Live migration of a VM from one PM to another makes it possible to react to the changing resource requirements of the VMs [9]. For example, in times of low demand, several VMs can be consolidated to one PM, so that other PMs can be switched off, thus saving energy. When the resource demand of the VMs increases, they can be migrated to other PMs with a lower load, thus avoiding SLA violations. For these reasons, VM migration is a key ingredient of dynamic VM placement schemes [93].

On the other hand, VM migrations take time, create overhead, and can have adverse impact on SLA fulfillment [85]. A VM migration may increase the load of both the source and the target PM, puts additional burden on the network, and makes the migrated VM less responsive during migration [51]. Therefore, it is important to keep the number of live migrations at a reasonable level.

Understanding the exact impact of live migration is a difficult problem on its own. A possible model for predicting the duration and overhead of live migration was presented by Verma et al. [102, 103]. According to their findings, migration increases the load of the source PM, but not the load of the target PM. In contrast, other researchers also measured increased load on the target PM [85]. The quest for a universally usable model of migration overhead is still ongoing [91].

3.8 Actions of the CP

The CP has to update the VMs' placement in several cases:

- To react to a customer request [88]
- To react to critical situations [41] and changes in system load [87]
- In the course of a regular evaluation of the current placement, in order to improve overall optimization objectives (see Section 3.9)

The first case is quite obvious. If a customer requests a new VM, it must be allocated on a PM or eCP. If a customer requests the cancellation of an existing VM, it must be removed from the hosting PM or eCP. Although rarely considered in the literature, but a customer may also request a change in the parameters of a VM (e.g., resizing). In all these cases, the CP must make a change to the current placement of VMs. This may also be a good occasion to review and re-optimize the placement of other VMs. For example, if a VM was removed upon the request of the customer, and the affected PM hosts only one more VM with a small load, then it may make sense to migrate that VM to another PM, so that this PM can be switched off.

Often, a customer request consists of multiple VMs, for example, VMs hosting the respective tiers of a multi-tier application [48]. Another important example is the case of elastic services: here, the number of VMs that take part in implementing the service changes automatically based on system-load (auto-scaling)

[62, 52]. In such cases, it is important to consider the placement of the affected VMs jointly, in order to avoid excessive communication costs [1].

The CP must also react to unplanned situations, like overloading of servers that may threaten SLA adherence [105], thermal anomalies [85], or breakdown of servers. Server unavailability may also be a planned situation (e.g., maintenance).

Beside the above reactive actions, a CP will also have to regularly review and potentially re-optimize the whole VM placement, in order to find a better fit to the changed demand of the existing VMs, modified eCP rental fees, modified electricity prices, or other changes that did not require immediate action but made the placement sub-optimal [87, 93]. Such a review may be carried out at regular times (e.g., every 10 minutes), or it may be triggered by specific events. For instance, a CP can continuously monitor the load of its servers or the performance of the VMs, and whenever some load or performance indicator goes below or above specified thresholds, this may be a reason to re-consider the VM placement.

Re-optimizing the VM placement may consist of one or more of the following actions:

- Migration of a VM from one host to another one
- Switching the state of a PM
- Starting/ending the rental of a VM from an eCP
- VM re-sizing

Increasing or decreasing the resource allotment of a VM ("VM re-sizing") can take multiple forms. In the case of VMs mapped on a PM owned by the CP, the VMM (Virtual Machine Monitor) can be instructed to set the resources allocated to the respective VMs as necessary [105, 42, 102]. In the case of VMs rented from eCPs, it may make sense to re-pack the application into VMs of different size, e.g., into a smaller number of larger VMs. This gives rise to an interesting balance between horizontal elasticity (number of VMs for the given service) and vertical elasticity (size of the VMs) [87].

3.9 Objectives

VM placement is inherently a multi-objective problem [38, 97, 107]. The following is a list of typical objectives for the CP:

- Monetary objectives:
  - Minimize fees paid to eCPs
  - Minimize operations costs
  - Amortize capital expenditures
  - Maximize income from customers
  - Avoid penalties
- Performance-related objectives:
  - Satisfy service-level objectives (availability, response time, makespan etc.)
  - Minimize number of SLA violations
- Energy-related objectives:
  - Minimize overall energy consumption
  - Minimize number of active PMs
  - Minimize carbon footprint
- Technical objectives:
  - Minimize number of migrations
  - Maximize utilization of resources

- Balance load among PMs
- Minimize network traffic
- Avoid overheating of hardware units

Of course, not all of these goals are independent from each other, e.g., several other objectives can be transformed to a monetary objective. Nevertheless, there are several independent or even conflicting objectives that VM placement should try to optimize. Given k objectives, in order to come to a well-defined optimization problem, one common technique is to constrain $k-1$ of the objectives and optimize the last one; another possibility is to optimize the weighted sum of the k objectives.

4 Problem models in the literature

A huge number of papers have been published about different versions of the VM allocation problem. In the following, we first give a categorization in Section 4.1, and then review the problem models of the most important works. Most existing works concentrate on either the Single-DC or the Multi-IaaS problem (defined in Section 4.1), which are quite different in nature; these problem formulations are discussed in Sections 4.2 and 4.3, respectively. Finally, some other problem models are described in Section 4.4.

4.1 Important special cases and subproblems

The problem described in Section 3 is very general. Most authors investigated special cases or subproblems, the most popular of which are presented next. It should be noted that these problem variants are not necessarily mutually exclusive, so that a given work may deal with a combination of them.

4.1.1 The Single-DC problem

The subproblem that has received the most attention is the Single-DC problem. In this case, the CP has a single DC with a number of PMs, and there are no eCPs. Usually, the number of PMs is assumed to be high enough to serve all customer requests. Typical objectives are optimizing the utilization of resources and minimizing overall energy consumption, subject to performance constraints (SLAs). Since all PMs are in the same DC, network bandwidth is often assumed to be uniform and sufficiently high so that it can be ignored.

4.1.2 The Multi-IaaS problem

In this case, the CP does not own any PMs, it uses only leased VMs from multiple IaaS providers. Since there are no PMs, all concerns related to them – states and state transitions, sharing of resources among multiple VMs, load-dependent power consumption – are void. Power consumption plays no role, the main goals are minimizing the monetary costs associated with VM rental and maximizing performance. Since data transfer between the different IaaS providers can become a bottleneck, this also has to be taken into account.

It is important to mention that the literature on the Multi-IaaS problem is mostly unrelated to the literature on the Single-DC problem. On one hand, this is natural because the two problems are quite different. On the other hand, a hybrid cloud provider must solve a combination of these two problems. This is why we include both of them in our paper, and we expect increased convergence between them in the future.

4.1.3 The One-dimensional VM placement problem

In this often-investigated special case, only the computational demands and computational capacities are considered, and no other resources. Moreover, the CPU is taken to be single-core, making the problem truly one-dimensional.

The question whether one or more dimensions are taken into account is independent on whether own PMs or eCPs are used. In other words, the one-dimensional VM placement problem can be a special case of the Single-DC, the Multi-IaaS, or other problem formulations.

4.1.4 The On/Off problem

In this case, each PM has only two states: On and Off. Furthermore, the power consumption of PMs that are Off is assumed to be 0, while the power consumed by PMs that are On is the same positive constant for each PM, and dynamic power consumption is not considered. The transition between the two power states is assumed to be instantaneous. As a consequence, the aim is simply to minimize the number PMs that are On. This is an often-investigated special case of the Single-DC problem.

4.1.5 Online vs. offline optimization

As mentioned in Section 3.8, the CP must react immediately to customer requests. This requires local modifications: allocating a new VM to a host, possibly turning on a new host if necessary, or deallocating a VM from a host, possibly switching the host to a low-energy state if it becomes empty. Finding the best reaction to the customer request in the given situation is an *online* optimization task.

On the other hand, the CP can also – e.g., on regular occasions – review the status of all VMs and hosts, and possibly make global modifications, e.g., migrating VMs between hosts. Finding the best new configuration is an *offline* optimization task.

These are two distinct tasks, for which a CP may use two different algorithms.

It should be noted that there is some ambiguity in the literature on the terminology used to differentiate between the above two cases, and the terms "online" and "offline" are used by some authors to describe other problem characteristics. We use these notions in this sense because this is in line with their generally accepted meaning in the theory of algorithms.

4.1.6 Placement tasks

Closely related to online vs. offline optimization is what we may call the placement task. On the one hand, (i) initial placement and (ii) placement re-optimization must be differentiated: the former determines a placement for a new set of VMs, whereas the latter optimizes an existing placement. (The key difference is that placement re-optimization must use migrations, which is not necessary for initial placement.) On the other hand, based on the set of VMs for which the placement is determined, the following three different levels can be distinguished: (i) all VMs of the CP, (ii) a set of coupled VMs, e.g., the VMs implementing a given service, or (iii) a single VM. Since these are two independent dimensions, we get 6 possible placement tasks; all of them are meaningful, although some are rather rare (e.g., initial placement of all VMs occurs only when a new DC starts its operation). It should also be noted that some works addressed multiple placement tasks, e.g., initial placement of a single VM and placement re-optimization of all VMs.

4.1.7 The Load prediction problem

When the CP makes some change in the mapping of VMs or the states of PMs at time instance t_0, it can base its decision only on its observations of VM behavior for the period $t \leq t_0$; however, the decision will have an effect only for $t > t_0$. The CP could make ideal decisions only if it knew the future resource utilization of the VMs. Since these are not known, it is an important subproblem to *predict* the resource utilization values of the VMs or their probability distributions, at least for the near future.

Load prediction is seen by some authors as an integral part of the VM placement problem, whereas others do not consider it, either because VM behavior is assumed to be constant (at least in the short run), or it is assumed that load prediction is done by a separate algorithm. Load prediction may or may not be considered, independently from the types of resources, i.e., also within the Single-DC or Multi-IaaS problem.

4.2 The Single-DC problem

The Single-DC problem has received a lot of attention also before the cloud computing age, with the main objective of achieving good utilization of physical resources in a DC. Early works include Muse, a resource management system for hosting centers [27], approaches to using Dynamic Voltage Scaling for power management of server farms [46] and to dynamic provisioning of multi-tier internet applications [98], as well as first results on consolidation using VM migration [54]. The term "load unbalancing" was coined to describe the objective of consolidating load on few highly utilized PMs instead of distributing them among many PMs

with low utilization [79]. From about 2007, as virtualized data centers have become ever more prevalent, the amount of research on resource management in DCs has seen significant growth [15, 5, 99]. These works already exhibited all of the important characteristics of the problem: consolidation of the VMs on fewer PMs using migrations, taking into account service levels and load fluctuations.

In recent years, the handling of SLA violations has become more sophisticated and energy consumption has become one of the most crucial optimization objectives. For example, the work of Beloglazov and Buyya [9, 7, 6] and Guazzone et al. [42, 43] has focused specifically on minimizing energy consumption.

Energy minimization can be primarily achieved by minimizing the number of active servers; it is thus no wonder that many works focused only on this and ignored the dynamic power consumption of PMs (leading to the special case of the On/Off problem). Exceptions include the work of Jung et al., which treated dynamic power consumption as a linear function of CPU load [51], as well as the non-linear function used by Guazzone et al. [42], and the table-based approach used in pMapper [99].

Most of the works on the Single-DC problem consider only the CPU capacity of the PMs and the computational demand of the VMs, but no other resources, reducing the problem to a single dimension. Several authors mentioned this deficiency as an area for future research [8, 42]. Only few works take into account also memory [82, 88] or memory and I/O as further dimensions [72, 95, 105]. Moreover, the sharing of cores of multi-core CPU-s was hardly addressed explicitly. For example, Beloglazov and Buyya model a multi-core CPU by means of a single-core CPU with capacity equal to the sum of the capacities of the cores of the original multi-core CPU [9]. Another extreme is the approach of Ribas et al., which does consider multi-core CPU-s, but only the number of cores is taken into account, their capacity is not [82].

The majority of these works did not address the Load prediction problem. A notable exception is the early work of Bobroff et al. [15], which uses a stochastic model to predict probabilistically the future load of a VM based on past observations. More recently, Guenter et al. used linear regression for similar purposes in a slightly different setting without virtualization [44]. Beloglazov and Buyya introduce a Markov chain approach for a related, but perhaps somewhat simpler problem: to detect when a PM becomes overloaded [10].

Concerning the investigated SLAs, most works consider the number of occasions when a server is overloaded [7, 6, 15], which indirectly lead to SLA violations. Only few works considered directly the response time [42] or waiting time [86] as specific metrics with quantitative QoS requirements.

The main characteristics of some representative works are summarized in Table 1. The meaning of the table's columns is explained below. A full circle means that the cited work explicitly deals with the given characteristic as part of their problem formulation and algorithms; an empty circle means that the given work does not explicitly address it.

- Resources: the types of resources of VMs and PMs that are taken into account
  - CPU: computational capacity of the PM and computational load of the VMs are taken into account.
  - Cores: individual cores of a multi-core processor are differentiated.
  - Other: at least one resource other than the CPU (e.g., memory) is also taken into account.
- Energy: the way energy optimization is supported by the given approach
  - Switch off: the given approach aims at emptying PMs so that they can be switched to a low-power state.
  - Dynamic power: also the dynamic power consumption of the PMs is taken into account.
- Placement: the kind of placement task addressed by the given work
  - Initial: the initial placement of the VMs is determined.
  - Reoptimization: an existing placement is optimized.
  - All VMs: the placement of all VMs in the DC is determined.
  - VM set: the placement of a set of coupled VMs that together form a service is determined.
  - One VM: the placement of a single VM is determined.
- SLA: the way SLAs are handled (see also Section 3.6)

Table 1: Characteristics of problem models in the Single-DC problem

	Resources			Energy		Placement					SLA			Other				
Paper	CPU	Cores	Other	Switch off	Dynam. power	Initial	Reoptimization	All VMs	VM set	One VM	Soft	User-level	Priorities	Different PMs	Migration	Migration cost	Data transfer	Load predict.
[5]	●	○	○	○	○	●	○	○	●	○	○	●	○	●	○	○	●	○
[6]	●	○	○	●	●	●	●	●	○	●	●	○	○	●	●	○	○	○
[9]	●	○	○	●	●	●	●	●	○	●	●	○	○	●	●	●	○	○
[10]	●	○	○	○	○	○	○	○	○	○	●	○	○	○	●	○	○	●
[11]	●	○	●	○	○	●	○	●	○	○	○	○	○	●	○	○	●	○
[15]	●	○	○	●	○	○	●	●	○	○	●	○	○	○	●	○	○	●
[17]	●	○	○	○	○	●	●	○	●	○	○	○	●	●	●	●	○	○
[31]	○	○	○	●	○	○	●	●	○	○	●	●	○	○	○	○	○	○
[42]	●	○	○	●	●	○	●	●	○	○	●	●	○	●	●	○	○	○
[44]	○	○	○	●	○	○	○	○	○	○	●	○	○	○	○	○	○	●
[45]	○	●	●	○	○	○	●	●	○	○	○	○	○	○	●	○	○	○
[49]	●	○	●	○	○	●	○	○	●	○	○	○	○	●	○	○	●	○
[51]	●	○	○	●	●	○	●	●	○	○	●	●	○	○	●	●	○	●
[65]	○	○	○	○	○	●	○	○	●	○	○	○	○	●	○	○	○	○
[72]	●	○	●	●	○	●	●	●	○	○	○	○	○	○	●	○	○	○
[82]	○	●	●	●	○	●	○	●	○	○	○	○	○	●	○	○	○	○
[86]	○	○	○	●	○	○	●	●	○	●	●	●	●	○	●	○	○	○
[88]	●	○	●	●	○	○	●	●	○	●	○	○	○	●	●	●	○	○
[89]	●	○	●	●	●	●	●	●	○	●	○	●	○	○	●	○	○	○
[95]	●	○	●	○	○	●	○	○	●	●	●	○	○	●	○	○	○	●
[99]	●	○	○	●	●	○	●	●	○	○	○	●	○	●	●	●	○	○
[101]	●	○	○	●	○	○	●	●	○	○	●	○	○	●	●	○	○	●
[105]	●	○	●	○	○	○	●	●	○	○	○	○	○	●	●	●	○	●

  - Soft: soft SLAs are supported.
  - User-level: user-level SLAs are supported.
  - Priorities: VMs may have different priorities.
- Other: some other important aspects
  - Different PMs: differences in the capacity and/or power consumption of PMs is leveraged to find the best VM-to-PM mapping.
  - Migration: the approach leverages migration of VMs between PMs.
  - Migration cost: migration costs are taken into account and must be minimized.
  - Data transfer: the communication between VMs is taken into account.
  - Load prediction: the future load of the VMs is predicted by the approach based on past observations.

As can be seen in Table 1, there are many differences between the approaches that were presented in the literature. In fact, it is hard to find two that address exactly the same problem. Of course, there are some basic properties that are typical of most approaches, e.g. the CPU is considered in almost all works, as well as the possibility to migrate VMs and to switch off unused PMs. Other characteristics, such as the sharing of individual cores of a multi-core CPU among VMs or communication between VMs are still largely unexplored.

Of course, Table 1 should not be seen as a valuation of these works (assuming that more filled circles indicate a higher "score"). Also approaches that tackle a limited version of the problem can be highly valuable if that problem is practically meaningful and the approach addresses it in an effective and efficient way. It is also important to mention that we focus here only on problem models and algorithms, but some works include many other aspects. Indeed, some works describe complete systems that are successfully applied in practice, such as Mistral [51], Muse [27], pMapper [99], and Sandpiper [105].

4.3 The Multi-IaaS problem

As already mentioned, the Multi-IaaS problem is quite different from the Single-DC problem. In the Multi-IaaS problem, the utilization and state of PMs, as well as their energy consumption, are not relevant. On the other hand, monetary costs related to the leasing of VMs from eCPs, appear as a new factor to consider. In fact, some works consider quite sophisticated leasing fee structures: e.g., VMs reserved for longer periods may be cheaper than on-demand VMs [39], or the costs may consist of a fixed rental fee and usage-based variable fees for the used resources [61].

In many formulations of the Multi-IaaS problem, the entities that need to be mapped to resources are not VMs but (computational) tasks. This is not really a significant conceptual difference though: also in the Single-DC problem, the actual goal is to map applications or components of applications to resources, and VMs are just wrappers that facilitate the safe co-location of applications or components of applications on the same resources and their migration.

More importantly, communication and dependencies among the tasks are often considered important ingredients of the Multi-IaaS problem [14, 39, 75] – in contrast to the Single-DC problem, where communication among VMs is hardly considered.

In the Multi-IaaS problem, the tasks and their dependencies are often given in the form of a directed acyclic graph (DAG), in which the vertices represent the tasks and the edges represent data transfer and dependencies at the same time. Scientific workflows are popular examples of complex applications that are well suited for a DAG representation [75, 106]. The resulting problem, often called "workflow scheduling problem" [3], has the advantage of solid mathematical formalism using graph theory; moreover, it is similar to other multi-resource scheduling problems (e.g., multiprocessor scheduling), so that a rich arsenal of available scheduling techniques can be applied to it [78]. Beside minimizing cost, the other objective of such scheduling problems is to minimize the makespan of the workflow, i.e., the time it takes from start of the first task to finish of the last task.

The main characteristics of some representative works are summarized in Table 2. The meaning of full versus empty circles is the same as in Table 1. The meaning of the table's columns, where different from those of Table 1, is explained below.

Table 2: Characteristics of problem models in the Multi-IaaS problem

	Resources			Scheduling		Costs			Other	
Paper	CPU	Cores	Other	Dependencies	Makespan	Long-term rental	On-demand	Usage-based	Migration	Load prediction
[14]	○	○	○	●	●	○	●	○	○	○
[22]	●	○	○	○	●	○	●	○	●	○
[39]	○	●	○	●	●	●	●	○	○	○
[61]	○	○	●	○	●	●	○	●	○	○
[64]	●	○	○	○	○	○	●	○	●	○
[75]	●	○	●	●	●	○	●	●	○	●
[76]	●	○	○	●	○	○	●	●	○	○
[96]	○	○	○	○	○	○	●	○	○	○
[97]	○	○	○	○	●	○	●	○	○	○
[104]	●	○	○	○	●	○	●	○	○	●

- Scheduling: the way scheduling-related aspects are modeled
  - Dependencies: dependencies between tasks arising from data transfer are considered.
  - Makespan: minimization of the workflow's makespan is either an explicit objective or there is an upper bound on the makespan.
- Costs: the kinds of monetary costs of leased VMs that the approach takes into account
  - Long-term rental: discounted fees for VMs that are rented for a long term (e.g., multiple months)
  - On-demand: fees that are either proportional to the time the VM is used or charged for small time quanta (e.g., hourly), based on the number of time quanta the VM is used
  - Usage-based: fees that are proportional to the used amount of some resource, e.g., the number of transferred bytes to/from a VM
- Other: some miscellaneous aspects
  - Migration: the approach leverages migration of tasks between VMs or between eCPs.
  - Load prediction: the future load of the tasks is predicted by the approach based on past observations.

As can be seen from Table 2, computational capacity and computational load, which are mostly considered one-dimensional (i.e., without accurate modeling of multi-core CPUs) is also the focus of most works in the Multi-IaaS context, just like in the case of the Single-DC problem. Makespan minimization and the minimization of on-demand rental costs are considered in most works. The other aspects are rarely handled. Again, it is interesting to note how different the used problem formulations are.

4.4 Other problem formulations

Although most of the relevant works fall into either the Single-DC or the Multi-IaaS category, there are a few works that address some other, more general problems.

4.4.1 Multi-DC

An important generalization of the Single-DC problem is the *Multi-DC* problem, in which the CP possesses multiple DCs. For an incoming VM request, the CP must first decide in which DC the new VM should be provisioned and then on which PM of the selected DC. While the second step is the same as the Single-DC problem, choosing the most appropriate DC may involve completely different decision-making [1]. A possibility is to consider the different power efficiency and carbon footprint of the different DCs, taking into account that different DCs may have access to different energy sources, e.g., some DCs may be able to better leverage renewing energy sources. In an attempt to optimize overall carbon footprint, the CP may prefer to utilize such "green" DCs as much as possible [55].

4.4.2 Hybrid cloud

In most works that address hybrid cloud setups, the CP owns one DC and also has some eCPs at its disposal. This can be seen as a common generalization of the Single-DC and the Multi-IaaS problems.

Casalicchio et al. address this problem with an emphasis on the Single-DC subproblem. That is, the PMs are explicitly modeled, migrations between PMs are allowed but incur a cost, there is a sophisticated handling of SLAs, but communication and dependencies among VMs are not handled, similarly to many formulations of the Single-DC problem [24].

In contrast, the approach of Bittencourt et al. shows more similarity to formulations of the Multi-IaaS problem. Here, dependencies among the tasks are given in the form of a DAG, there is a hard deadline on the makespan, and the objective is to minimize the total VM leasing costs, as is common in workflow scheduling. The own DC of the CP is modeled as a special eCP, offering free resources, but only in limited quantity [12, 13].

Bossche et al. use a similar approach, which is largely based on the Multi-IaaS problem, and own DCs are modeled as special eCPs offering free resources in limited quantity [16]. They explicitly allow to have more than one own DC, so that this can be seen as a common generalization of the Multi-DC and Multi-IaaS problems. On the other hand, the model uses a number of restrictions, e.g., communication and dependencies between VMs are not supported, nor migration of VMs or aspects related to power consumption.

5 Overview of proposed algorithms

From a theoretical point of view, we must differentiate between exact algorithms that are guaranteed to always deliver the optimum, and heuristics that do not offer such a guarantee. Although the majority of the proposed algorithms are heuristics, also some exact algorithms have been proposed, so it makes sense to review the two groups separately.

As already mentioned, most of the literature deals with either the Single-DC problem or the Multi-IaaS problem, and these two are quite different. Interestingly, the exact methods proposed for the two problems are very similar, hence we review them together. On the other hand, the heuristics proposed for the two problems are quite different, so we review them separately.

5.1 Exact algorithms

In most cases, the exact algorithm consists of formulating the problem in terms of some mathematical programming formalism and using an existing solver to solve the mathematical program.

Integer Linear Programming (ILP) seems to be by far the most popular way to express both the Single-DC [5, 44, 65] and the Multi-IaaS problem [39, 61, 64], or even their common generalization [16] as a mathematical program. Several authors found that even the special case of ILP in which each variable is binary (BIP – Binary Integer Programming) is sufficient to express the constraints of the problem in a natural way [16, 61, 64, 65].

Some authors preferred to use non-linear constraints, leading to a Mixed Integer Non-Linear Programming (MINLP) formulation [43] or a Pseudo-Boolean (PB) formulation with binary variables and a combination of linear and non-linear constraints [82].

For all these mathematical programs, appropriate solvers are available, both as commercial and as open-source software packages. In each case, the solver will deliver optimal results, but its worst-case runtime is

exponential with respect to the size of the input, so that solving large-scale problem instances takes much too long. Most researchers turned to heuristics for this reason.

It is important to mention that an ILP formulation can be useful in devising a heuristic. Removing the integrality constraint, the resulting Linear Programming (LP) formulation can be solved in polynomial time. The result obtained this way may not be integer, but in some cases a rounding method can be used to turn it into an integer solution with near-optimal cost [5, 39].

5.2 Heuristics for the Single-DC problem

Several authors observed the similarity between the VM placement problem and the well-known bin-packing problem, in which objects of given weight must be packed into a minimum number of unit-capacity bins. Indeed, if only one dimension, e.g., the computational demand of the VMs and the computational capacity of the PMs is considered, and the aim is to minimize the number of PMs that are turned on, the resulting problem is very similar to bin-packing. There are some simple but effective heuristics for bin-packing, like First Fit (FF), in which each object is placed into the first bin where it fits, Best Fit (BF), in which each object is placed in the bin where it fits and the remaining spare capacity is minimal, and Worst Fit (WF), in which each object is placed in the bin where it fits and the remaining spare capacity is maximal. Despite their simplicity, these algorithms are guaranteed to deliver results that are at most 70% off the optimum [34, 35]. This approximation ratio can be improved if the objects are first sorted in decreasing order of their weights, leading to the modified algorithms First Fit Decreasing (FFD), Best Fit Decreasing (BFD) etc. Specifically, if OPT denotes the optimal number of bins, then FFD is guaranteed to use no more than $11/9OPT + 6/9$ bins [33].

These simple bin-packing heuristics can be easily adapted to the VM placement problem. Indeed, the usage of FF has been suggested [15], just like BF [9], WF [51, 65, 95], FFD [99, 101] and BFD [7, 6, 42]. It should be noted though that the approximation results concerning these algorithms on bin packing do not automatically carry over to the more complicated VM placement problem [68].

Metaheuristics have also been suggested, e.g., simulated annealing [47], genetic algorithms [40], and ant colony optimization [38].

Some authors proposed proprietary heuristics. Some of them are simple greedy algorithms [86, 105] or straight-forward selection policies [5, 7, 6, 88]. Others are rather complex: for example, the algorithm of Jung et al. first determines a target mapping by means of a Worst-Fit-like heuristic, but then uses an A^* tree traversal algorithm to create a reconfiguration plan, taking into account not only the adaptation costs, but also the cost of running the algorithm itself (which means that search space exploration is restricted if the algorithm has already run for a long time); moreover, this algorithm is carried out in a hierarchical manner, on multiple levels [51]. Mishra and Sahoo categorize both PMs and VMs according to what kind of resource is used by them most (from the three investigated dimensions, which are CPU, memory, and I/O) into so-called resource triangles, and attempt to match them on the basis of complementary resource triangles (e.g., a VM that uses the CPU most should be mapped on a PM where the CPU is the least used resource), at the same time also taking into account the utilization levels [72]. Verma et al. devised an algorithm that starts by analyzing the workload time series of the applications to determine an envelope of the time series that captures the bulk and the peak of the distribution, which information is then used to cluster the applications on the basis of correlating peaks, and then the application clusters are spread evenly on the necessary number of PMs [101]. However, none of these sophisticated heuristics offer performance guarantees in terms of approximation factors.

A different approach is to regard the VM placement problem as a *control* task, in which a controller tries to balance the utilization of the PMs between the conflicting objectives of minimizing power consumption and keeping performance levels, and to apply control-theoretic methods. This includes fuzzy control techniques [86] and distributed PID controllers [95].

5.3 Heuristics for the Multi-IaaS problem

The heuristic algorithms that have been proposed for the Multi-IaaS problem are quite heterogeneous. The simplest algorithms include list scheduling [14], greedy provisioning and allocation policies [104], greedy scheduling and clustering algorithms [75], and simple proprietary heuristics [22]. Metaheuristics have also

been suggested, e.g., particle swarm optimization [76]. Also, more sophisticated algorithms have been proposed, e.g., based on existing algorithms for the knapsack problem [61].

The above algorithms, whether simple or sophisticated, offer no performance guarantees, or at least, none has been proven. An exception is the work of Tsamoura et al., addressing a multi-objective optimization problem, in which makespan and cost are minimized simultaneously. That is, the aim is to find Pareto-optimal solutions in the time–cost space, and also the bids of eCPs are in the form of time–cost functions. Drawing on earlier results [77], an approximation algorithm with pseudo-polynomial runtime can be devised [97].

5.4 Algorithms for other problem formulations

As already mentioned in Section 4.4, there are few works considering other problem formulations, like the Multi-DC problem or hybrid cloud setups, and these works are similar to either the Single-DC or the Multi-IaaS problem. Accordingly, the algorithms that have been proposed for these problem formulations are also similar to the ones for the other problem variants.

In particular, Binary Integer Programming has been suggested to optimally solve the task allocation problem in a hybrid cloud scenario [16]. Hill climbing has also been used as a simple heuristic [24], as well as proprietary heuristics [12]. Heuristics inspired by bin-packing play a role here as well, e.g., First Fit [55], and Mills et al. compare several bin-packing-style heuristics in a multi-DC setup [71].

5.5 Evaluation of algorithms

Most papers also provide some evaluation of the algorithms they propose. In most cases, this evaluation is done empirically, but there are also some examples of rigorous mathematical analysis.

5.5.1 Rigorous analysis

Tsamoura et al. proved the correctness and complexity of their algorithms: an exact polynomial-time algorithm for a special case and an approximation algorithm with pseudo-polynomial runtime for the general case, albeit for a rather uncommon problem formulation [97].

For some restricted problem versions, polynomial-time approximation algorithms have been presented with rigorously proven approximation guarantees [1, 2, 18, 68].

Guenter et al. proved an important property of the linear program that they proposed: that its optimal solution will be integral, without explicit integrality constraints, thus allowing the use of an LP solver instead of a – much slower – ILP solver [44].

5.5.2 Empirical evaluation

In many cases, the evaluation was carried out using simulation. There are simulators specifically for cloud research, for example CloudSim [21], but many researchers used their own simulation environments. Relatively few researchers tested their algorithms in a real environment [31, 51, 57, 66, 70, 74, 101, 103, 105] or using a combination of real hardware and simulation [85, 93, 95, 104, 109]. It has to be added though that in most of these cases, the "real" environment used for evaluation was rather small (e.g., just a handful of PMs and VMs). Apparently, most researchers do not have the possibility to make experiments on large-scale real systems.

As a compromise between pure simulation and a real evaluation environment, several researchers used traces from real applications and real servers. Some research groups of industry players used traces from their own infrastructure [40, 44, 99, 109]. Others used publicly available workload traces, e.g., from the Parallel Workloads Archive (`http://www.cs.huji.ac.il/labs/parallel/workload/`) of the Hebrew University [36, 50, 86], the Grid Observatory (`http://grid-observatory.org/`) [85], PlanetLab (`http://www.planet-lab.org/`) [9, 10], or workload traces made available by Google (`https://code.google.com/p/googleclusterdata/`) [81, 82]. A related approach, taken by several researchers, was to use a web application with real web traces: for example, RUBiS, a web application for online auctions [25], has been used my multiple researchers with various web server traces [51, 95]; Wikipedia traces were also used [37]. Other benchmark applications used include the NAS Parallel Benchmarks (`http://www.nas.nasa.gov/publications/`

npb.html) [58, 73, 96, 103], the BLAS linear algebra package (http://www.netlib.org/blas/) [103] and the related Linpack benchmark (http://netlib.org/benchmark/hpl/) [31, 99].

6 Details of some selected works

Because of the sheer volume, it is impossible to provide a detailed description of all works in the field. However, we selected some of the most influential and most interesting works and give more details about them in the following. "Most influential" has been determined based on the yearly average number of citations that the given paper has received according to Google Scholar (http://scholar.google.com), as of February 2015, and this list has been extended with some other works that are – in our opinion – also of high importance to the field.

According to the above metric, the most influential papers are those of Beloglazov and Buyya from the University of Melbourne (one of those papers is joint work with Abawajy). They address the Single-DC problem, focusing on the single dimension of CPU capacity of PMs and CPU load of VMs. The main optimization objective is to consolidate the workload on the minimal number of PMs with the aim of minimizing energy consumption. As a secondary objective, also the number of migrations should be kept low. The authors' early works focus on analyzing the context of and requirements towards such an optimization framework, as well as architectural considerations and preliminary results on some efficient optimization heuristics [7, 8].

Those heuristics are presented in more detail in a later paper [6]. The main idea is to first remove all VMs from lightly used PMs so that they can be switched off and also remove some VMs from overloaded PMs so that they will not be overloaded. In a second phase, a new accommodating PM is searched for the removed VMs. The latter subproblem is seen as a special version of the bin-packing problem, in which the bins may have differing sizes (different PM capacities) and prices (different energy efficiency of the PMs). For this problem, the authors developed the Modified Best Fit Decreasing (MBFD) heuristic, which considers the VMs in decreasing order of load and allocates each of them to the PM with best energy efficiency that has sufficient capacity to host it. For the problem of selecting some VMs to migrate off an overloaded PM, the authors consider several heuristics. The Minimization of Migrations (MM) policy selects the minimum number of VMs that must be removed to let the PM's load go back to the normal range. The Highest Potential Growth (HPG) policy selects the VMs that have the lowest ratio of current load to requested load. Finally, the Random Choice (RC) policy selects the VMs to be removed randomly. The authors used the CloudSim framework to simulate a DC with 100 PMs and 290 VMs to evaluate the presented heuristics and compare them to a Non-Power Aware method (NPA), one using DVFS only, and a Single-Threshold (ST) VM selection algorithm. The simulation results, accompanied by a detailed statistical analysis, demonstrate the superiority of the presented methods with respect to energy consumption, number of SLA violations, and number of migrations. From the presented VM selection methods, the MM heuristic proved best.

In a related paper, the same authors provide a mathematical analysis of some rather restricted special cases or sub-problems of the single-DC problem [9]. In particular, they provide optimal offline and online algorithms for the problem of when to migrate a VM off from a PM, and prove an upper bound for the competitive ratio of online algorithms for the case of n homogeneous PMs. Besides, they also consider some adaptive heuristics for dynamic VM consolidation. The problem is the same as the one considered in the other works of the authors, and the algorithms are also similar, but are now adaptive: instead of using fixed thresholds for determining under-utilization and over-utilization, the thresholds now adapt to the variability of the VMs' load. For this, several methods are considered: Median Absolute Deviation (MAD), Interquartile Range (IQR), Local Regression (LR), and Robust Local Regression (RLR). The performance of the algorithms is evaluated again using CloudSim, but this time with a simulated DC with 800 heterogeneous PMs and real workload traces from PlanetLab. Also here, the authors carried out a very thorough statistical analysis to come to the conclusion that the LR method outperforms the others in terms of energy consumption and SLA violations.

Finally, yet another paper of the same authors looks at one specific sub-problem of VM consolidation: how to decide when a PM is overloaded, and consequently, when VMs should be removed [10]. Two conflicting goals are taken into account: on the one hand, the time when the host is overloaded should be minimized in order to avoid performance degradation and SLA violation; on the other hand, the overload detection method should signalize an overload only if absolutely necessary, in order to keep the utilization high and avoid unnecessary migrations. The authors devise a method using Markov chains for stationary workloads,

which can also be applied to non-stationary workloads by using the Multisize Sliding Window workload estimation technique. Simulation results on PlanetLab traces demonstrate the good performance of the proposed method.

Bobroff, Kochut, and Beaty from IBM Research investigated a similar problem: they also aim at minimizing the number of active PMs and the number of SLA violations by carefully consolidating VMs to PMs in a single DC [15]. Here, also, the problem is one-dimensional, with the CPU being the single investigated resource, and SLA violations are assumed to happen if the CPU of a PM is overloaded. As a generic solution framework, the authors propose the Measure-Forecast-Remap cycle, in which the workload consumption of VMs is measured, based on which their future resource demand is forecast, and a new VM-to-PM mapping is generated. This cycle is iterated at regular time intervals of length τ. (In the practical examples, τ is 15 minutes.) The Remap phase is based on the similarity to the bin-packing problem and makes use of a First-Fit heuristic. The strength of the paper lies in the solution for the Forecast phase (the Load prediction problem, in our terminology). It is based on a sophisticated time series analysis, aiming to identify the principal periodic components of the load distribution based on past data. As a result, the future load can be estimated along with the distribution of the prediction error. This allows consolidation with a given upper limit on the allowed probability of server overload.

Another similar work is pMapper by Verma, Ahuja, and Neogi from IBM India and IIT Delhi [99]. Also here, the aim is to optimize the mapping of VMs to PMs with respect to energy consumption and number of migrations. Another similarity is the one-dimensional nature of the problem, considering only CPU capacity and CPU load. Beside trying to switch off PMs, the authors emphasize dynamic power consumption. Interestingly, they find that utilization does not determine power consumption and argue that this prohibits the use of global optimization techniques. Instead, local power efficiency characteristics are formulated that seem to hold in practice and can be exploited for local optimization techniques. Based on these insights, three algorithms are presented. The first one, called mPP (min Power Parity), is a variation of the FFD heuristic: it considers VMs in decreasing order of CPU load, and puts each VM into the PM with sufficient capacity that offers the best energy efficiency. The weakness of this method is that it can lead to a prohibitively large number of migrations. Hence, the second algorithm, mPPH (min Power Placement with History) enhances mPP by taking into account the starting allocation, so that unnecessary migrations can be avoided. The third algorithm, pMaP, goes one step further in decreasing the number of migrations: it uses mPPH to generate a recommended new placement, but actually performs only those migrations that improve the overall energy – migrations tradeoff. The algorithms were implemented in the framework of the pMapper system, and tested with a simulator using server utilization traces from a real data center. The authors' algorithms were compared to a non-power-aware load-balancer and a static placement approach. The results show that at high levels of utilization, the difference between the algorithms' results is not so significant, but at lower utilization, the proposed algorithms perform significantly better, with pMaP being the best. Finally, it is important to note that the paper contains several other aspects beyond the algorithmic part, for example, the pMapper architecture and practical experience about (deficiencies of) the performance isolation provided by virtualization.

A closely related paper, also from the IBM India Research Lab, investigates the opportunities for static placement in more detail. The work of Verma, Dasgupta, Nayak, De, and Kothari starts with a very detailed empirical assessment of server traces from a real data center [101]. Among other findings, they establish that VMs' actual resource requirements are most of the time less than half of the maximum value and that there is significant correlation between the load of some pairs of VMs – especially those that belong to the same application. The authors then use these insights for designing optimization algorithms with the aim of minimizing energy consumption and the number of PM overloads, taking into account the single dimension of CPU load. The authors argue that for static (long-term) placement, the correlation between the loads of VMs, especially between their peak loads, is key: VMs with correlating peak load should not be placed on the same PM. They propose two new algorithms. CBP (Correlation Based Placement) is an extension of pMapper's placement algorithm, using some given percentile of the load distribution of a VM (e.g., the size at 90% of the cumulative distribution function) as its size, and avoiding the co-location of VMs with a correlation of their loads higher than a given limit. The other algorithm (Peak Clustering Based Placement – PCP) is completely new and works by clustering the VMs based on correlation between their peak loads and distributing VMs of the same cluster evenly among the PMs. The algorithms were evaluated using simulation, based on server traces from the data center where the first experiments were carried out. A

comparison with pMapper's placement algorithm (which is optimized for dynamic placement) shows the superiority of the newly proposed algorithms for static placement, with PCP performing best in most cases. Finally, the authors show how the two new algorithms can be tuned, as they are quite sensitive to the used cutoff parameters and the training period.

Another paper that also starts with empirical investigations is the work of Srikantaiah (Pennsylvania State University) and Kansal and Zhao (Microsoft Research). Their focus is on the minimization of energy consumption by means of consolidation, subject to performance constraints [89]. In contrast to the works described above, they consider two kinds of resources: CPU and disk. The main finding of the paper is the observation that consolidation impacts performance and energy consumption in a highly non-trivial manner. Up to some point, increasing the utilization leads to higher energy efficiency, as expected. However, at some point, some resource of the PM saturates, so that further increase in the utilization leads to performance degradation; since jobs take longer to complete, the energy consumption per job increases. As a result, energy consumption per job is a U-shaped function of utilization, yielding an optimal level of utilization. The authors propose to aim for this optimal utilization, which should be determined in an offline profiling phase. Afterwards, a two-dimensional packing heuristic is used, where the bin sizes correspond to the optimal utilization of the PMs. The heuristic is a variation of Worst-Fit, aiming at maximizing the remaining free capacity of PMs. This heuristic can be used both for accommodating new VMs and for optimizing the current placement of the VMs.

A quite different problem formulation was addressed by Meng, Pappas, and Zhang from IBM Research: they also consider the Single-DC problem, but with the aim of minimizing network communication costs [70]. The PMs' resources are not considered in details, but it is assumed that some capacity planning approach has been used to define a number of *slots* on each PM, and the task is to map the VMs to the slots, under the assumption that each VM fits into any slot. As input, the communication intensity is given for all pairs of VMs as well as the communication cost for all pairs of slots. The VM-to-slot mapping should minimize the resulting total communication cost. The authors classify this problem as a quadratic assignment problem, and prove its NP-hardness by reduction from Balanced Minimum k-Cut. They propose a multi-level clustering algorithm: it clusters the VMs based on communication intensity, it clusters also the slots based on communication costs, maps VM clusters to slot clusters, and then calls itself recursively for each of the generated VM cluster – slot cluster pairs. The next part of the paper is quite uncommon: for two special communication matrices and four network topologies, the authors try to determine the optimal cost (or, if this is not successful, a lower bound) and the expected cost of random placement, in order to assess the optimization opportunities. The authors also evaluate the practical performance of their proposed algorithm and compare it with two general-purpose quadratic assignment heuristics, using a combination of real server traces and synthetic additions. The results show that the proposed algorithm finds slightly better results with significantly shorter running time than the other heuristics.

Communication costs also play a vital role in the DAG scheduling approaches that are common for Multi-IaaS problem formulations. A representative example is the work of Pandey, Wu, Guru, and Buyya from the University of Melbourne [76]. Here, the aim is to map the tasks of a scientific workflow on cloud resources. For each task and each compute resource, it is given how long it would take and how much it would cost to execute the task on the resource. Dependencies between the tasks are given in the form of a DAG. For each edge of the DAG, the amount of transferred data is given; similarly, for each pair of compute resources, the cost of communication between them is given. For a mapping of tasks to resources, the total cost of a resource is defined as the execution cost of the task on this resource plus the sum of the data access costs along the incident edges; the objective is to minimize the maximum cost of a resource. For this optimization problem, the authors propose the use of Particle Swarm Optimization (PSO), a popular metaheuristic. Each particle encodes a task-resource mapping. The optimization is carried out in an online manner: first, the source tasks of the DAG are allocated; when some tasks have finished executing, the allocation of the tasks that are ready to be executed is again optimized using PSO and so on. The algorithm was evaluated using simulation on a rather small problem instance (3 compute resources, 5 tasks) and compared to an algorithm that always selects the fastest but most expensive resource. Unsurprisingly, the solution found by the proposed algorithm incurs lower costs.

Another approach to the Multi-IaaS problem is presented by Tordsson (Umea University) and Montero, Moreno-Vozmediano, and Llorente (Universidad Complutense de Madrid) [96]. Here, we can select from a list of possible VM types, where each VM type is associated with a capacity. There are multiple cloud

providers, and for each pair of VM type and cloud provider, the hourly rental fee is given. The aim is to select a set of altogether n VMs from the cloud providers, such that the total price is below a given limit and the total capacity is maximal. This optimization problem is formulated as an integer program and solved using CPLEX, a commercial off-the-shelf solver. The authors also show how some further constraints can be incorporated, e.g., the number of VMs of a given type can be constrained. The algorithm is evaluated using three real cloud providers and four VM types, allocating a total of 16 VMs to run a distributed benchmark application. Experimenting with different cost limits, an interesting observation is that in most cases the optimal allocation involves more than one cloud provider, highlighting the practicality of such a multi-cloud setup. Beyond the algorithmic part, the authors also discuss some other aspects, like the role of a cloud broker in ensuring interoperability between different cloud providers.

7 Areas in need of further research

Despite the large amount of work already presented in the literature, there are still several aspects that, in our opinion, have not yet been addressed satisfactorily. This is true both for problem formulations and for algorithms. Moreover, the state of the art concerning the evaluation of algorithms also needs improvement. We elaborate on these topics in the next subsections.

7.1 Problem formulations

We see the following issues as the most important deficits in the prevalent problem formulations.

- **Hybrid cloud**. As already mentioned, most works address either the Single-DC or the Multi-IaaS problem. Very few works address hybrid clouds, and even those have usually a strong bias in the modeling either towards the Single-DC or the Multi-IaaS subproblem, modeling the other parts only rudimentarily. Yet, hybrid clouds play an increasingly important role in practice [13, 73]. Especially in enterprise environments, hybrid clouds are becoming the standard, and so enterprise IT executives face decisions every day that relate to both in-house and cloud resources [23]. Hence, in the future, we expect to see more research about genuine hybrid cloud problem formulations.

- **Task–VM–PM mapping**. In the Single-DC problem, the usual formulation is about mapping VMs to PMs. In the Multi-IaaS problem, it is more common to investigate the mapping of tasks to VMs. However, these are just two sides of the same coin: users actually want to get their tasks mapped to PMs, and VMs are just a tool that is used to enable this mapping in a safe and efficient way. This becomes especially clear in a hybrid cloud setting, where the users' tasks are either wrapped into VMs assigned to local PMs or they are directly mapped to eCPs' VMs. Hence, in the future, we expect to see a converged model of the trilateral task–VM–PM assignment.

- **Co-optimization**. VM placement is just one level where power consumption is optimized. But power consumption optimization techniques are also implemented on the server level (e.g., Dynamic Voltage and Frequency Scaling), the level of individual components (e.g., switching unused cores, cache ways, memory banks, disks etc. to a low-energy state), and in network equipment (routers, switches), making up altogether a very complex system. In particular, it is not clear how these different optimization techniques interact, possibly interfere with each other [80]. It is not clear whether the optimal decision in the VM placement problem, if it does not account for the other optimization levels, is indeed the best choice for the overall system's power consumption. More research is needed to better understand these interactions.

- **Multi-core CPUs**. The existing problem formulations in the literature either do not model multi-core CPUs at all, or model them in a very simplistic way. This compromises the practical applicability of such approaches, because multi-core CPUs are now omnipresent.

- **Communication**. Data transfer among VMs is another aspect severely missing from many existing problem formulations, especially in the case of the Single-DC problem, although it can impact overall system performance substantially. In the literature about the Multi-IaaS problem, communication

among tasks is more frequently taken into account; however, almost exclusively coupled with the assumption that all dependencies are given in the form of a DAG. In practice though, there are often cyclic communication scenarios (e.g., two applications regularly exchanging information in both directions), and in many cases, the communication paths are not static but change at runtime depending on real-time information. Also, the modeling of workflows as DAGs usually assumes finish-to-start dependencies between adjacent tasks, but in practice, the second task can usually start its execution once some partial results of the first task are available. For these reasons, although the DAG scheduling approach is tempting because of its theoretical clarity, its applicability is limited to some narrow domains. Further research is needed on more widely applicable models of communication among VMs.

- **Co-location interference**. When deciding to place a set of VMs on a PM, many works only check that the total size of the VMs does not exceed the PM's capacity. However, in practice, there are also other attributes of the VMs that influence how suitable they are for co-location. One factor to consider is correlation: how likely it is that the resource demand of several of the VMs will increase at the same time [101]. Another factor is the "noisy neighbor" effect: since current virtualization technologies do not offer complete performance isolation of the co-located VMs, if one of the VMs uses a resource excessively, this may degrade the performance of the others [56]. Up to now, very few works addressed these issues.

7.2 Algorithms

The most important deficiency in terms of algorithms is that mostly heuristic algorithms (in most cases, quite simple heuristics) have been proposed, without any performance guarantee. This is problematic because even if they perform well in controlled experiments, they may yield poor results in real settings, especially for large and highly constrained problem instances.

- **Exact algorithms**. Since the VM placement problem contains the bin-packing problem as special case, which is NP-hard in the strong sense [69], there is no hope for an exact algorithm with polynomial or even pseudo-polynomial runtime. Nevertheless, there is still much that could be done in the context of exact algorithms, e.g., efficiently solvable special cases, fixed-parameter tractability, randomized algorithms with limited error probability, algorithms with low typical-case complexity [67]. Those authors that did experiment with exact solutions, usually used *off-the-shelf solvers* for different classes of mathematical programming; the fact that those solvers took a long time to solve even mid-sized problem instances does not mean that it is not possible to come up with better exact algorithms, tailored specifically to the given problem.

- **Approximation algorithms**. Another logical possibility that is largely unexplored as yet would be to use approximation algorithms, i.e., polynomial-time algorithms that are guaranteed to deliver a result with cost at most constant times the optimum. Since there are good approximation algorithms for the bin-packing problem, this may suggest that similar results could also be achieved for the VM placement problem.

- **Coping with uncertainty**. Most algorithms assume that all parameters of the problem are fixed and precisely known constants. (Even the approaches that attack the Load prediction problem assume that parameters other than the VMs' load are fixed and precisely known.) However, in reality, cloud data centers are very complex and highly dynamic systems, so that a real cloud management system must cope with estimation errors (e.g., PMs' background load is not constant, so that the estimate of a PM's available capacity may turn out be incorrect) and unforeseen events (e.g., a PM may be damaged or may need to be restarted because of an urgent operating system patch). The algorithms presented so far in the literature are usually not robust enough to handle such situations.

7.3 Evaluation of algorithms

Beside the problem formulations and the proposed algorithms, we feel a need for improvement also in the way algorithms for VM allocation are usually evaluated.

- **Analytic evaluation**. Most papers in the literature completely lack an analytic evaluation of the proposed algorithms. As a minimum, an estimation of the asymptotic worst-case runtime and memory consumption of the algorithms should be given. If mathematically feasible, an estimation of the asymptotic *average-case behavior* of the algorithms (using some appropriate probability distribution of the input parameters) would be even more interesting. Alternatively, an analysis of some more easily handled special cases would also contribute to a better understanding of and thus to an increased confidence in the proposed algorithms.
- **Empirical evaluation**. In absence of a detailed analytic evaluation, the empirical evaluation of the proposed algorithms is very important. Ideally, each new paper should show the advantages of the proposed method by means of a systematic comparison to previously suggested methods on a large number of different, practically relevant benchmark instances. Unfortunately, this is hardly ever done. One problem is that there are no widely accepted benchmarks for the VM placement problem (and its special cases), another issue is the co-existence of many different problem formulations, making meaningful comparisons difficult. But independently from these issues, researchers often compare their approaches to trivial algorithms or to algorithms that do not take into account some important characteristic of the problem, compare different versions of their own algorithm to each other, or do not do any comparison at all. As a result, we have currently no way to tell which of the proposed algorithms works best. The community will need to develop more rigor concerning the empirical evaluation of algorithms in order to better support the future development of the field.

8 Conclusions

We presented a survey of the state of the art in the VM allocation problem. Because of the large number of papers in this field, we could not describe all of them, but we tried to show a representative selection of the most important works. As we have seen, most papers deal with either the Single-DC or the Multi-IaaS problem, but also within those two big clusters, there are significant differences between the problem formulations used in each paper. Currently, the literature on these two subproblems is mostly disjoint, with only few works addressing a combination of the two. However, we argued that in order to capture hybrid cloud scenarios, a convergence of these two fields will be necessary in the future.

Given the diversity of the available approaches to VM placement, a natural question that arises is: which method is best, or, more realistically, when to use which method. Unfortunately, the heterogeneity of the considered problem formulations and the lack of meaningful algorithm comparison studies make it very hard to answer these questions. We see here definitely the need for future work comparing the real-world performance of algorithms under different scenarios. Also, a regular competition would be very helpful for the community, similarly to competitions of other fields, like the Competition on Software Verification (`http://sv-comp.sosy-lab.org/`).

For now, we can make recommendations mainly based on problem formulations. That is, in order to find out which approaches may be most suitable in a given situation, one should first determine if the Single-DC, the Multi-IaaS, or some of the other variants apply. Then, the main characteristics should be identified according to Tables 1 or 2. For example, in the case of communication-intensive workloads, one should consult the approaches that take inter-VM communication costs into account; likewise, if there are stringent SLAs on response time, then one should focus on approaches that support such user-level SLOs etc. This way, the search can be narrowed down to a small number of works that need to be evaluated in detail.

We hope that our survey will help practitioners select the most appropriate existing works and that it will also contribute to the maturation of this important and challenging field by demonstrating both the previous achievements and the areas for future research.

Acknowledgements

This work was partially supported by the Hungarian Scientific Research Fund (Grant Nr. OTKA 108947) and the János Bolyai Research Scholarship of the Hungarian Academy of Sciences.

References

[1] Mansoor Alicherry and T.V. Lakshman. Network aware resource allocation in distributed clouds. In *Proceedings of IEEE Infocom*, pages 963–971, 2012.

[2] Mansoor Alicherry and T.V. Lakshman. Optimizing data access latencies in cloud systems by intelligent virtual machine placement. In *Proceedings of IEEE Infocom*, pages 647–655, 2013.

[3] Amid Khatibi Bardsiri and Seyyed Mohsen Hashemi. A review of workflow scheduling in cloud computing environment. *International Journal of Computer Science and Management Research*, 1(3):348–351, 2012.

[4] Luiz André Barroso, Jimmy Clidaras, and Urs Hölzle. *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*. Morgan & Claypool, 2nd edition, 2013.

[5] Daniel M. Batista, Nelson L. S. da Fonseca, and Flavio K. Miyazawa. A set of schedulers for grid networks. In *Proceedings of the 2007 ACM Symposium on Applied Computing (SAC'07)*, pages 209–213, 2007.

[6] Anton Beloglazov, Jemal Abawajy, and Rajkumar Buyya. Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing. *Future Generation Computer Systems*, 28:755–768, 2012.

[7] Anton Beloglazov and Rajkumar Buyya. Energy efficient allocation of virtual machines in cloud data centers. In *10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing*, pages 577–578, 2010.

[8] Anton Beloglazov and Rajkumar Buyya. Energy efficient resource management in virtualized cloud data centers. In *10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing*, pages 826–831, 2010.

[9] Anton Beloglazov and Rajkumar Buyya. 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*, 24(13):1397–1420, 2012.

[10] Anton Beloglazov and Rajkumar Buyya. 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*, 24(7):1366–1379, 2013.

[11] Ofer Biran, Antonio Corradi, Mario Fanelli, Luca Foschini, Alexander Nus, Danny Raz, and Ezra Silvera. A stable network-aware VM placement for cloud systems. In *Proceedings of the 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (Ccgrid 2012)*, pages 498–506. IEEE Computer Society, 2012.

[12] Luiz F. Bittencourt and Edmundo R.M. Madeira. HCOC: a cost optimization algorithm for workflow scheduling in hybrid clouds. *Journal of Internet Services and Applications*, 2(3):207–227, 2011.

[13] Luiz F. Bittencourt, Edmundo R.M. Madeira, and Nelson L.S. da Fonseca. Scheduling in hybrid clouds. *IEEE Communications Magazine*, 50(9):42–47, 2012.

[14] Luiz F. Bittencourt, Rizos Sakellariou, and Edmundo R. M. Madeira. Using relative costs in workflow scheduling to cope with input data uncertainty. In *Proceedings of the 10th International Workshop on Middleware for Grids, Clouds and e-Science*, 2012.

[15] Norman Bobroff, Andrzej Kochut, and Kirk Beaty. Dynamic placement of virtual machines for managing SLA violations. In *10th IFIP/IEEE International Symposium on Integrated Network Management*, pages 119–128, 2007.

[16] Ruben Van den Bossche, Kurt Vanmechelen, and Jan Broeckhove. Cost-optimal scheduling in hybrid IaaS clouds for deadline constrained workloads. In *IEEE 3rd International Conference on Cloud Computing*, pages 228–235, 2010.

[17] David Breitgand and Amir Epstein. SLA-aware placement of multi-virtual machine elastic services in compute clouds. In *12th IFIP/IEEE International Symposium on Integrated Network Management*, pages 161–168, 2011.

[18] David Breitgand and Amir Epstein. Improving consolidation of virtual machines with risk-aware bandwidth oversubscription in compute clouds. In *Proceedings of IEEE Infocom 2012*, pages 2861–2865, 2012.

[19] Rajkumar Buyya, Anton Beloglazov, and Jemal Abawajy. Energy-efficient management of data center resources for cloud computing: a vision, architectural elements, and open challenges. In *Proceedings of the 2010 International Conference on Parallel and Distributed Processing Techniques and Applications*, pages 6–17. CSREA Press, 2010.

[20] Rajkumar Buyya, Chee Shin Yeo, Srikumar Venugopal, James Broberg, and Ivona Brandic. Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility. *Future Generation Computer Systems*, 25(6):599–616, 2009.

[21] Rodrigo N. Calheiros, Rajiv Ranjan, Anton Beloglazov, César A. F. De Rose, and Rajkumar Buyya. CloudSim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms. *Software: Practice and Experience*, 41(1):23–50, 2011.

[22] David Candeia, Ricardo Araújo, Raquel Lopes, and Francisco Brasileiro. Investigating business-driven cloudburst schedulers for e-science bag-of-tasks applications. In *2nd IEEE International Conference on Cloud Computing Technology and Science*, pages 343–350, 2010.

[23] Capgemini. Simply. business cloud. `http://www.capgemini.com/resource-file-access/resource/pdf/simply._business_cloud_where_business_meets_cloud.pdf` (last accessed: February 10, 2015), 2013.

[24] Emiliano Casalicchio, Daniel A. Menasc, and Arwa Aldhalaan. Autonomic resource provisioning in cloud systems with availability goals. In *Proceedings of the 2013 ACM Cloud and Autonomic Computing Conference*, 2013.

[25] Emmanuel Cecchet, Anupam Chanda, Sameh Elnikety, Julie Marguerite, and Willy Zwaenepoel. Performance comparison of middleware architectures for generating dynamic web content. In *Proceedings of the ACM/IFIP/USENIX 2003 International Conference on Middleware*, pages 242–261, 2003.

[26] Sivadon Chaisiri, Bu-Sung Lee, and Dusit Niyato. Optimal virtual machine placement across multiple cloud providers. In *IEEE Asia-Pacific Services Computing Conference, (APSCC 2009)*, pages 103–110, 2009.

[27] Jeffrey S. Chase, Darrell C. Anderson, Prachi N. Thakar, and Amin M. Vahdat. Managing energy and server resources in hosting centers. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles*, pages 103–116, 2001.

[28] Yuan Chen, Subu Iyer, Xue Liu, Dejan Milojicic, and Akhil Sahai. Translating service level objectives to lower level policies for multi-tier services. *Cluster Computing*, 11:299–311, 2008.

[29] Ludmila Cherkasova, Diwaker Gupta, and Amin Vahdat. When virtual is harder than real: Resource allocation challenges in virtual machine based IT environments. Technical report, HP Laboratories Palo Alto, 2007.

[30] Navraj Chohan, Claris Castillo, Mike Spreitzer, Malgorzata Steinder, Asser Tantawi, and Chandra Krintz. See spot run: using spot instances for mapreduce workflows. In *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing (HotCloud'10)*, 2011.

[31] Rajarshi Das, Jeffrey O. Kephart, Charles Lefurgy, Gerald Tesauro, David W. Levine, and Hoi Chan. Autonomic multi-agent management of power and performance in data centers. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems: Industrial Track*, pages 107–114, 2008.

[32] Digital Power Group. *The cloud begins with coal – Big data, big networks, big infrastructure, and big power*. 2013.

[33] György Dósa. The tight bound of first fit decreasing bin-packing algorithm is $FFD(I) \leq 11/9OPT(I)+6/9$. In *Combinatorics, Algorithms, Probabilistic and Experimental Methodologies*, pages 1–11. Springer, 2007.

[34] György Dósa and Jiří Sgall. First fit bin packing: A tight analysis. In *30th Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 538–549, 2013.

[35] György Dósa and Jiří Sgall. Optimal analysis of best fit bin packing. In *41st International Colloquium on Automata, Languages and Programming (ICALP)*, pages 429–441, 2014.

[36] Dror G. Feitelson, Dan Tsafrir, and David Krakov. Experience with using the parallel workloads archive. *Journal of Parallel and Distributed Computing*, 74(10):2967–2982, 2014.

[37] Ana Juan Ferrer, Francisco Hernndez, Johan Tordsson, Erik Elmroth, Ahmed Ali-Eldin, Csilla Zsigri, Ral Sirvent, Jordi Guitart, Rosa M. Badia, Karim Djemame, Wolfgang Ziegler, Theo Dimitrakos, Srijith K. Nair, George Kousiouris, Kleopatra Konstanteli, Theodora Varvarigou, Benoit Hudzia, Alexander Kipp, Stefan Wesner, Marcelo Corrales, Nikolaus Forg, Tabassum Sharif, and Craig Sheridan. OPTIMIS: A holistic approach to cloud service provisioning. *Future Generation Computer Systems*, 28(1):66–77, 2012.

[38] Yongqiang Gao, Haibing Guan, Zhengwei Qi, Yang Hou, and Liang Liu. A multi-objective ant colony system algorithm for virtual machine placement in cloud computing. *Journal of Computer and System Sciences*, 79:1230–1242, 2013.

[39] Thiago A. L. Genez, Luiz F. Bittencourt, and Edmundo R. M. Madeira. Workflow scheduling for SaaS/PaaS cloud providers considering two SLA levels. In *Network Operations and Management Symposium (NOMS)*, pages 906–912. IEEE, 2012.

[40] Daniel Gmach, Jerry Rolia, Ludmila Cherkasova, Guillaume Belrose, Tom Turicchi, and Alfons Kemper. An integrated approach to resource pool management: Policies, efficiency and quality metrics. In *IEEE International Conference on Dependable Systems and Networks*, pages 326–335, 2008.

[41] Daniel Gmach, Jerry Rolia, Ludmila Cherkasova, and Alfons Kemper. Resource pool management: Reactive versus proactive or let's be friends. *Computer Networks*, 53(17):2905–2922, 2009.

[42] Marco Guazzone, Cosimo Anglano, and Massimo Canonico. Exploiting VM migration for the automated power and performance management of green cloud computing systems. In *First International Workshop on Energy Efficient Data Centers (E2DC 2012)*, pages 81–92. Springer, 2012.

[43] Marco Guazzone, Cosimo Anglano, and Massimo Canonico. Exploiting vm migration for the automated power and performance management of green cloud computing systems. Technical Report TR-INF-2012-04-02-UNIPMN, University of Piemonte Orientale, 2012.

[44] Brian Guenter, Navendu Jain, and Charles Williams. Managing cost, performance, and reliability tradeoffs for energy-aware server provisioning. In *Proceedings of IEEE INFOCOM*, pages 1332–1340. IEEE, 2011.

[45] Sijin He, Li Guo, Moustafa Ghanem, and Yike Guo. Improving resource utilisation in the cloud environment using multivariate probabilistic models. In *IEEE 5th International Conference on Cloud Computing*, pages 574–581, 2012.

[46] Tibor Horvath, Tarek Abdelzaher, Kevin Skadron, and Xue Liu. Dynamic voltage scaling in multi-tier web servers with end-to-end delay control. *IEEE Transactions on Computers*, 56(4):444–458, 2007.

[47] Chris Hyser, Bret McKee, Rob Gardner, and Brian J. Watson. Autonomic virtual machine placement in the data center. Technical report, HP Laboratories, 2008.

[48] Waheed Iqbal, Matthew N. Dailey, and David Carrera. SLA-driven dynamic resource management for multi-tier web applications in a cloud. In *10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing (CCGrid)*, pages 832–837, 2010.

[49] Deepal Jayasinghe, Calton Pu, Tamar Eilam, Malgorzata Steinder, Ian Whalley, and Ed Snible. Improving performance and availability of services hosted on iaas clouds with structural constraint-aware virtual machine placement. In *IEEE International Conference on Services Computing (SCC)*, pages 72–79, 2011.

[50] Joe Wenjie Jiang, Tian Lan, Sangtae Ha, Minghua Chen, and Mung Chiang. Joint VM placement and routing for data center traffic engineering. In *Proceedings of IEEE Infocom 2012*, pages 2876–2880, 2012.

[51] Gueyoung Jung, Matti A. Hiltunen, Kaustubh R. Joshi, Richard D. Schlichting, and Calton Pu. Mistral: Dynamically managing power, performance, and adaptation cost in cloud infrastructures. In *IEEE 30th International Conference on Distributed Computing Systems (ICDCS)*, pages 62–73, 2010.

[52] Gabor Kecskemeti, Gabor Terstyanszky, Peter Kacsuk, and Zsolt Nemeth. An approach for virtual appliance distribution for service deployment. *Future Generation Computer Systems*, 27(3):280–289, 2011.

[53] Matthias Keller and Holger Karl. Response time-optimized distributed cloud resource allocation. In *Proceedings of the 2014 ACM SIGCOMM workshop on Distributed cloud computing*, pages 47–52, 2014.

[54] Gunjan Khanna, Kirk Beaty, Gautam Kar, and Andrzej Kochut. Application performance management in virtualized server environments. In *10th IEEE/IFIP Network Operations and Management Symposium*, pages 373–381, 2006.

[55] Atefeh Khosravi, Saurabh Kumar Garg, and Rajkumar Buyya. Energy and carbon-efficient placement of virtual machines in distributed cloud data centers. In *Euro-Par 2013 Parallel Processing*, pages 317–328. Springer, 2013.

[56] Shin-gyu Kim, Hyeonsang Eom, and Heon Y. Yeom. Virtual machine consolidation based on interference modeling. *The Journal of Supercomputing*, 66(3):1489–1506, 2013.

[57] Ricardo Koller, Akshat Verma, and Anindya Neogi. WattApp: an application aware power meter for shared data centers. In *Proceedings of the 7th international conference on Autonomic computing*, pages 31–40, 2010.

[58] Ricardo Koller, Akshat Verma, and Raju Rangaswami. Estimating application cache requirements for provisioning caches in virtualized systems. In *IEEE 19th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)*, pages 55–62, 2011.

[59] Madhukar Korupolu, Aameek Singh, and Bhuvan Bamba. Coupled placement in modern data centers. In *IEEE International Symposium on Parallel and Distributed Processing (IPDPS 2009)*, pages 1–12, 2009.

[60] Daniel Guimaraes do Lago, Edmundo R. M. Madeira, and Luiz Fernando Bittencourt. Power-aware virtual machine scheduling on clouds using active cooling control and DVFS. In *Proceedings of the 9th International Workshop on Middleware for Grids, Clouds and e-Science*, 2011.

[61] Ulrich Lampe, Melanie Siebenhaar, Ronny Hans, Dieter Schuller, and Ralf Steinmetz. Let the clouds compute: cost-efficient workload distribution in infrastructure clouds. In *Proceedings of the 9th International Conference on Economics of Grids, Clouds, Systems, and Services (GECON 2012)*, pages 91–101. Springer, 2012.

[62] Wubin Li, Petter Svärd, Johan Tordsson, and Erik Elmroth. A general approach to service deployment in cloud environments. In *2nd International Conference on Cloud and Green Computing (CGC)*, pages 17–24, 2012.

[63] Wubin Li, Petter Svärd, Johan Tordsson, and Erik Elmroth. Cost-optimal cloud service placement under dynamic pricing schemes. In *Proceedings of the 6th IEEE/ACM International Conference on Utility and Cloud Computing*, pages 187–194, 2013.

[64] Wubin Li, Johan Tordsson, and Erik Elmroth. Modeling for dynamic cloud scheduling via migration of virtual machines. In *Proceedings of the 3rd IEEE International Conference on Cloud Computing Technology and Science*, pages 163–171, 2011.

[65] Wubin Li, Johan Tordsson, and Erik Elmroth. Virtual machine placement for predictable and time-constrained peak loads. In *Proceedings of the 8th International Conference on Economics of Grids, Clouds, Systems, and Services (GECON 2011)*, pages 120–134. Springer, 2011.

[66] Liang Liu, Hao Wang, Xue Liu, Xing Jin, Wen Bo He, Qing Bo Wang, and Ying Chen. GreenCloud: A new architecture for green data center. In *Proceedings of the 6th International Conference on Autonomic Computing and Communications*, pages 29–38, 2009.

[67] Zoltán Ádám Mann. *Optimization in computer engineering – Theory and applications.* Scientific Research Publishing, 2011.

[68] Zoltán Ádám Mann. Rigorous results on the effectiveness of some heuristics for the consolidation of virtual machines in a cloud data center. `http://www.cs.bme.hu/~mann/publications/Preprints/Mann_Rigorous_VM_Consolidation.pdf`, 2015.

[69] Silvano Martello and Paolo Toth. *Knapsack problems: algorithms and computer implementations.* John Wiley & Sons, 1990.

[70] Xiaoqiao Meng, Vasileios Pappas, and Li Zhang. Improving the scalability of data center networks with traffic-aware virtual machine placement. In *IEEE INFOCOM 2010 proceedings*, pages 1–9, 2010.

[71] Kevin Mills, James Filliben, and Christopher Dabrowski. Comparing vm-placement algorithms for on-demand clouds. In *Proceedings of the 3rd IEEE International Conference on Cloud Computing Technology and Science*, pages 91–98, 2011.

[72] Mayank Mishra and Anirudha Sahoo. On theory of vm placement: Anomalies in existing methodologies and their mitigation using a novel vector based approach. In *IEEE International Conference on Cloud Computing*, pages 275–282, 2011.

[73] Rafael Moreno-Vozmediano, Ruben S. Montero, and Ignacio M. Llorente. Multicloud deployment of computing clusters for loosely coupled MTC applications. *IEEE Transactions on Parallel and Distributed Systems*, 22(6):924–930, 2011.

[74] Ripal Nathuji and Karsten Schwan. VirtualPower: coordinated power management in virtualized enterprise systems. In *Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles (SOSP)*, pages 265–278, 2007.

[75] Daniel de Oliveira, Kary A. C. S. Ocana, Fernanda Baiao, and Marta Mattoso. A provenance-based adaptive scheduling heuristic for parallel scientific workflows in clouds. *Journal of Grid Computing*, 10:521–552, 2012.

[76] Suraj Pandey, Linlin Wu, Siddeswara Mayura Guru, and Rajkumar Buyya. A particle swarm optimization-based heuristic for scheduling workflow applications in cloud computing environments. In *24th IEEE International Conference on Advanced Information Networking and Applications (AINA)*, pages 400–407. IEEE, 2010.

[77] Christos H. Papadimitriou and Mihalis Yannakakis. Multiobjective query optimization. In *Proceedings of the 20th ACM Symposium on Principles of Database Systems*, pages 52–59. ACM, 2001.

[78] Michael L. Pinedo. *Scheduling: Theory, Algorithms, and Systems*. Springer, 3rd edition, 2008.

[79] Eduardo Pinheiro, Ricardo Bianchini, Enrique V. Carrera, and Taliver Heath. Load balancing and unbalancing for power and performance in cluster-based systems. In *Workshop on Compilers and Operating Systems for Low Power*, pages 182–195, 2001.

[80] Ramya Raghavendra, Parthasarathy Ranganathan, Vanish Talwar, Zhikui Wang, and Xiaoyun Zhu. No "power" struggles: Coordinated multi-level power management for the data center. In *Proceedings of the 13th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 48–59, 2008.

[81] Charles Reiss, Alexey Tumanov, Gregory R. Ganger, Randy H. Katz, and Michael A. Kozuch. Heterogeneity and dynamicity of clouds at scale: Google trace analysis. In *ACM Symposium on Cloud Computing (SoCC)*, 2012.

[82] Bruno Cesar Ribas, Rubens Massayuki Suguimoto, Razer A. N. R. Montano, Fabiano Silva, Luis de Bona, and Marcos A. Castilho. On modelling virtual machine consolidation to pseudo-boolean constraints. In *13th Ibero-American Conference on AI*, pages 361–370, 2012.

[83] Suzanne Rivoire, Parthasarathy Ranganathan, and Christos Kozyrakis. A comparison of high-level full-system power models. In *Workshop on Power Aware Computing and Systems (HotPower '08)*, 2008.

[84] Benny Rochwerger, David Breitgand, Eliezer Levy, Alex Galis, Kenneth Nagin, Ignacio Llorente, Ruben Montero, Yaron Wolfsthal, Erik Elmroth, Juan Caceres, M. Ben-Yehuda, Wolfgang Emmerich, and Fermin Galn. The Reservoir model and architecture for open federated cloud computing. *IBM Journal of Research and Development*, 53(4):1–11, 2009.

[85] Ivan Rodero, Hariharasudhan Viswanathan, Eun Kyung Lee, Marc Gamell, Dario Pompili, and Manish Parashar. Energy-efficient thermal-aware autonomic management of virtualized HPC cloud infrastructure. *Journal of Grid Computing*, 10(3):447–473, 2012.

[86] Mohsen Amini Salehi, P. Radha Krishna, Krishnamurty Sai Deepak, and Rajkumar Buyya. Preemption-aware energy management in virtualized data centers. In *5th International Conference on Cloud Computing*, pages 844–851. IEEE, 2012.

[87] Mina Sedaghat, Francisco Hernandez-Rodriguez, and Erik Elmroth. A virtual machine re-packing approach to the horizontal vs. vertical elasticity trade-off for cloud autoscaling. In *Proceedings of the 2013 ACM Cloud and Autonomic Computing Conference*, 2013. article nr. 6.

[88] Lei Shi, John Furlong, and Runxin Wang. Empirical evaluation of vector bin packing algorithms for energy efficient data centers. In *IEEE Symposium on Computers and Communications (ISCC)*, pages 9–15, 2013.

[89] Shekhar Srikantaiah, Aman Kansal, and Feng Zhao. Energy aware consolidation for cloud computing. *Cluster Computing*, 12:1–15, 2009.

[90] Steve Strauch, Oliver Kopp, Frank Leymann, and Tobias Unger. A taxonomy for cloud data hosting solutions. In *Proceedings of the IEEE International Conference on Cloud and Green Computing (CGC 2011)*, pages 577–584, 2011.

[91] Anja Strunk. Costs of virtual machine live migration: A survey. In *8th IEEE World Congress on Services*, pages 323–329, 2012.

[92] Rick Sturm and Wayne Morris. *Foundations of Service Level Management*. SAMS Publishing, 2000.

[93] Petter Svärd, Wubin Li, Eddie Wadbro, Johan Tordsson, and Erik Elmroth. Continuous datacenter consolidation. Technical report, Umea University, 2014.

[94] Vanish Talwar, Dejan Milojicic, Qinyi Wu, Calton Pu, Wenchang Yan, and Gueyoung Jung. Approaches for service deployment. *IEEE Internet Computing*, 9(2):70–80, 2005.

[95] Luis Tomás and Johan Tordsson. An autonomic approach to risk-aware data center overbooking. *IEEE Transactions on Cloud Computing*, 2(3):292–305, 2014.

[96] Johan Tordsson, Rubén S. Montero, Rafael Moreno-Vozmediano, and Ignacio M. Llorente. Cloud brokering mechanisms for optimized placement of virtual machines across multiple providers. *Future Generation Computer Systems*, 28(2):358–367, 2012.

[97] Efthymia Tsamoura, Anastasios Gounaris, and Kostas Tsichlas. Multi-objective optimization of data flows in a multi-cloud environment. In *Proceedings of the Second Workshop on Data Analytics in the Cloud*, pages 6–10, 2013.

[98] Bhuvan Urgaonkar, Prashant Shenoy, Abhishek Chandra, and Pawan Goyal. Dynamic provisioning of multi-tier internet applications. In *Second International Conference on Autonomic Computing*, pages 217–228, 2005.

[99] Akshat Verma, Puneet Ahuja, and Anindya Neogi. pMapper: power and migration cost aware application placement in virtualized systems. In *Middleware 2008*, pages 243–264, 2008.

[100] Akshat Verma, Puneet Ahuja, and Anindya Neogi. Power-aware dynamic placement of HPC applications. In *Proceedings of the 22nd Annual International Conference on Supercomputing*, pages 175–184, 2008.

[101] Akshat Verma, Gargi Dasgupta, Tapan Kumar Nayak, Pradipta De, and Ravi Kothari. Server workload analysis for power minimization using consolidation. In *Proceedings of the 2009 USENIX Annual Technical Conference*, pages 355–368, 2009.

[102] Akshat Verma, Gautam Kumar, and Ricardo Koller. The cost of reconfiguration in a cloud. In *Proceedings of the 11th International Middleware Conference*, pages 11–16, 2010.

[103] Akshat Verma, Gautam Kumar, Ricardo Koller, and Aritra Sen. CosMig: Modeling the impact of reconfiguration in a cloud. In *19th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)*, pages 3–11. IEEE, 2011.

[104] David Villegas, Athanasion Antoniou, Seyed Masoud Sadjadi, and Alexandru Iosup. An analysis of provisioning and allocation policies for infrastructure-as-a-service clouds. In *12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, pages 612–619, 2012.

[105] Timothy Wood, Prashant Shenoy, Arun Venkataramani, and Mazin Yousif. Sandpiper: Black-box and gray-box resource management for virtual machines. *Computer Networks*, 53(17):2923–2938, 2009.

[106] Zhangjun Wu, Zhiwei Ni, Lichuan Gu, and Xiao Liu. A revised discrete particle swarm optimization for cloud workflow scheduling. In *International Conference on Computational Intelligence and Security*, pages 184–188, 2010.

[107] Jing Xu and Jose A. B. Fortes. Multi-objective virtual machine placement in virtualized data center environments. In *Green Computing and Communications (GreenCom), 2010 IEEE/ACM International Conference on Cyber, Physical and Social Computing (CPSCom)*, pages 179–188, 2010.

[108] Qi Zhang, Lu Cheng, and Raouf Boutaba. Cloud computing: state-of-the-art and research challenges. *Journal of Internet Services and Applications*, 1(1):7–18, 2010.

[109] Xiaoyun Zhu, Donald Young, Brian J. Watson, Zhikui Wang, Jerry Rolia, Sharad Singhal, Bret McKee, Chris Hyser, Daniel Gmach, Robert Gardner, Tom Christian, and Ludmila Cherkasova. 1000 islands: an integrated approach to resource management for virtualized data centers. *Cluster Computing*, 12(1):45–57, 2009.