

Copyright Notice

The document is provided by the contributing author(s) as a means to ensure timely dissemination of scholarly and technical work on a non-commercial basis. This is the author's version of the work. The final version can be found on the publisher's webpage.

This document is made available only for personal use and must abide to copyrights of the publisher. Permission to make digital or hard copies of part or all of these works for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage. This works may not be reposted without the explicit permission of the copyright holder.

Permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the corresponding copyright holders. It is understood that all persons copying this information will adhere to the terms and constraints invoked by each copyright holder.

IEEE papers: © IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The final publication is available at <http://ieeexplore.ieee.org>

ACM papers: © ACM. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The final publication is available at <http://dl.acm.org/>

Springer papers: © Springer. Pre-prints are provided only for personal use. The final publication is available at <link.springer.com>

Distributed Virtual Machine Consolidation: A Systematic Mapping Study

Adnan Ashraf*, Benjamin Byholm, Ivan Porres

Faculty of Natural Sciences and Technology

Åbo Akademi University, Finland

Abstract

Background: Virtual Machine (VM) consolidation is an effective technique to improve resource utilization and reduce energy footprint in cloud data centers. It can be implemented in a centralized or a distributed fashion. Distributed VM consolidation approaches are currently gaining popularity because they are often more scalable than their centralized counterparts and they avoid a single point of failure.

Objective: To present a comprehensive, unbiased overview of the state-of-the-art on distributed VM consolidation approaches.

Method: A Systematic Mapping Study (SMS) of the existing distributed VM consolidation approaches.

Results: 19 papers on distributed VM consolidation categorized in a variety of ways. The results show that the existing distributed VM consolidation approaches use four types of algorithms, optimize a number of different objectives, and are often evaluated with experiments involving simulations.

Conclusion: There is currently an increasing amount of interest on developing and evaluating novel distributed VM consolidation approaches. A number of research gaps exist where the focus of future research may be directed.

Keywords: Cloud computing, Data center, Virtual machine, Consolidation, Placement, Energy-efficiency

1 Introduction

Energy footprint of cloud data centers is a matter of great concern for cloud providers [39]. According to a recent report [51], data centers in the United States consumed an estimated 70 billion kilowatt-hours of electricity in 2014, which corresponds to 1.8% of total United States electricity consumption. High energy consumption not only translates into a high operating cost, but also leads to huge carbon emissions. The ever increasing demand for computing resources to provide highly scalable and reliable services has caused an energy crisis [30]. The high energy consumption of data centers can partly be attributed to the large-scale installations of computing and cooling infrastructures, but more importantly it is due to the inefficient use of the computing resources [5]. Production servers seldom operate near their full capacity [4]. However, even at the completely idle state, they consume a substantial proportion of their peak power [10]. Therefore, under-utilized servers are highly inefficient.

Hardware virtualization technologies allow to share a Physical Machine (PM) among multiple, performance-isolated platforms called Virtual Machines (VMs) to improve resource utilization. Further improvement in resource utilization and reduction in energy consumption can be achieved by consolidating VMs on under-utilized PMs. The basic idea of VM consolidation is to migrate and place the VMs on as few PMs as possible and then release the remaining, unused PMs for termination or for switching to a low-power mode to conserve energy. A VM consolidation approach uses live VM migration to consolidate VMs on a reduced set of PMs. VM consolidation has emerged as one of the most effective and promising techniques to reduce energy footprint of cloud data centers [5, 11].

Figure 1 presents a simple hypothetical scenario to illustrate the VM consolidation process. The first half of Figure 1 shows three PMs where each PM hosts multiple VMs and every VM uses a certain amount of the PM resources. It is assumed that due to some significant load variations, PM 2 and PM 3 have become under-utilized. The under-utilized PMs in such a scenario may continue to remain under-utilized for hours, days, or even weeks unless the existing VMs require more resources or some new VMs are placed on the under-utilized PMs. Therefore, it is difficult to provide a resource and energy efficient allocation of VMs without consolidation of VMs on the under-utilized PMs. The second half of Figure 1 shows that after migrating all VMs from PM 2 to PM 3, PM 2 can be turned-off or switched to a low-power mode.

There is currently an increasing amount of interest on developing and evaluating efficient VM consolidation approaches for cloud data centers. Over the past few years, researchers have used a multitude of ways to

*Corresponding author e-mail: adnan.ashraf@abo.fi

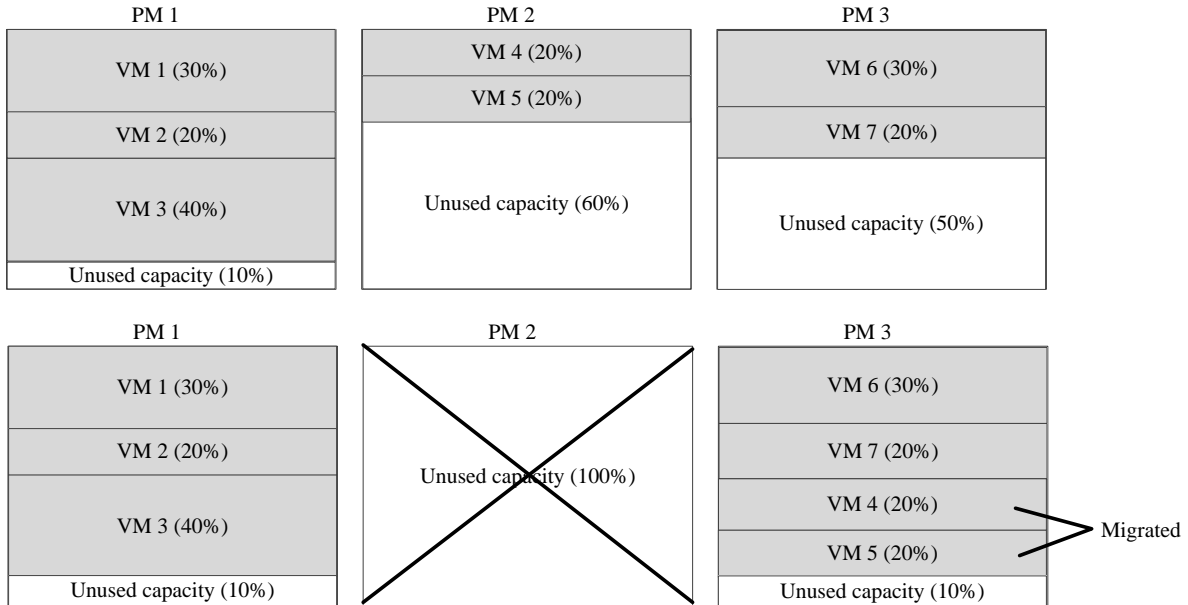


Figure 1: A simple example to illustrate the consolidation of VMs on under-utilized PMs to conserve energy

develop novel VM consolidation approaches. Some of these approaches have been recently reported in the form of nonsystematic literature reviews such as [3] and [46]. However, the drawback of these existing nonsystematic studies is that they provide a partial and possibly biased overview of the state-of-the-art on VM consolidation. For a comprehensive and unbiased coverage of the existing literature on VM consolidation, there is a need to study the existing VM consolidation approaches in a systematic way.

VM consolidation can be implemented in a centralized or a distributed fashion. Traditional VM consolidation approaches, such as [5, 12, 13, 16–18, 25], tend to be centralized. A centralized VM consolidation approach uses a centralized algorithm on a centralized architecture and does not provide support for multiple, geographically distributed data centers. The main drawbacks of centralized VM consolidation approaches include limited scalability and lack of robustness due to a single point of failure. On the other hand, a distributed or decentralized VM consolidation approach uses a distributed algorithm or a distributed architecture for PMs [11, 14] or provides support for multiple, geographically distributed data centers [35, 62]. Distributed VM consolidation is a recurring theme in recent VM consolidation approaches such as [11, 14, 36]. Distributed approaches are gaining popularity because they have benefits over centralized approaches. They are often more scalable than their centralized counterparts and they avoid a single point of failure [36, 38]. Feller et al. [14] showed that their proposed VM consolidation algorithm does not compute a solution (in a reasonable amount of time) on a centralized architecture, but finds a good solution on a distributed architecture. Lucanin and Brandic [35] reported that their VM consolidation algorithm for geographically distributed data centers finds a good solution for a large-scale problem comprising ten thousand VMs. Sedaghat et al. [49] showed that their proposed distributed VM consolidation algorithm scales to tens of thousands of PMs and VMs without compromising on the quality of the solution. Sedaghat et al. [50] reported that their proposed distributed VM consolidation algorithm finds a near-optimal solution for 100,000 PMs in a reasonable amount of time. Marzolla et al. [36] showed that their proposed distributed VM consolidation algorithm is resilient to major failures and outages involving a thousand PMs. Therefore, distributed VM consolidation approaches are more suitable for large-scale data centers involving thousands of VMs and PMs.

We present a systematic study of the existing distributed VM consolidation approaches. The objective is to present a comprehensive, unbiased overview of the state-of-the-art on distributed VM consolidation approaches. Considering the broad nature of the research objective, it was not appropriate to launch a Systematic Literature Review (SLR). Therefore, we launched a Systematic Mapping Study (SMS) [6, 42, 43]. A SMS follows the same principled process as a SLR, but: (1) it has a broader scope, (2) it uses different criteria for inclusions/exclusions and quality assessments, and (3) its data collection and synthesis tend to be more qualitative than for a SLR [61]. It is “intended to ‘map out’ the research that has been undertaken rather than to answer a detailed research question” [6].

We proceed as follows. Section 2 presents the design and schedule of our study. The results of the SMS are presented in Section 3. In Section 4, we discuss major threats to the validity of the results presented in this paper. Finally, we present our conclusions in Section 5.

2 Study Design

One of the most important differences between a nonsystematic literature review and a SMS is that a SMS follows an unbiased and repeatable process. Moreover, the process is documented as a review protocol. Therefore, we defined a review protocol for our SMS on distributed VM consolidation approaches. In this section, we present the design of our study and the review process.

2.1 Research Questions

The Research Questions (RQs) are as follows:

- RQ1: What approaches have been developed for distributed VM consolidation?
- RQ2: What kinds of algorithms are being used in the existing distributed VM consolidation approaches?
- RQ3: What objectives are being optimized in the existing distributed VM consolidation approaches?
- RQ4: How are the existing distributed VM consolidation approaches being evaluated?
- RQ5: What are the most popular publication forums for distributed VM consolidation papers and how have they changed over time?

RQ1 is the basic question for obtaining an overview of the state-of-the-art on distributed VM consolidation. RQ2 is aimed at obtaining the types of algorithms which are being used in the existing distributed VM consolidation approaches. Possible types include heuristics, metaheuristics, and machine learning algorithms. Moreover, the algorithms may also be categorized into offline and online optimization algorithms.

RQ3 concerns the objectives which are being optimized. Possible objectives include minimizing energy consumption, minimizing the number of active PMs, minimizing Service Level Agreement (SLA) violations, minimizing the number of VM migrations, minimizing cost, minimizing network traffic, maximizing performance, maximizing reliability, and minimizing resource utilization. RQ3 also deals with the number of objectives which are being optimized and how the optimization problem is formulated. Possible problem formulations include single-objective, multi-objective (two or three objectives) with an Aggregate Objective Function (AOF), pure multi-objective, and many-objective (four or more objectives) [57].

RQ4 concerns the evaluation method. The most common evaluation method in the VM consolidation literature is experiment. An experiment may involve the use of prototype implementations or simulations. Moreover, the experiment design may involve realistic, synthetic, or hybrid load patterns. Similarly, an experimental evaluation may or may not include a comparison of the results with other existing VM consolidation approaches. Finally, a comparison of the results may or may not include statistical tests to assess the statistical significance of the results.

RQ5 is a typical question for SMSs in software engineering [42]. The objective is to identify the most popular, peer-reviewed publication forums with respect to distributed VM consolidation papers. The publication forums may include journals, conferences, and workshops. In addition, the second part of RQ5 concerns the frequencies of published papers in popular forums over time to see the trends.

Based on the RQs, the Population, Intervention, Comparison, Outcomes, and Context (PICOC) [33] is presented in Table 1.

Table 1: PICOC

Aspect	Value
Population (P)	Energy-aware and cost-effective data center management techniques
Intervention (I)	Distributed VM consolidation/placement approaches/ algorithms/methods/heuristics
Comparison (C)	No comparison intervention
Outcomes (O)	An overview of the state-of-the-art on distributed VM consolidation approaches
Context (C)	Cloud data centers

2.2 Search Strategy for Primary Studies

This section presents our search strategy. It is based on the SLR and SMS guidelines described by Kitchenham and Charters [33] and Wohlin et al. [61].

2.2.1 Search Terms

Table 2 presents the most important search terms along with their alternate spellings. The search terms are primarily based on the RQs and PICOC in Section 2.1. Moreover, they are also in line with recent and prominent works on VM consolidation such as [5, 11–14, 16–18, 25, 36].

Table 2: Search terms

#	Search term	Alternate spellings
1	Consolidat*	Consolidate, consolidating, consolidation
2	Plac*	Place, placing, placement
3	Virtual machine*	Virtual machine, virtual machines
4	VM*	VM, VMs
5	Server*	Server, servers
6	Algorithm*	Algorithm, algorithms
7	Approach*	Approach, approaches
8	Method*	Method, methods
9	Heuristic*	Heuristic, heuristics
10	Cloud	None
11	Data center	Data center, datacenter, data centre, datacentre
12	Distributed	None
13	Decentralized	None

2.2.2 Search Strings

The search strings are presented in Table 3. They are formed by making appropriate combinations of the search terms presented in Table 2. We used Boolean *AND* and Boolean *OR* operators to make the search strings. The two search strings in Table 3 can also be combined into one search string by using the Boolean *OR* operator. Therefore, the papers which contain any of the two search strings were retrieved. The search strings were validated against a set of known papers [11, 14, 36].

Table 3: Search strings

#	Search string
1	(Distributed OR Decentralized) AND Consolidat* AND ("virtual machine*" OR VM* OR server*) AND (algorithm* OR approach* OR method* OR heuristic*) AND (cloud OR "data center" OR datacenter OR "data centre" OR datacentre)
2	(Distributed OR Decentralized) AND Plac* AND ("virtual machine*" OR VM*) AND (algorithm* OR approach* OR method* OR heuristic*) AND (cloud OR "data center" OR datacenter OR "data centre" OR datacentre)

2.2.3 Databases

The search strings in Table 3 were searched in the publication title, abstract, and keywords. The following digital libraries were searched: (1) Institute of Electrical and Electronics Engineers (IEEE) Xplore, (2) Association for Computing Machinery (ACM) Digital Library, (3) ScienceDirect, and (4) SpringerLink. The search strings were customized for each digital library. Moreover, since using multiple digital libraries creates duplicates, the search results were analyzed to identify and remove the duplicates.

2.3 Study Selection Criteria

This section presents our inclusion and exclusion criteria for primary studies.

2.3.1 Inclusion Criteria

The inclusion criteria for primary studies are as follows:

- Distributed or decentralized VM consolidation approach or algorithm or method or heuristic *AND*
- For cloud data center(s) *AND*
- Written in English *AND*
- Published in a peer-reviewed journal, conference, or workshop of computer science, computer engineering, or software engineering

In addition, if several papers presented the same VM consolidation approach, only the most recent was included.

2.3.2 Exclusion Criteria

The exclusion criteria are the inverse of the inclusion criteria. If several papers presented the same VM consolidation approach, all except the most recent were excluded.

2.4 Study Selection Procedure

The study selection procedure was applied on the search results to remove false positives. It comprises two phases, namely (1) title and abstract level screening and (2) full-text level screening.

2.4.1 Title and Abstract Level Screening

In this phase, the inclusion/exclusion criteria in Section 2.3 were applied to publication title and abstract. To minimize researcher bias, two researchers (first and second author) independently analyzed the search results. Afterwards, the results were compared and disagreements were resolved through discussions. Moreover, for any unresolved disagreements consensus meetings [8] involving all three researchers (first, second, and third author) were arranged. The short-listed studies from this phase were used as input for the second phase.

2.4.2 Full-text Level Screening

In this phase, the selected studies from the first phase were further analyzed on the basis of full-text. Two researchers (first and second author) independently applied the inclusion/exclusion criteria in Section 2.3 on the full-text. In this phase, the researchers also documented a reason for each excluded study [55]. The results were compared in a similar way as in the first phase and disagreements between the researchers were resolved through discussions and consensus meetings. The output of this phase was the final short-listed set of primary studies.

2.5 Study Quality Assessment Checklist and Procedure

The final selected studies based on the selection procedure in Section 2.4 were assessed for their quality merit. To minimize researcher bias, two researchers (first and second author) independently assessed the quality merit of the final selected studies. Any studies not meeting the minimum quality requirements were excluded from the final set of primary studies.

Table 4 presents the checklist for study quality assessment. For each question in the checklist, a three-level, numeric scale was used [55]. The levels are: yes (2 points), no (0 point), and partial (1 point). Based on the checklist and the numeric scale, a study could score a maximum of 34 and a minimum of 0 points. Moreover, for each study, the two independent scores from the two researchers were aggregated by computing arithmetic mean. We used the first quartile ($34/4 = 8.5$) as the cutoff point for the inclusion of studies. Therefore, if a study scored less than 8.5 points, it was excluded due to its lack of quality merit. The researchers documented the obtained score of each included/excluded study.

Table 4: Study quality assessment checklist

#	Question
Theoretical contribution	
1	Is the VM consolidation approach clearly described?
2	Is the pseudocode of the proposed algorithm included and clearly described?
3	Is a concise mathematical notation used?
4	Is the underlying theory clearly described?
5	Are the assumptions clearly stated?
6	Are the optimization objectives clearly stated?
Experimental evaluation	
7	Is the evaluation method clearly described?
8	Is a prototype implementation presented?
9	Is a simulation presented?
10	Is the experimental design clearly described?
11	Is the experimental setup clearly stated?
12	Are realistic load patterns used in the experimental design?
13	Are results from multiple different experiments included?
14	Are results from multiple runs of each experiment included?
15	Are the experimental results compared with other state-of-the-art VM consolidation approaches?
16	Is a statistical test used to assess the statistical significance of the results?
17	Are the limitations or threats to validity clearly stated?

2.6 Data Extraction Strategy

Table 5 presents the data extraction form which was used to record the data that the researchers extracted from the primary studies. Two researchers (first and second author) independently extracted data from all of the primary studies. The extracted data were used to classify the primary studies in a number of different ways. If a paper contained more than one VM consolidation approaches, it was classified as coming under more than one heading [6].

Budgen et al. [6] reported that differences in the classification of papers is a recurring theme in SMSs, even for experienced researchers. The two researchers differed considerably when classifying the papers according to the data extraction form in Table 5. The extracted data were compared and the differences were resolved in consensus meetings and by referring back to the original papers [8].

Table 5: Data extraction form

Data item	Value	Additional notes
General		
Data extractor name		
Data extraction date		
Study ID (S1, S2, S3, ...)		
Bibliographic reference (title, authors, year, journal/conference/workshop name)		
Author affiliations and countries		
Publication type (journal, conference, or workshop)		
VM consolidation related		
Type of algorithm used (e.g., heuristic, metaheuristic)		
Specific algorithm used (e.g., first-fit decreasing, genetic algorithm)		
Online or offline optimization		
Number of optimization objectives		
Name(s) of optimization objective(s)		
Problem formulation (e.g., single-objective, multi-objective with AOF)		
Name of evaluation method (e.g., analytical, experiment)		
Evaluation tool (e.g., simulation, prototype implementation)		
Type of load pattern(s) used (realistic, synthetic, or hybrid)		
Name(s) of other VM consolidation approach(es) used for the comparison of the results		
Study ID(s) of other VM consolidation approach(es) used for the comparison of the results		
Name of the statistical test used (e.g., Wilcoxon Signed-Rank Test)		
Statistically significant results (yes, no)		

2.7 Synthesis of the Extracted Data

The extracted data based on the data extraction strategy in Section 2.6 were synthesized separately for each RQ. The papers were categorized in a variety of dimensions and counts of the number of papers in different categories were recorded [61].

Perhaps the most important result of a SMS is a systematic map, which allows to identify evidence clusters and evidence deserts to direct the focus of future SLRs and to highlight areas where more primary studies should be performed [33]. It is important to present the systematic map in an appropriate visual format that provides a quick overview of the field and supports better analyses [42]. Therefore, a visual representation of the systematic map was created by using appropriate chart types including a pie chart and several bubble charts [43]. The visual representation of the systematic map allowed thematic analysis [61] to see which categories were well investigated and to identify research gaps [42].

2.8 Schedule of the Study

The initial draft of the review protocol was written on 09.3.2016. From 10.3.2016 to 18.3.2016, the protocol went through internal and external reviews resulting in several major and minor revisions. The internal reviewers were the authors themselves, while the external reviewer was Muhammad Usman¹ from the Department of Software Engineering at Blekinge Institute of Technology.

Table 6 presents the schedule of the SMS. The duration for title and abstract level screening was based on approximately 40 abstracts per day. Similarly, the duration for full-text level screening was based on approximately 5 full-text papers a day. The durations in Table 6 implicitly also include approximately 10%

¹<https://www.bth.se/eng/staff/muhammad-usman-muu/>

time for consensus meetings to resolve disagreements among the researchers. It should be noted that the search for primary studies was completed on 24.3.2016. Therefore, any papers published after this date are not covered in the mapping study.

Table 6: Schedule of the study

#	Activity	Duration	Due date
1	Protocol review and revision	07 business days	18.3.2016
2	Search for primary studies	04 business days	24.3.2016
3	Title and abstract level screening	07 business days	06.4.2016
4	Full-text level screening	07 business days	15.4.2016
5	Quality assessment	02 business days	19.4.2016
6	Data extraction	05 business days	26.4.2016
7	Data synthesis	05 business days	03.5.2016
8	Initial draft of paper	10 business days	18.5.2016

3 Results

In this section, we present the results of the SMS on distributed VM consolidation approaches. Table 7 presents the number of papers in different stages of the SMS. The results show that the initial search retrieved 202 results. However, out of 202, 86 were found duplicate and were subsequently removed. The main reason for such a large number of duplicate results is that the results from the ACM Digital Library were based on the ACM Guide to Computing Literature, which provides an expanded search that includes papers from the ACM full-text collection as well as from a number of other digital libraries including IEEE Xplore, ScienceDirect, and SpringerLink. The advantage of using the ACM Guide to Computing Literature is that it often finds more papers, but the disadvantage is that some of those papers are also found in the IEEE Xplore, ScienceDirect, and SpringerLink digital libraries. There were also a few cases where a VM consolidation approach was published in several papers. In such cases, only the most recent paper was included.

Table 7: Number of papers in different stages

SMS Stage	Number of papers
Initial search results	202
After removing duplicates	116
After title and abstract level screening	41
After full-text level screening	21
After quality assessment	19

From the remaining 116 papers, 75 papers were removed in the title and abstract level screening, resulting in 41 papers. Another 20 papers were removed in the full-text level screening. Out of these, 15 did not present a VM consolidation approach [1, 9, 19, 26–29, 32, 45, 47, 52, 54, 58, 60, 63], 2 presented a centralized VM consolidation approach [15, 23], 2 were published in several papers [20, 37], and 1 had a duplicate record [49]. Finally, from the remaining 21 papers, 2 papers were removed in the quality assessment stage because their aggregate scores were below the cutoff point [7, 34]. It resulted in a total of 19 papers for data extraction and synthesis. Table 8 presents study identifiers (IDs) and references of the final selected primary studies. Each primary study in Table 8 presents a distributed VM consolidation approach.

Before presenting detailed results, we start with a simple, graphical overview of the topic. Figure 2 presents a word cloud of distributed VM consolidation approaches generated from the titles and abstracts of the 19 primary studies. Common English words and those appearing only once were removed from the word list. Moreover, different forms and alternate spellings of a word were aggregated. The word cloud shows that some of the most frequent words include energy, cloud, data, and cost, which appear 56, 55, 46, and 45 times, respectively.

3.1 RQ1: Distributed VM Consolidation Approaches

The 19 distributed VM consolidation approaches can be classified into three distinct categories: (1) pure distributed VM consolidation algorithms, (2) centralized algorithms with a distributed architecture for VM consolidation, and (3) VM consolidation algorithms for geographically distributed data centers. Table 9 presents the three categories of distributed VM consolidation approaches. A pure distributed VM consolidation approach uses a distributed rather than a centralized algorithm to find a migration plan that optimizes the placement of VMs on PMs. In contrast, a distributed architecture approach provides a decentralized schema or multi-agent architecture for PMs [11, 14], but implements a centralized VM consolidation algorithm. The approaches in

3.2 RQ2: Algorithm Types and Names

Table 10 presents algorithm types and algorithm names along with the study IDs and counts. The results show that the most frequently used algorithm in distributed VM consolidation approaches is distributed or coordinated local search heuristic, which was used in six primary studies. The second most used algorithms are Ant Colony Optimization (ACO) and greedy, which were both used in two primary studies. All other algorithms were used in only one primary study. The results also show that S6 appears under two headings, namely greedy and genetic algorithm. This is because S6 presents a hybrid, two-stage optimization approach that combines a genetic algorithm with a greedy, best-fit approach for local improvement.

Table 10: Types of algorithms and specific algorithms

Type	Algorithm name	Study IDs	Count
Heuristic	Distributed or coordinated local search	S1, S5, S7, S11, S12, S19	6
	Greedy	S6, S18	2
	Weighted maximum bipartite matching	S13	1
	Static threshold-based	S14	1
	Distributed bartering algorithm	S17	1
Metaheuristic	Local search	S2	1
	ACO	S3, S4	2
	Genetic algorithm	S6	1
	Constraint programming-based large neighborhood search	S15	1
Machine learning	Fuzzy Q-learning algorithm	S8	1
Statistical	Probabilistic search	S9	1
	Gibbs sampling and the alternating direction method of multipliers	S10	1
	Metropolis-Hastings algorithm	S16	1

Figure 3 presents a summary of the four algorithm types from Table 10. It shows that 11 primary studies presented a heuristic approach, 5 presented a metaheuristic, 3 presented a statistical approach, and 1 presented a machine learning approach. Therefore, the results indicate that heuristics and metaheuristics are currently the most popular algorithm types for distributed VM consolidation, which collectively account for 80% of all algorithm types.

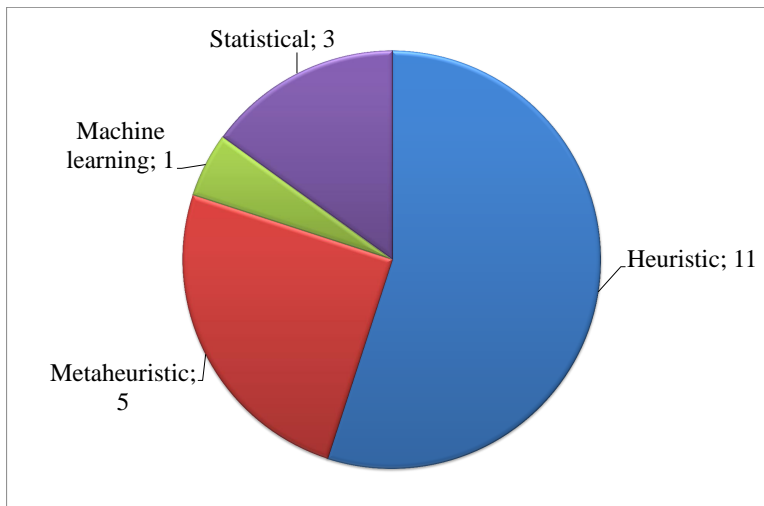


Figure 3: Algorithm types

Table 11 presents a classification of the primary studies with respect to online and offline optimization techniques. The results show that only 2 primary studies presented an online optimization approach, while 14 primary studies are based on offline optimization. Moreover, 3 primary studies were classified as inconclusive because the papers do not contain sufficient information in this regard. Therefore, the results indicate that 89% of the primary studies either use offline optimization or are classified as inconclusive. This is an important result

and is in line with past research in this area. Other researchers have also found that most of the optimization problems in cloud computing are currently being addressed with offline optimization techniques [22]. However, a drawback of the offline optimization techniques is that they require complete knowledge of the problem including possible future events, which is often difficult in a real world setting. In contrast, an online optimization algorithm receives input and produces output in an online manner [5]. Therefore, to be able to better cope with workload variability of different types of modern applications, VM consolidation should be performed continuously in an online manner [5, 11]. Hence, there is a clear need to develop online optimization approaches for distributed VM consolidation.

Table 11: Online or offline optimization

Online or offline	Study IDs	Count
Online	S3, S9	2
Offline	S2, S4, S5, S6, S7, S10, S12, S13, S14, S15, S16, S17, S18, S19	14
Inconclusive	S1, S8, S11	3

3.3 RQ3: Objectives

Table 12 categorizes the primary studies with respect to the different problem formulations and the number of optimization objectives. The results show that 11 studies proposed a multi-objective (two or three objectives) problem formulation with an AOF. Similarly, 2 studies presented a many-objective (four or more objectives) [57] problem formulation with an AOF. The remaining 6 studies proposed a single-objective problem formulation. The results also indicate that 74% of the primary studies optimize either one or two objectives. Hence, there is currently more published work on approaches that optimize fewer objectives. It should be noted that, although 13 approaches optimize two or more objectives, none of the studies proposed a pure multi-objective or many-objective problem formulation. The benefit of the AOF approach is that it reduces complexity and often improves runtime of the algorithm by limiting the search to a subspace of the feasible solutions. However, the drawback is that a correct combination of the objectives requires certain weights to be assigned to each objective, which often requires an in-depth knowledge of the problem domain [46]. Therefore, the assignment of the weights is essentially subjective [24]. In contrast, pure multi-objective and many-objective approaches do not require weights and often provide a more elaborate search of the search space. Hence, there exists a research gap. There is a need to develop novel distributed VM consolidation approaches that formulate the problem as pure multi-objective or many-objective.

Table 12: Problem formulations and number of optimization objectives

Problem formulation	Objectives	Study IDs	Count
Single-objective	1 objective	S7, S12, S13, S14, S17, S19	6
Multi-objective with AOF	2 objectives	S2, S4, S6, S8, S9, S10, S11, S16	8
	3 objectives	S3, S5, S18	3
Many-objective with AOF	4 objectives	S15	1
	5 objectives	S1	1

Figure 4 presents the number of primary studies (as bubble size) with respect to the four different algorithm types and the five different number of objectives. Again, S6 appears twice. The results show that in 6 out of 11 studies involving heuristics, a simpler problem formulation was proposed because it optimizes a single objective. Whereas, none of the other algorithm types was used for single-objective optimization. Therefore, although heuristics are currently the most published algorithm type for distributed VM consolidation, they are mostly used for simpler, single-objective problem formulations.

Table 13 categorizes the primary studies on the basis of different optimization objectives. The results show that minimizing energy consumption is the most commonly found optimization objective in the currently published distributed VM consolidation approaches as it appears in 11 out of 19 studies. Other popular optimization objectives include minimizing SLA violations and minimizing the number of VM migrations, which were found in 6 and 4 studies, respectively. According to Table 13, the existing distributed VM consolidation approaches optimize a total of twelve different objectives, which can be classified into four higher-level categories, namely network, performance, energy, and cost. However, there is a clear emphasis on minimizing energy consumption.

Figure 5 presents the number of primary studies (as bubble size) with respect to the four different algorithm types and the twelve different optimization objectives. Once again, S6 appears twice under minimizing the number of VM migrations and under minimizing energy consumption. The results in Figure 5 show that minimizing energy consumption and minimizing SLA violations are the two common objectives in all four algorithm types. The bubble plot in Figure 5 also highlights a number of research gaps for potential future




Figure 4: Number of primary studies (as bubble size) with respect to algorithm types and number of optimization objectives

Table 13: Names of optimization objectives

Category	Name of objective	Study IDs	Count
Network	Minimizing migration cost	S1, S15, S16	3
	Minimizing network communication cost	S1, S17, S18	3
	Minimizing the number of VM migrations	S3, S4, S5, S6	4
Performance	Minimizing the risk of resource contention	S1	1
	Minimizing SLA violations	S2, S3, S8, S9, S15, S18	6
	Minimizing delay cost	S10	1
	Maximizing load balancing	S1, S11, S16	3
Energy	Minimizing energy consumption	S3, S6, S8, S9, S10, S11, S12, S13, S14, S15, S18	11
	Maximizing the number of released PMs	S4, S7	2
	Maximizing resource efficiency	S1, S5	2
Cost	Minimizing the overall operational cost	S15, S19	2
	Minimizing rental cost	S2	1

primary studies. It shows that 9 out of 12 optimization objectives are currently being addressed with only one or two algorithm types. For example, minimizing network communication cost appears only in heuristic approaches. So far, the other three algorithm types have not been investigated for this objective. Hence, there is an opportunity to develop new distributed VM consolidation approaches to bridge these research gaps.

3.4 RQ4: Evaluation Methods and Tools

We also extracted data to categorize the primary studies on the basis of the evaluation methods used. We found that experiment is the most common evaluation method for distributed VM consolidation approaches as it was used in all 19 primary studies. Moreover, analytical models were used in 4 studies to additionally provide a proof of a theoretical approximation or limit. Therefore, experiment is the main evaluation method.

Table 14 categorizes the primary studies on the basis of the different types of evaluation tools and load patterns used in the experiments. The results show that the most common evaluation tool in distributed VM consolidation approaches is simulation as it was used in 16 studies. Moreover, 3 studies presented a prototype implementation. Therefore, simulation is currently the most popular evaluation tool for distributed VM consolidation approaches. According to Table 14, synthetic, realistic, and hybrid load patterns were used in 11, 7 and, 3 studies, respectively. It should be noted that the experiments in S2 and S3 were based on both synthetic and realistic load patterns. The results illustrate that synthetic load patterns are the most common




Figure 5: Number of primary studies (as bubble size) with respect to algorithm types and optimization objectives

type of load patterns used in the existing distributed VM consolidation approaches. Moreover, the experiments in 10 studies used simulations involving synthetic load patterns. On the other hand, there is currently only one study under each load pattern type that uses a prototype implementation. Therefore, there is clearly a need for more primary studies involving prototype implementations under all three types of load patterns.

Table 14: Types of evaluation tools and load patterns used

Evaluation tool	Load pattern	Study IDs	Count
Simulation	Synthetic	S1, S2, S3, S5, S7, S12, S14, S15, S16, S19	10
	Realistic	S2, S3, S8, S9, S10, S13	6
	Hybrid	S6, S11	2
Prototype	Synthetic	S4	1
	Realistic	S18	1
	Hybrid	S17	1

Table 15 categorizes the primary studies based on the names of the other VM consolidation approaches that were used in the primary studies for a comparison of the results. The results show that 9 studies did not contain a comparison of the results. Moreover, each of first-fit decreasing, static threshold approach, median absolute deviation, interquartile range, and best-fit decreasing was used in 2 studies. Each of the other 15 approaches was used in only one study. Therefore, the results show that a total of 20 different approaches were used for a comparison of the results in 10 primary studies. Moreover, most of the approaches were used in only one study. Hence, there is currently no general agreement in the cloud computing community concerning good approaches for a comparison of the results.

Table 16 presents study IDs of the other distributed VM consolidation approaches that were used for the comparison of the results in the primary studies. The results show that S4 is the only primary study in which the results of the proposed approach were compared with another distributed VM consolidation approach (S7). It is an important result. It shows that all approaches used for the comparison of the results except V-MAN (S7) are centralized approaches. Consequently, in 9 out of 10 studies that contain a comparison of the results, the results of the proposed distributed VM consolidation approach are compared with centralized VM consolidation approaches. Therefore, there exists little evidence on how the different distributed VM consolidation approaches compare to one another. For more meaningful comparisons of the results in future primary studies on distributed VM consolidation, we recommend that one or more of the 19 approaches studied in this SMS should be considered.

Another important question concerning the evaluation of a distributed VM consolidation approach is that

Table 15: Names of other VM consolidation approaches used for a comparison of the results

Name of other approach	Study IDs	Count
None	S5, S7, S9, S10, S13, S15, S16, S18, S19	9
First-fit decreasing	S1, S4	2
Centralized greedy algorithm	S2	1
Ant colony optimization with vector algebra	S3	1
Static threshold approach	S3, S8	2
Median absolute deviation	S3, S8	2
Interquartile range	S3, S8	2
Local regression	S3	1
Sercon	S4	1
V-MAN	S4	1
Best-fit decreasing	S6, S12	2
Random	S11	1
Incremental consolidation using thresholds	S11	1
One-shot merge	S11	1
Genetic algorithm	S12	1
Energy and migration-cost-aware approach (pMapper)	S12	1
Probability-based approach	S12	1
Energy-aware hierarchical scheduling	S14	1
Energy-aware scheduling for infrastructure clouds	S14	1
No migration	S17	1
Optimal	S17	1

Table 16: Study IDs of other VM consolidation approaches used for the comparison of the results

Study ID of other approach	Study IDs	Count
None	S1, S2, S3, S5, S6, S7, S8, S9, S10, S11, S12, S13, S14, S15, S16, S17, S18, S19	18
S7	S4	1

whether or not a statistical test was used to assess the statistical significance of the results. Therefore, we extracted this information from each primary study. The results show that none of the 19 studies includes a statistical test. Therefore, it is not a common practice to use statistical testing.

3.5 RQ5: Publication Forums

Table 17 presents publication forum names for the 19 distributed VM consolidation papers over time. The results show that the papers are published in 17 different publication forums. There are only two publication forums in which more than one papers have been published. These include IEEE Transactions on Services Computing and IEEE International Conference on Cloud Computing Technology and Science. Therefore, it is difficult to conclude whether or not there are any popular publication forums for distributed VM consolidation approaches.

Figure 6 presents the number of primary studies with respect to the three different publication types over time. The results show that 8 primary studies were published in journals, 10 were published in conferences, and 1 was published in a workshop. Moreover, from 2010 to 2014 (five years), a total of 10 primary studies were published. Whereas, 6 studies were published in 2015 alone. Similarly, 3 studies are published already in 2016. Furthermore, all 8 journal papers were published between 2014 and 2016. Hence, there is currently an increasing amount of interest on developing and evaluating novel distributed VM consolidation approaches and publishing them in more rigorous publication forums.

4 Validity Evaluation

In this section, we discuss major threats to the validity of the results presented in this paper. The first main threat is related to the coverage of the relevant literature. To mitigate this threat, we designed a comprehensive search strategy based on the SLR and SMS guidelines in [33, 61]. The search terms were extracted from the RQs and were validated against a set of recent and prominent works on VM consolidation including [5, 11–14, 16–18, 25, 36]. Similarly, the search strings were validated against a set of known studies on distributed VM consolidation including [11, 14, 36]. The search was performed in four major computer science digital libraries. Finally, the search in the ACM Digital Library was performed by using the ACM Guide to Computing Literature, which provides an expanded search.

Table 17: Publication forums for distributed VM consolidation papers over time

Publication forum name	Year	IDs
Future Generation Computer Systems	2016	S1
Proceedings of the 9th IFIP TC 6 International Conference on Networking	2010	S2
IEEE Transactions on Services Computing	2015	S3, S5
IEEE 4th International Conference on Cloud Computing Technology and Science	2012	S4
IEEE Transactions on Cloud Computing	2015	S6
IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks	2011	S7
IEEE 4th Symposium on Network Cloud Computing and Applications	2015	S8
IEEE 27th International Parallel and Distributed Processing Symposium Workshops PhD Forum	2013	S9
IEEE Communications Letters	2016	S10
IEEE 6th International Conference on Cloud Computing Technology and Science	2014	S11
IEEE Transactions on Systems, Man, and Cybernetics: Systems	2016	S12
IEEE 21st International Conference on Parallel and Distributed Systems	2015	S13
IEEE 10th International Conference on High Performance Computing and Communications & IEEE International Conference on Embedded and Ubiquitous Computing	2013	S14
IEEE/ACM 6th International Conference on Utility and Cloud Computing	2013	S15
9th Asia-Pacific Symposium on Information and Telecommunication Technologies	2012	S16
39th International Conference on Parallel Processing	2010	S17
Computer Communications	2014	S18
Computer Networks	2015	S19

The second threat is related to the selection of the primary studies. The results show that 75 out of 116 papers were excluded in the title and abstract level screening. It is possible that some relevant papers were erroneously excluded during the initial screening phase. To mitigate this threat, two researchers (first and second author) independently screened the titles and abstracts of each paper. The results were compared and disagreements were resolved through discussions. Moreover, for any unresolved disagreements, consensus meetings [8] were arranged. A similar approach was used in the full-text level screening phase.

The third major threat is related to data extraction and classification of studies. As reported by Budgen et al. [6], differences in the classification of papers is a recurring theme in SMSs, even for experienced researchers. To mitigate this threat, two researchers (first and second author) independently extracted data from all 19 primary studies. The extracted data were compared and the differences were resolved in consensus meetings and by referring back to the original papers.




Figure 6: Number of primary studies (as bubble size) with respect to publication types over time

5 Conclusions

In this paper, we presented a Systematic Mapping Study (SMS) of distributed Virtual Machine (VM) consolidation approaches. We used Systematic Literature Review (SLR) and SMS guidelines in the literature to design a comprehensive search strategy. The initial search returned 202 results from four major computer science digital libraries. After the removal of duplicate results and the application of the inclusion/exclusion criteria at two levels, 21 primary studies were selected. Finally, 2 studies were excluded in the quality assessment stage, which left 19 primary studies for data extraction and synthesis.

The objective of the SMS was to provide a comprehensive, unbiased overview of the state-of-the-art on distributed VM consolidation approaches. The SMS comprises five Research Questions (RQs) concerning: (1) existing approaches, (2) types of algorithms being used, (3) objectives being optimized, (4) evaluation methods and tools being used, and (5) popular publication forums over time. The results of the first RQ showed that 14 out of 19 studies presented pure distributed VM consolidation algorithms, while 2 studies presented centralized algorithms with a distributed architecture for VM consolidation and 3 studies presented VM consolidation algorithms for geographically distributed data centers.

The answer to the second RQ showed that the existing distributed VM consolidation approaches use four different types of algorithms, namely heuristics, metaheuristics, machine learning algorithms, and statistical approaches. Moreover, heuristics and metaheuristics are currently the most popular algorithm types. The most frequently used algorithm is distributed or coordinated local search heuristic, while Ant Colony Optimization (ACO) and greedy are the second most used algorithms. Only a small fraction of the existing distributed VM consolidation approaches can be categorized as using an online optimization technique. Hence, online optimization techniques are currently not sufficiently investigated for distributed VM consolidation.

For the third RQ concerning optimization objectives, we categorized the primary studies with respect to number and name of objectives and problem formulation. The results showed that nearly $\frac{3}{4}$ of the primary studies optimize either one or two objectives and only a few approaches optimize more than two objectives. The existing distributed VM consolidation approaches optimize a total of 12 different objectives. The most popular optimization objective is minimizing energy consumption. Other popular objectives include minimizing Service Level Agreement (SLA) violations and minimizing the number of VM migrations. 9 out of 12 objectives are currently being addressed with only one or two algorithm types. Hence, the focus of future primary studies may be directed to investigate the remaining algorithm types for optimizing these objectives. About $\frac{2}{3}$ of the studies presented a multi-objective or many-objective problem formulation with an Aggregate Objective Function (AOF), while the rest of the studies presented a single-objective problem formulation. None of the studies presented a pure multi-objective or many-objective problem formulation. Hence, future research may be directed to develop pure multi-objective and many-objective distributed VM consolidation approaches.

The results of the fourth RQ showed that experiment is the most common evaluation method for distributed VM consolidation approaches. The most common evaluation tool is simulation. Moreover, synthetic load patterns are the most common type of load patterns. Therefore, simulations involving synthetic load patterns are currently the most common evaluation tool. We also extracted data with respect to the other VM consolidation approaches that were used for a comparison of the results. The results showed that 9 studies did not contain a comparison of the results. Moreover, a total of 20 different approaches were used for a comparison of the results in the remaining 10 studies. In 9 out of 10 studies that contained a comparison of the results, the results of the proposed distributed VM consolidation approach were compared with centralized VM consolidation approaches. Therefore, there exists little evidence on how the different distributed VM consolidation approaches compare to one another. Hence, there is a need for more comparative studies involving multiple distributed VM consolidation approaches. We recommend that one or more of the 19 approaches studied in this SMS should be considered for more meaningful comparisons of the results in future studies.

The answer to the fifth RQ showed that the 19 studies were published in 17 different publication forums. Therefore, it is difficult to conclude whether or not there are any popular publication forums for distributed VM consolidation approaches. There is currently an increasing amount of interest on developing and evaluating novel distributed VM consolidation approaches for cloud data centers and publishing them in more rigorous publication forums.

Acknowledgements

The authors would like to thank Muhammad Usman at Blekinge Institute of Technology for his valuable comments on the review protocol. The work was supported by the Need for Speed (N4S) Program (<http://www.n4s.fi>). Adnan Ashraf also received a research grant from the Ulla Tuominen Foundation.

References

- [1] A. Abdul, R. Iqbal, A. Jame, M. Odetayo, and N. Shah. Nora: Network oriented resource allocation for data intensive applications in the cloud environment. In *Computer Supported Cooperative Work in Design (CSCWD), Proceedings of the 2014 IEEE 18th International Conference on*, pages 295–300, May 2014.
- [2] T. Adhikary, A. K. Das, M. A. Razzaque, and A. M. J. Sarkar. Energy-efficient scheduling algorithms for data center resources in cloud computing. In *High Performance Computing and Communications 2013 IEEE International Conference on Embedded and Ubiquitous Computing (HPCC_EUC), 2013 IEEE 10th International Conference on*, pages 1715–1720, Nov 2013.
- [3] R. W. Ahmad, A. Gani, S. H. A. Hamid, M. Shiraz, A. Yousafzai, and F. Xia. A survey on virtual machine migration and server consolidation frameworks for cloud data centers. *Journal of Network and Computer Applications*, 52:11 – 25, 2015.
- [4] L. A. Barroso and U. Hölzle. The case for energy-proportional computing. *Computer*, 40(12):33–37, Dec 2007.
- [5] A. Beloglazov and R. 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, Sept. 2012.
- [6] D. Budgen, M. Turner, P. Brereton, and B. Kitchenham. Using mapping studies in software engineering. In *Proceedings of PPIG 2008*, pages 195–204. Lancaster University, 2008.
- [7] P. Donadio, G. B. Fioccola, R. Canonico, and G. Ventre. Combining it and network orchestration for green clouds. In *Intelligent Networking and Collaborative Systems (INCoS), 2014 International Conference on*, pages 539–544, Sept 2014.
- [8] T. Dybå and T. Dingsøy. Empirical studies of agile software development: A systematic review. *Information and Software Technology*, 50(9-10):833 – 859, 2008.
- [9] E. Elmroth, J. Tordsson, F. Hernández, A. Ali-Eldin, P. Svärd, M. Sedaghat, and W. Li. Self-management challenges for multi-cloud architectures. In *Proceedings of the 4th European Conference on Towards a Service-based Internet, ServiceWave’11*, pages 38–49, Berlin, Heidelberg, 2011. Springer-Verlag.
- [10] X. Fan, W.-D. Weber, and L. A. Barroso. Power provisioning for a warehouse-sized computer. *SIGARCH Computer Architecture News*, 35(2):13–23, June 2007.
- [11] F. Farahnakian, A. Ashraf, T. Pahikkala, P. Liljeberg, J. Plosila, I. Porres, and H. Tenhunen. Using ant colony system to consolidate VMs for green cloud computing. *Services Computing, IEEE Transactions on*, 8(2):187–198, March 2015.
- [12] F. Farahnakian, P. Liljeberg, and J. Plosila. Energy-efficient virtual machines consolidation in cloud data centers using reinforcement learning. In *Parallel, Distributed and Network-Based Processing (PDP), 2014 22nd Euromicro International Conference on*, pages 500–507, Feb 2014.
- [13] F. Farahnakian, T. Pahikkala, P. Liljeberg, and J. Plosila. Energy aware consolidation algorithm based on K-nearest neighbor regression for cloud data centers. In *Utility and Cloud Computing (UCC), 6th IEEE/ACM International Conference on*, 2013.
- [14] E. Feller, C. Morin, and A. Esnault. A case for fully decentralized dynamic VM consolidation in clouds. *Cloud Computing Technology and Science, IEEE International Conference on*, pages 26–33, 2012.
- [15] E. Feller, L. Rilling, and C. Morin. Energy-aware ant colony based workload placement in clouds. In *Grid Computing (GRID), 2011 12th IEEE/ACM International Conference on*, pages 26–33, September 2011.
- [16] M. Ferdaus, M. Murshed, R. Calheiros, and R. Buyya. Virtual machine consolidation in cloud data centers using ACO metaheuristic. In F. Silva, I. Dutra, and V. Santos Costa, editors, *Euro-Par 2014 Parallel Processing*, volume 8632 of *Lecture Notes in Computer Science*, pages 306–317. Springer International Publishing, 2014.
- [17] T. C. Ferreto, M. A. Netto, R. N. Calheiros, and C. A. D. Rose. Server consolidation with migration control for virtualized data centers. *Future Generation Computer Systems*, 27(8):1027–1034, 2011.
- [18] Y. Gao, H. Guan, Z. Qi, Y. Hou, and L. Liu. A multi-objective ant colony system algorithm for virtual machine placement in cloud computing. *Journal of Computer and System Sciences*, 79(8):1230 – 1242, 2013.

- [19] L. Gu, D. Zeng, S. Guo, and B. Ye. Joint optimization of vm placement and request distribution for electricity cost cut in geo-distributed data centers. In *Computing, Networking and Communications (ICNC), 2015 International Conference on*, pages 717–721, Feb 2015.
- [20] J. O. Gutierrez-Garcia and A. Ramirez-Nafarrate. Policy-based agents for virtual machine migration in cloud data centers. In *Services Computing (SCC), 2013 IEEE International Conference on*, pages 603–610, June 2013.
- [21] J. O. Gutierrez-Garcia and A. Ramirez-Nafarrate. Collaborative agents for distributed load management in cloud data centers using live migration of virtual machines. *IEEE Transactions on Services Computing*, 8(6):916–929, Nov 2015.
- [22] M. Harman, K. Lakhota, J. Singer, D. R. White, and S. Yoo. Cloud engineering is search based software engineering too. *Journal of Systems and Software*, 86(9):2225–2241, 2013.
- [23] F. Hermenier, S. Demasse, and X. Lorca. Bin repacking scheduling in virtualized datacenters. In *Proceedings of the 17th International Conference on Principles and Practice of Constraint Programming, CP’11*, pages 27–41, Berlin, Heidelberg, 2011. Springer-Verlag.
- [24] X. B. Hu, M. Wang, and E. D. Paolo. Calculating complete and exact pareto front for multiobjective optimization: A new deterministic approach for discrete problems. *IEEE Transactions on Cybernetics*, 43(3):1088–1101, June 2013.
- [25] I. Hwang and M. Pedram. Hierarchical virtual machine consolidation in a cloud computing system. In *Cloud Computing (CLOUD), 2013 IEEE Sixth International Conference on*, pages 196–203, 2013.
- [26] M. A. Islam, S. Ren, N. Pissinou, A. H. Mahmud, and A. V. Vasilakos. Distributed temperature-aware resource management in virtualized data center. *Sustainable Computing: Informatics and Systems*, 6:3 – 16, 2015. Special Issue on Selected Papers from 2013 International Green Computing Conference (IGCC).
- [27] C. Jing, Y. Zhu, and M. Li. Customer satisfaction-aware scheduling for utility maximization on geo-distributed data centers. *Concurr. Comput. : Pract. Exper.*, 27(5):1334–1354, Apr. 2015.
- [28] G. Jung, M. A. Hiltunen, K. R. Joshi, R. D. Schlichting, and C. Pu. Mistral: Dynamically managing power, performance, and adaptation cost in cloud infrastructures. In *Distributed Computing Systems (ICDCS), 2010 IEEE 30th International Conference on*, pages 62–73, June 2010.
- [29] M. H. Kabir, G. C. Shoja, and S. Ganti. Vm placement algorithms for hierarchical cloud infrastructure. In *Proceedings of the 2014 IEEE 6th International Conference on Cloud Computing Technology and Science, CLOUDCOM ’14*, pages 656–659, Washington, DC, USA, 2014. IEEE Computer Society.
- [30] T. Kaur and I. Chana. Energy efficiency techniques in cloud computing: A survey and taxonomy. *ACM Computing Surveys*, 48(2):22:1–22:46, Oct. 2015.
- [31] E. Kavvadia, S. Sagiadinos, K. Oikonomou, G. Tsioutsoulouklis, and S. Aïssa. Elastic virtual machine placement in cloud computing network environments. *Computer Networks*, 93, Part 3:435 – 447, 2015. Cloud Networking and Communications {II}.
- [32] A. Khosravi, S. K. Garg, and R. Buyya. Energy and carbon-efficient placement of virtual machines in distributed cloud data centers. In *Proceedings of the 19th International Conference on Parallel Processing, Euro-Par’13*, pages 317–328, Berlin, Heidelberg, 2013. Springer-Verlag.
- [33] B. Kitchenham and S. Charters. Guidelines for performing Systematic Literature Reviews in Software Engineering (version 2.3). Technical Report EBSE-2007-01, Keele University and University of Durham, 2007.
- [34] L. Larsson, D. Henriksson, and E. Elmroth. Scheduling and monitoring of internally structured services in cloud federations. In *Proceedings of the 2011 IEEE Symposium on Computers and Communications, ISCC ’11*, pages 173–178, Washington, DC, USA, 2011. IEEE Computer Society.
- [35] D. Lucanin and I. Brandic. Pervasive cloud controller for geotemporal inputs. *IEEE Transactions on Cloud Computing*, PP(99):1–1, 2015.
- [36] M. Marzolla, O. Babaoglu, and F. Panzieri. Server consolidation in clouds through gossiping. In *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2011 IEEE International Symposium on a*, pages 1–6, June 2011.

- [37] S. S. Masoumzadeh and H. Hlavacs. Integrating vm selection criteria in distributed dynamic vm consolidation using fuzzy q-learning. In *Network and Service Management (CNSM), 2013 9th International Conference on*, pages 332–338, Oct 2013.
- [38] S. S. Masoumzadeh and H. Hlavacs. A cooperative multi agent learning approach to manage physical host nodes for dynamic consolidation of virtual machines. In *Network Cloud Computing and Applications (NCCA), 2015 IEEE Fourth Symposium on*, pages 43–50, June 2015.
- [39] T. Mastelic, A. Oleksiak, H. Claussen, I. Brandic, J.-M. Pierson, and A. V. Vasilakos. Cloud computing: Survey on energy efficiency. *ACM Computing Surveys*, 47(2):33:1–33:36, Dec. 2014.
- [40] C. Mastroianni, M. Meo, and G. Papuzzo. Analysis of a self-organizing algorithm for energy saving in data centers. In *Parallel and Distributed Processing Symposium Workshops PhD Forum (IPDPSW), 2013 IEEE 27th International*, pages 907–914, May 2013.
- [41] D. Mehta, B. O’Sullivan, and H. Simonis. Energy cost management for geographically distributed data centres under time-variable demands and energy prices. In *Utility and Cloud Computing (UCC), 2013 IEEE/ACM 6th International Conference on*, pages 26–33, Dec 2013.
- [42] K. Petersen, R. Feldt, S. Mujtaba, and M. Mattsson. Systematic mapping studies in software engineering. In *12th International Conference on Evaluation and Assessment in Software Engineering (EASE), University of Bari, Italy*, 2008.
- [43] K. Petersen, S. Vakkalanka, and L. Kuzniarz. Guidelines for conducting systematic mapping studies in software engineering: An update. *Information and Software Technology*, 64:1–18, 2015.
- [44] C. Pham, N. H. Tran, C. T. Do, E. N. Huh, and C. S. Hong. Joint consolidation and service-aware load balancing for datacenters. *IEEE Communications Letters*, 20(2):292–295, Feb 2016.
- [45] C. Pham, N. H. Tran, M. N. H. Nguyen, J. H. Son, and C. S. Hong. Hosting virtual machines on distributed datacenters. In *Proceedings of the 10th International Conference on Ubiquitous Information Management and Communication, IMCOM ’16*, pages 85:1–85:5, New York, NY, USA, 2016. ACM.
- [46] F. L. Pires and B. Barán. A virtual machine placement taxonomy. In *Cluster, Cloud and Grid Computing (CCGrid), 2015 15th IEEE/ACM International Symposium on*, pages 159–168, May 2015.
- [47] V. Ravi and H. S. Hameed. Reinforcement learning based service provisioning for a greener cloud. In *Eco-friendly Computing and Communication Systems (ICECCS), 2014 3rd International Conference on*, pages 85–90, Dec 2014.
- [48] Y. Sakumoto, M. Aida, and H. Shimonishi. Autonomous decentralized mechanisms for generating global order in large-scale system: Using metropolis-hastings algorithm and applying to virtual machine placement. In *Information and Telecommunication Technologies (APSITT), 2012 9th Asia-Pacific Symposium on*, pages 1–6, Nov 2012.
- [49] M. Sedaghat, F. Hernández-Rodríguez, and E. Elmroth. Decentralized cloud datacenter reconsolidation through emergent and topology-aware behavior. *Future Generation Computer Systems*, 56:51 – 63, 2016.
- [50] M. Sedaghat, F. Hernández-Rodríguez, E. Elmroth, and S. Girdzijauskas. Divide the task, multiply the outcome: Cooperative vm consolidation. In *Cloud Computing Technology and Science (CloudCom), 2014 IEEE 6th International Conference on*, pages 300–305, Dec 2014.
- [51] A. Shehabi, S. J. Smith, D. A. Sartor, R. E. Brown, M. Herrlin, J. G. Koomey, E. R. Masanet, N. Horner, I. L. Azevedo, and W. Lintner. United states data center energy usage report. Technical Report LBNL-1005775, Ernest Orlando Lawrence Berkeley National Laboratory, 2016.
- [52] M. Simonin, E. Feller, A.-C. Orgerie, Y. Jégou, and C. Morin. Snooze: An autonomic and energy-efficient management system for private clouds. In *Revised Selected Papers of the COST IC0804 European Conference on Energy Efficiency in Large Scale Distributed Systems - Volume 8046*, EE-LSDS 2013, pages 114–117, New York, NY, USA, 2013. Springer-Verlag New York, Inc.
- [53] J. Sonnek, J. Greensky, R. Reutiman, and A. Chandra. Starling: Minimizing communication overhead in virtualized computing platforms using decentralized affinity-aware migration. In *Parallel Processing (ICPP), 2010 39th International Conference on*, pages 228–237, Sept 2010.
- [54] H. Teyeb, A. Balma, N. B. H. Alouane, and S. Tata. Optimal virtual machine placement in large-scale cloud systems. In *Proceedings of the 2014 IEEE International Conference on Cloud Computing, CLOUD ’14*, pages 424–431, Washington, DC, USA, 2014. IEEE Computer Society.

- [55] M. Usman, E. Mendes, F. Weidt, and R. Britto. Effort estimation in agile software development: A systematic literature review. In *Proceedings of the 10th International Conference on Predictive Models in Software Engineering*, PROMISE '14, pages 82–91, New York, NY, USA, 2014. ACM.
- [56] L. Velasco, A. Asensio, J. Berral, E. Bonetto, F. Musumeci, and V. López. Elastic operations in federated datacenters for performance and cost optimization. *Computer Communications*, 50:142 – 151, 2014. Green Networking.
- [57] C. von Lüken, B. Barán, and C. Brizuela. A survey on multi-objective evolutionary algorithms for many-objective problems. *Computational Optimization and Applications*, 58(3):707–756, 2014.
- [58] R. Wang and N. Kandasamy. A distributed control framework for performance management of virtualized computing environments: Some preliminary results. In *Proceedings of the 1st Workshop on Automated Control for Datacenters and Clouds*, ACDC '09, pages 7–12, New York, NY, USA, 2009. ACM.
- [59] W. Wang, Y. Jiang, and W. Wu. Multiagent-based resource allocation for energy minimization in cloud computing systems. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, PP(99):1–16, 2016.
- [60] G. Wen, J. Hong, C. Xu, P. Balaji, S. Feng, and P. Jiang. Energy-aware hierarchical scheduling of applications in large scale data centers. In *Cloud and Service Computing (CSC), 2011 International Conference on*, pages 158–165, Dec 2011.
- [61] C. Wohlin, P. Runeson, M. Höst, M. C. Ohlsson, B. Regnell, and A. Wesslén. *Experimentation in Software Engineering*. Springer-Verlag Berlin Heidelberg, 1 edition, 2012.
- [62] Z. Xu, W. Liang, and Q. Xia. Electricity cost minimization in distributed clouds by exploring heterogeneity of cloud resources and user demands. In *Parallel and Distributed Systems (ICPADS), 2015 IEEE 21st International Conference on*, pages 388–395, Dec 2015.
- [63] K. yin Chen, Y. Xu, K. Xi, and H. J. Chao. Intelligent virtual machine placement for cost efficiency in geo-distributed cloud systems. In *Communications (ICC), 2013 IEEE International Conference on*, pages 3498–3503, June 2013.
- [64] Q. Zhang, J. Xiao, E. Gürses, M. Karsten, and R. Boutaba. Dynamic service placement in shared service hosting infrastructures. In *Proceedings of the 9th IFIP TC 6 International Conference on Networking*, NETWORKING'10, pages 251–264, Berlin, Heidelberg, 2010. Springer-Verlag.