Energy Efficient Algorithms based on VM Consolidation for Cloud Computing: Comparisons and Evaluations

Qiheng Zhou[1], Minxian Xu*[1], Sukhpal Singh Gill[2], Chengxi Gao[1],
Wenhong Tian[3], Chengzhong Xu[1,4], and Rajkumar Buyya[5]

[1]Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, China
[2]School of Electronic Engineering and Computer Science, Queen Mary University of London, UK
[3]School of Software and Information Engineering, University of Electronic Science and Technology of China, China
[4]Faculty of Science and Technology, Macao University, China
[5]School of Computing and Information Systems, University of Melbourne, Australia

Abstract—**Cloud Computing paradigm has revolutionized IT industry and be able to offer computing as the fifth utility. With the pay-as-you-go model, cloud computing enables to offer the resources dynamically for customers anytime. Drawing the attention from both academia and industry, cloud computing is viewed as one of the backbones of the modern economy. However, the high energy consumption of cloud data centers contributes to high operational costs and carbon emission to the environment. Therefore, Green cloud computing is required to ensure energy efficiency and sustainability, which can be achieved via energy efficient techniques. One of the dominant approaches is to apply energy efficient algorithms to optimize resource usage and energy consumption. Currently, various virtual machine consolidation-based energy efficient algorithms have been proposed to reduce the energy of cloud computing environment. However, most of them are not compared comprehensively under the same scenario, and their performance is not evaluated with the same experimental settings. This makes users hard to select the appropriate algorithm for their objectives. To provide insights for existing energy efficient algorithms and help researchers to choose the most suitable algorithm, in this paper, we compare several state-of-the-art energy efficient algorithms in depth from multiple perspectives, including architecture, modelling and metrics. In addition, we also implement and evaluate these algorithms with the same experimental settings in CloudSim toolkit. The experimental results show the performance comparison of these algorithms with comprehensive results. Finally, detailed discussions of these algorithms are provided.**

Index Terms—**Cloud Computing, Data Centers, Energy Efficient Algorithms, Virtual Machine Consolidation, QoS**

I. INTRODUCTION

The emergence of cloud computing has contributed the computing resources as a new utility, like electricity and gas, and shaped the way how IT resources can be used by customers. Since its emergence, cloud computing has developed rapidly into one of the backbones of modern economy. Cloud consumers including government, research institutes and industry enterprises have all embraced and benefited from

cloud computing significantly. Cloud computing also enables new business to be established within a short time, facilitates the expansion of enterprise globally, accelerates the progress of scientific research, and promotes the creation of various models and applications. Cloud service providers are also offering a variety of cloud services for customers with on-demand access to resources based on the pay-as-you-go model [1] [2] [3].

The infrastructure of cloud computing is cloud data centers. Currently, various cloud providers, including Amazon, Google, Microsoft, establish large-scale cloud data centers to fulfill the resources and services demands for customers. To ensure the availability and reliability, cloud data centers are required to be running 24/7. Nowadays, the majority of data centers spread of 300-4500 square meters containing hundreds to thousands of physical machines. A typical data center can consume up to 25,000 KWh per day. It is estimated that US data centers can consume 140 million kWh by 2020, which equals to 50 coal-based power plants. It is also claimed that the carbon footprint will reach 2-3% of global emission [4]. To relieve the huge energy consumption and carbon emission, energy efficient techniques are required in cloud data centers [5] [6] [7].

Virtualization technology is an important part of green cloud data centers to support energy efficiency via Virtual Machine (VM) consolidation, where a VM is the software implementation of a computer that running with an operating system and applications [2]. VM consolidation is referred as the process the VMs can be reallocated from one physical machine to another without affecting the execution of users' requests. VM consolidation has been identified as one of the dominant energy efficient solutions to reduce the energy consumption of cloud data centers [5] [8]. As VMs are packed on fewer physical machines via consolidation, idle physical machines can be turned off or switched to the low-power mode [9].

VM consolidation has been proven to be an effective approach to reduce data center energy consumption, and various energy efficient algorithms based on VM consolidation have been proposed. These algorithms aim to optimize the placement of VMs to reduce the energy consumption while ensuring other constraints, e.g. Service Level Agreement (SLA) violations. Due to the uncontrolled network traffic, it is hard to conduct experiments under large-scale environment and reproduce results. Therefore, running experiments with validated simulation toolkit is an acceptable and reasonable way. Simulation toolkit can be easily applied to construct large-scale environment and generate reproducible results [10].

Among all existing cloud data centers simulation toolkits, CloudSim [11] is the most widely used one. CloudSim supports both system and behavior for cloud data centers, including physical machines, virtual machines, workloads and resource scheduling policies. The resource provisioning model is also generic so that it can be extended with ease and limited efforts. These attractive features have attracted users from hundreds of universities and research institutes, and some extended simulators complementing CloudSim have also been developed, e.g. CloudAnalyst [12] and NetworkCloudSim [13].

The motivation of this article is to analyze the problem of rising energy consumption in depth through a comparison of some state-of-the-art energy efficient algorithms in cloud data centers. This article is particularly motivated from the following facts:

- The need and demand for understanding existing VM-based energy efficient algorithms for cloud data centers.
- The proposed algorithms were evaluated in different scenarios and configurations, thus their advantage and disadvantage are not carefully examined.
- The requirement for selecting the best suitable algorithm based on different priorities.

In this paper, we select various well-known VM consolidation-based energy efficient algorithms and implement experiments using CloudSim toolkit. The evaluated algorithms are picked from state-of-the-art energy efficient algorithms that have shown good performance in energy efficiency.

The main **contributions** of this work are as follows:

- Offering a cross-sectional view of the investigated VM consolidation-based energy efficient algorithms, which present outstanding performance in cloud computing area.
- Presenting a unified simulation-based analysis framework based on CloudSim that allows evaluation and comparison of energy-efficient VM consolidation algorithms in a unified and unbiased way.
- Discussing the advantages and disadvantages of the investigated algorithms to provide suggested algorithms for different scenarios.

The rest of the paper is organized as: We provide an overview on VM consolidation-based energy efficient algorithm for cloud computing in Section 2. The investigated algorithms are introduced in Section 3. Section 4 discusses the modelling of the investigated algorithms, and the adopted metrics in the investigated algorithms are summarized in Section 5. Section 6 shows the performance comparison of the investigated algorithms. Finally, conclusions and future research trends are given.

II. RELATED WORK

A. Energy Efficient Algorithms in Clouds

A few articles have conducted surveys or taxonomies on energy efficient algorithms based on VM consolidation for cloud data centers. Mansouri et al. [14] introduced a survey on resource management in cloud environment, and discussed VM consolidation-based energy efficient algorithms from cloud management system level. Kaur et al. [2] presented a comprehensive survey for energy efficient scheduling approach in clouds, and compared some consolidation-based techniques, including VM consolidation-based approach. Orgerie et al. [15] conducted a survey on techniques for improving the energy efficiency for large-scale distributed systems, and surveyed on VM migration algorithms for Clouds. Gill et al. [5] proposed a taxonomy for sustainable cloud computing and introduced a taxonomy for VM consolidation-based algorithms without focusing on energy efficiency. Mann et al. [16] introduced a survey on VM allocation in cloud data centers from problem modelling and optimization algorithms perspectives. Ahmad et al. [17] conducted a survey on VM migration and server consolidation framework for cloud data centers, in which the commonalities and differences of investigated VM migration algorithms are highlighted.

However, these surveys and taxonomy focused on the comparison with high-level comparison of VM consolidation-based energy efficient algorithms without evaluating the performance under experimental environments. Thus, our work advances the previous works by evaluating the state-of-the-art algorithms not only from modelling perspective but also with experimental comparisons. It also identifies the merits and demerits of investigated algorithms to provide suggestions for research in related areas.

B. VM consolidation-based Energy Efficient Algorithms

Beloglazov et al. [18] proposed an energy-aware allocation algorithm of data center resources that focuses on VM scheduling named Modified Best Fit Decreasing (MBFD). Their objective is to reduce the energy consumption of data centers while ensuring SLA. Mastroianni et al. [19] proposed energy efficient scheduling policy based on probabilistic procedures. Li et al. [20] presented a holistic VM scheduling algorithm capable of minimizing total data center energy consumption, including computing energy and cooling energy. Ranjbari et al. [21] introduced an algorithm based on learning automata for energy and SLA efficient consolidation of VMs in cloud data centers. Farahnkian et al. [22] proposed a novel dynamic VM consolidation approach based on Ant Colony Optimization to reduce the energy consumption of data centers.

In the performance evaluation of these algorithms, MBFD has been evaluated as the baseline, however, these algorithms

are not compared and evaluated together, thus it is hard to identify the performance of these VM consolidation-based energy efficient algorithms. In this paper, we aim to compare these algorithms in depth and evaluate them under the same configurations to show the performance comparison.
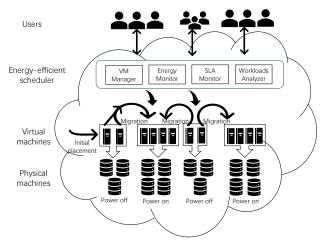
III. OVERVIEW OF THE INVESTIGATED ALGORITHMS

Fig. 1: Energy efficient scheduling based on VM consolidation

Fig. 1 shows the high-level architecture of VM consolidation-based energy efficient algorithms in cloud data centers. This architecture is generic and our investigated algorithms all conform it. The architecture mainly has 4 entities:

1. Users: Cloud users submit their requests to cloud, and the requests will be processed by cloud services.

2. Energy efficient scheduler: It can contain the following components to support the energy efficient scheduling and act as the interface between users and provisioned resources from service providers:

(a) VM manager: This component tracks the resource usage of VMs and makes decisions on VMs behaviors, including when and where to consolidate VMs. To achieve this objective, it requires the energy and SLA information from Energy Monitor and SLA monitor.

(b) Energy Monitor: It monitors the energy usage in the system caused by VMs and physical machines. It also provides information for the decisions of the VM manager.

(c) SLA Monitor: It monitors how SLA is influenced by the operations in the system. It can also represent performance constraints when reducing energy consumption in the system.

(d) Workloads Analyzer: This component interprets users' requests and service requirements before processing them. The workloads can be distributed to different VMs based on resource characteristics.

3. Virtual machines: The workloads are deployed and executed on VMs. The VMs can be managed according to the incoming workloads with two phases: initial placement and VM migration. VMs are placed by initial placement algorithm to be allocated to physical machines initially. According to

workloads, the placement can be optimized via VM consolidation algorithm, thus the unused machines can be temporarily turned off or switched to low-power mode.

4. Physical machines: The infrastructure offers physical machines to provision virtualized resources to meet users' requests.

Based on the following criteria, we carefully select 5 state-of-the-art VM consolidation-based energy efficient algorithms for our comparisons and evaluations:

- To make the comparison more persuasive, the algorithms were published in prominent journals or conferences, and the algorithms can be representative of a category of algorithms.
- To ensure the evaluation results reproducible, the algorithms were implemented in CloudSim or can be easily evaluated in CloudSim.
- To make the algorithm comparable, the algorithms should have been evaluated with the same baseline.

In the following subsections, we will provide the overview of our investigated algorithms. Table I shows the comparison of the investigated algorithms based on different parameters such as application type, operational environment, objective function, scheduling mechanism, scheduling criteria, type of workloads and identified merits and demerits of each algorithm.

A. MBFD

Modified Best Fit Decreasing (MBFD) [18] aims to reduce the energy consumption of data centers while ensuring SLA. It solves the VM initial placement phase by regarding it as a bin packing problem. MBFD is designed to assign the VMs to the hosts that produce the minimum incrementation of energy consumption. In the VM consolidation phase, the algorithm optimizes the VM allocation via consolidation for better energy efficiency, the target host is also selected as in the initial placement. The motivation of the proposed algorithm is due to the uncertainty of task lifecycles, some VMs are likely to host an over-provisioned of applications while others run with low resource utilization. Unbalanced workloads in cloud data centers result in a severe waste of resources and performance degradation. The VM consolidation in this work is a proactive process and can be applied with heterogeneous workloads.

MBFD has been evaluated as the baseline for many VM consolidation-based energy efficient algorithm. Various recent algorithms were proposed to improve the performance of this algorithm. The advantage of this algorithm is that it considers the trade-offs between the reduced energy and SLA violation and it is easy to implement. The limitation is that it does not consider the holistic resource management, which is complemented by recent research.

B. EcoCloud

EcoCloud [19] is an energy efficient scheduling policy based on probabilistic procedures. VM assignment and VM migration are the two phases in the algorithm. In the VM

TABLE I: High-level Comparison of Investigated Algorithms

Algorithm	Application Type	Operational Environment	Objective Function	Power Model	Power Component	Scheduling Mechanism	Workloads	Merits	Demerits
MBFD [18]	Dynamic workloads (web service)	Distributed and Heterogeneous	To optimize energy consumption	Linear	CPU and memory	Dynamic Consolidation (Proactive)	Heterogeneous workloads	Reduced energy consumption and SLA violation rate	There is a need of holistic resource management
EcoCloud [19]	Dynamic workloads	Distributed and homogeneous	To improve consolidation of VMs	Linear	Memory and CPU	Bounouli-based scheduling (Proactive)	Planetlab workloads traces	Evaluated the effect of power consumption on SLA violations	To investigate the impact of number of VM migrations on SLA violations
GRANITE [20]	Heterogeneous	Distributed	To investigate the temperature distribution of airflow and server CPU	Linear	Cooling, storage, memory, CPU and network	2D-Computational Fluid Dynamics (CFD) modelling-based scheduling (Proactive)	Google datacenter trace logs	Minimizing total datacenter energy consumption (cooling and computing)	To improve accuracy, 2-dimensional CFD model can be enhanced to 3-dimensional CFD model
LOAD [21]	Dynamic workloads	Homogeneous	To optimize energy consumption, number of VM migrations and SLA violations	Linear	CPU	Learning automata based scheduling (proactive)	CPU utilization	Improved CPU utilization and reduced SLA violations and energy consumption	Under-utilization of resource is not considered
ACS [22]	CPU and memory-intensive workloads	Distributed and Heterogeneous	To investigate the relationship between energy consumption, VM migrations and QoS	Linear	CPU and, Memory	Ant colony optimization and (reactive)	Heterogeneous workloads	Reduced energy and VM migrations	The impact of VM migration on network bandwidth can be investigated to reduce power consumption further

assignment phase, different from the MBFD introduced in the previous section, the VM manager sends an invitation to a subset of all active hosts to obtain a list of hosts that accept the incoming VM. After receiving the invitation, the host performs a Bernoulli trial, which computes the value of an overall assignment function to decide whether to accept a VM or not. The assignment function considers the resource utilization and utilization threshold, then calculates the probability to accept the incoming VM. If the trial succeeds, the host sends back a message to the manager. The manager collects all messages and selects one available host to allocate the incoming VM. The second part of EcoCloud is VM migration, which utilizes two Bernoulli trial based functions to handle the over-utilized and under-utilized situations separately. If the trial of the under-utilized situation is successful, a randomly selected VM will be migrated. On the other hand, migrations with over-utilized will migrate the VM that decreases resource utilization to be lower than the over-utilized threshold. The target host to accept the VM is decided by another assignment function with a slight modification of a parameter.

EcoCloud can be the representative algorithm aiming to reduce energy while minimizing VM migrations. The main advantage of EcoCloud is that the approach can reduce downtime duration and transmission bandwidth. Furthermore, hosts can make migration decisions by themselves, thus the pressure of VM manager is relieved. The limitation of this work is that it investigates the effect of power consumption on SLAs while the effects of VM migrations on SLAs are not evaluated.

C. GRANITE

GRANITE [20] is a holistic virtual machine scheduling algorithm capable of minimizing total data center energy consumption, including computing energy and cooling energy. This work considers Computer Room Air Conditioner (CRAC) as the only cooling devices and constructed their models. Based on server models and cooling models, GRANITE utilizes a greedy algorithm to conduct the VM initial placement and dynamic migration. They assume that the users' resource demand can be predicted. In the initial placement stage, the greedy algorithm is applied in GRANITE for all VMs to select the host with the least increment of total energy after the placement. The CRAC will be adjusted if the CPU temperature is above the threshold. In the dynamic VM consolidation stage, the algorithm aims to balance the workloads and cooling energy consumption. The GRANITE defines a dynamic temperature threshold and checks host status. If the host temperature is above the temperature threshold, a set of VMs will be migrated to other hosts. The target host selection for migration is the same as the greedy algorithm in the initial placement stage.

GRANITE can present the algorithms that consider the holistic management of energy in cloud data centers. The core idea of the algorithm is similar to MBFD algorithm, while it considers cooling power, which leads to more accurate and holistic scheduling results. The advantage of this work is that it combines the server status and data center thermal control to form the fine-grained energy efficient scheduling. However, the model accuracy can be further improved by using a 3-Dimensional computational fluid dynamics model rather than 2-Dimensional one.

D. LOAD

LOAD [21] is an algorithm based on learning automata for energy and SLA efficient consolidation of VMs in cloud data centers. The proposed algorithm considers the demanded resource of users to predict overloaded hosts. By preventing overloaded hosts and shutting down idle hosts, the proposed algorithm aims to save energy consumption of data centers. The learning automata-based overload detection enhances the VM consolidation by predicting CPU usage of hosts upon resource usage history. Each VM is equipped with one automata including 3 actions, increasing CPU utilization, reducing CPU utilization and no changing of CPU utilization. At the beginning, the 3 actions are with equal probability. After the beginning stage, the reward and penalty of learning automata will be updated based on the environment. The automata selects one of the three actions in each iteration based on the probability. And in the next iteration, if the automata's decision is right, the action will be rewarded, otherwise, the action will be penalized. The learning automata is applied to estimate usage of VMs on the host. If the estimation shows that the host may be overloaded, the VMs will be migrated and other VMs will be prevented to be migrated to the current host. The migrated destination is based on the Best Fit Decreasing (BFD) algorithm [23]. The simulation results show that leveraging

learning-based prediction can reduce energy consumption of data centers.

LOAD is a typical algorithm applying learning techniques to optimize VM consolidation. This work advances the existing work by considering the dynamic prediction for resource usage. But the limitation is that this work predicts the overloaded situations and does not handle the under-utilized situations.

E. ACS

ACS [22] is an online optimization meta-heuristic algorithm based on ant colony optimization and VM consolidation to achieve the near-optimal solution. Its objective is to balance the energy consumption, the number of VM migrations and QoS concerning performance. In this approach, the authors formulate the energy efficient VM consolidation as a multi-objective optimization problem to optimize multiple metrics simultaneously. To leverage ACO, the necessary entities, including pheromone updating rules and probabilistic decision rules are defined. If one solution has more pheromone trails, the VM has a larger probability to be placed on the host. ACS also has updating rules for local and global pheromone, which are applied in each iteration. In iterations, the local pheromone is updated by ant when they perform a movement. After all ant construct their solutions locally, the global pheromone update is conducted as the migration process, and only the dominated place will be kept. The process is repeated until it reaches the maximum iteration number.

ACS represents a set of meta-heuristic algorithms proposed to balance multiple objectives. The results based on simulations show that the proposed algorithm can reduce energy consumption and VM migrations while guaranteeing QoS. The performance can be further improved by investigating the impact of VM migration on network.

In summary, the investigated algorithms all follow the two-phase energy efficient scheduling process as shown in Fig. 1. The investigated algorithms apply different techniques to optimize the placement of VMs. Most of the investigated algorithms focus on the optimization of VM consolidation phase except EcoCloud spends more effort on the initial placement based on probabilistic approach. In the algorithms focusing on the VM consolidation optimization, MBFD applies heuristic algorithms based on bin-packing modelling, while EcoCloud utilizes a probabilistic approach to find the host with the highest probability to accept migrated VM. GRANITE considers energy and performance together by modelling cooling energy consumption and VM performance degradation. ACS applies a meta-heuristic algorithm to find the optimized consolidation solutions.

IV. ARCHITECTURE AND MODELLING

This section discusses the investigated algorithms from the architecture and modelling perspectives, and the algorithm complexity is also discussed.
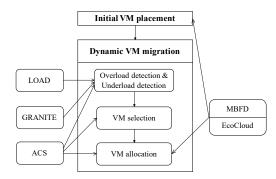

Fig. 2: The focused scheduling phase comparison of investigated algorithms

A. MBFD

Architecture: The four main components of the green cloud architecture are: broker, green service allocator, Virtual Machines (VMs) and Physical Machines (PMs). Broker enables the user interaction module to submit the workloads and their QoS requirements from any geographical distribution allocation. PM based hardware infrastructure creates virtualized resources (VMs). VMs are consolidated to fulfill the demand of workloads dynamically using DVFS. Green service allocator incorporates the energy manager and VM manager to allocate the virtual resources to user workloads based on their requirements for their execution at runtime.

Model: Equation (1) defines the power model for this research work.

$$P(u) = k \times P_{max} + (1 - k) \times P_{max} \times u \qquad (1)$$

where P_{max} is the maximum consumption of power while fully utilization of server; k is the small amount of power used by idle server; and u is the CPU utilization. The value of energy consumption E of a PM is defined in Equation (2). $u(t)$ is CPU utilization, which is a function of time.

$$E = \int_{t2}^{t1} P(u(t))dt \qquad (2)$$

where $t1$ is start time of Task T and $t2$ is end time of Task T.

B. EcoCloud

Architecture: The architecture of EcoCloud contains two main probabilistic procedures: 1) VM assignment and VM migration. Based on the availability of CPU and RAM on different servers of cloud data centers, VM allocation is performed. EcoCloud allocates the VM to the newly submitted application and sends an invitation to all the servers participated to find the best one to place this VM.

Model: Equation (3) defines the linear model used in this research for power consumption and it is expressed as:

$$P(u) = P_{idle} + (P_{max} - P_{idle}) \times u \qquad (3)$$

where P_{max} is the maximum consumption of power while fully utilization of server; P_{idle} is the small amount of power used by idle server; and u is the CPU utilization.

C. GRANITE

Architecture: It comprises three sub-components including workload manager, scheduling manager and cooling manager. Workload manager manages the workloads submitted by users and process for scheduling based on their requirements. Scheduling manager schedules the resources for the execution of workloads while maximizing the performance of data centers and minimizes the consumption of energy. Cooling manager maintains the temperature of data centers and saves cooling energy by performing VM placement and dynamic migration in an efficient manner.

Model: Equation (4) defines a linear power model to find the energy consumption of data centers, which is a combination of computing energy and cooling energy.

$$Energy = Energy_{computing} + Energy_{cooling} \quad (4)$$

D. LOAD

Architecture: The system architecture contains four sub-components such as user portal (to submit workloads), global manager (which is an intermediate between user portal and local manager), local manager (which manages the PMs and VMs, and it is controlled by a single global manager) and VM manager (which manages the virtual resources).

Model: Equation (1) is used to calculate the energy consumption for this research work, which represents the linear model of power and CPU utilization.

E. ACS

Architecture: In the architecture of ACS, two types of agents are included: local and global agents. The local agent is deployed in a host to solve the host status detection sub-problem by monitoring the host resource utilization. The global agent is responsible for supervising and optimizing the VM placement by taking advantage of the ACO-based algorithm.

Model: ACS utilizes the linear power model as MBFD in Equation 2.

We note that the investigated algorithms have different focuses on the phases of VM consolidation, and we show the focuses in Fig. 2. As introduced, the VM consolidation process can be mainly divided into two phases: the initial VM placement and dynamic VM migration. And for the VM migration, the overloads and under-utilized detection, VM selection and VM allocation are included. MBFD and EcoCloud spend more effort on the VM placement phase, including the initial VM placement and VM allocation in Dynamic VM migration. On the other hand, LOAD focuses on overloads detection by predicting the utilization via learning automaton to optimize overload detection. GRANITE takes the CPU temperature into consideration and optimizes the overload detection policy. ACS improves all the phases in dynamic VM consolidation. It utilizes LiRCUP [24] to forecast overloads of servers, and adopts Ant Colony System to find the near-optimal solutions for VM selection and allocation.

In summary, from the perspective of architecture, the investigated algorithms all adopt the layered architecture, and the layers can be mainly divided into three parts. The bottom layer is the resource provisioner, which provides physical or virtual resources. The middle layer is responsible for the energy efficient scheduling, which handles the VM management and provides the energy efficient scheduling algorithms. At the top layer, the users' request and optimization goals are configured.

From the modelling perspective, all algorithms adopt the linear power model. As for power components, all algorithms include the CPU energy consumption, MBFD, EcoCloud and ACS consider the memory part, and GRANITE utilizes a more comprehensive model including storage, network and cooling.

Algorithm complexity analysis: MBFD, GRANITE and LOAD are based on heuristic algorithms, and the complexity of them are all $M \times N$, where N is the number of hosts and M is the number of VMs. The complexity of EcoCloud equals to the number of hosts to accept VMs, which is N. Based on meta-heuristic with iterations, the complexity of ACS is $M \times N \times A \times I$, where A is the number of ants that concurrently build their migration plans, and I is the number of iterations.

V. METRICS

For the objective of energy efficiency, energy consumption is the major metric to be evaluated. However, the algorithms are also making trade-offs between the energy consumption and other metrics, such as SLA violations. In this section, we discuss the adopted metrics in our investigated energy efficient algorithms. Note that the investigated algorithms use some similar metrics while having some other additional metrics. Here we aim to cover the metrics applied in these algorithms and identify the differences between these metrics. Table II summaries the algorithms with their corresponding adopted metrics.

Metrics for energy efficiency

Total energy consumption: It is the total energy consumption consumed by physical machines in the data centers. It is derived from the power model in Equation (2).

Number of active servers: It represents how many servers are running as active during the observation time. The value should be minimized, and thus more idle serves can be switched into low-power mode.

Metrics for SLA

SLA violation percentage (SLAV) [18]: It is the percentage of SLA violations events relatively to the total number of the processed time frame. The SLA violation is identified when a given VM cannot get the amount of MIPS as requested.

VM migration times: The number of migrations triggered by the algorithm during the VM scheduling process.

Average SLA violation: It is the average CPU performance that has not been allocated to an application when requested, resulting in performance degradation.

In summary, we can notice that several metrics have been adopted for evaluations by more than one algorithm, including total energy, SLA violation percentage, VM migrations, average SLA violations and number of active hosts. To make our

TABLE II: Metrics adopted in compared algorithms

Metrics	Optimization Objective	Adopted Algorithms
Total Energy Consumption	Minimization	All
SLA violation percentage	Minimization	All
VM migrations	Minimization	BMDP, LOAD, EcoCloud, ACS
Average SLA violation	Minimization	BMDP, LOAD, ACS
Number of active hosts	Minimization	EcoCloud, GRANITE

evaluations more comparable from the metrics perspective, we evaluate these metrics in the performance evaluations section.

VI. PERFORMANCE EVALUATIONS

In this section, to compare the performance of investigated algorithms, we conduct experiments for the five well-known and investigated algorithms based on different performance metrics and two traces. We also include one provided algorithm in CloudSim as baselines, the Interquartile Range (IQR) [11] that manages a dynamic threshold for overload detection.

A. Experiments Settings

For host capacity, each host has two CPU cores with millions of instructions per second (MIPS) of 1880 or 2660, 4 GB of RAM and 1 TB of storage. We use the power model derived from HP ProLiant ML110 G4 or HP ProLiant ML110 G5, which has been used in [18] and [22]. According to the model, the energy consumption of the host with different utilization is shown in Table III. For VM configurations, we consider four types of VMs with MIPS of 500, 1000, 1500 and 2500, and the VMs number of each type is randomly generated. The detailed specifications of hosts and VMs are shown in Table IV.

We first carry out several experiments under synthetic workloads, and then, to simulate the real cloud data center, we utilize the real workload data of the CoMon project provided by PlanetLab [25]. The workload includes CPU utilization data of thousands of VMs allocated to servers that are located in more than 500 places around the world. And the data is collected every five minutes for 10 days, which represents the workload in the real cloud environment.

We select four metrics to evaluate the performance of these algorithms, including Energy Consumption, SLA violation percentage (SLAV), VM migration times and the number of active hosts. We choose these metrics as they have been adopted widely and also used in more than one algorithms as we discussed in Section V. Due to the page limitation, we evaluate SLAV to represent the SLA violations instead of average SLA violations.

B. Implementation Details

The configurations in EcoCloud include a shape parameter p, α and β in its probability function. We set $p = 3$, $\alpha = \beta = 0.25$ in our experiments, following the configurations of the original paper. For LOAD, the reward and penalty parameters for updating learning automaton, a and b, are both configured to 0.1. We utilize the parameters in the original paper of ACS, however, we don't have a training data set to estimate the initial pheromone level. Therefore, we set M to be the total

number of selected VMs, and P is configured as the number of under-utilized servers. The configured parameters of MBFD and GRANITE are the same as the original paper.

C. Synthetic Workloads

To show the performance under synthetic workloads, the first part of the evaluations focuses on the experiments under synthetic workloads, which has been used in MBFD and ACS. We firstly generate the same number of hosts and VMs and vary the lower utilization threshold that defines when the host is under-utilized. And we varied the threshold from 0.1 to 0.5 with increment as 0.1. The higher utilization threshold is set as 0.4 more than the lower utilization threshold. We follow the configuration of the utilization threshold interval in [18]. Under each configuration, we randomly generate workloads to run the experiments and repeat 10 times. Both the number of hosts and VMs are set as 50.

Fig. 3 shows the average results of the experiments under synthetic workloads. Based on the results, the higher values of lower utilization threshold can enable all the algorithms to achieve better energy efficiency, fewer VM migrations and more SLAV. To be more specific, ACS performs the best on energy consumption by reducing 21.1% power compared with MBFD. EcoCloud requires less than 600 migration times, which are much fewer than other algorithms. The reason that ACS achieves the best energy efficiency is that it has the lowest number of active hosts. For SLA violation comparison, these algorithms perform better when the lower threshold is set to be 0.1.

As setting the lower utilization threshold with 0.5 can achieve the best performance for all the algorithms, we fix the lower utilization threshold as 0.5, and run the experiments repeatedly with 10 times to show the variance of results, which are shown in Fig. 4. We can notice that ACS can achieve the best performance in energy consumption with 31.2 kWh, and EcoCloud can reduce the VM migrations as 420.5 in average.

To investigate the impacts of different numbers of hosts and VMs, we also fix the lower utilization threshold as 0.5, but configure the ratio of the number of servers to the number of VMs to be 1:1, 1:1.25, 1:1.5 and 1:1.75 respectively. With each ratio, the experiments are repeated 10 times, and the results as shown in Fig 5. We can observe that energy consumption increases significantly when there are more VMs. ACS is the most energy efficient algorithm and consumes 31.2 KWh and 31.5 kWh when the ratio is 1:1 and 1:1.75 respectively, which is about 21.2-34.4% less compared to MBFD. EcoCloud achieves the best results in VM migrations. Although the investigated algorithms can reduce more energy than IQR, more SLAV happens in these algorithms.

D. PlanetLab Workloads

To show the algorithm performance under realistic traces, we also conduct experiments under PlanetLab workloads. We vary the lower CPU utilization threshold varies from 0.1 to 0.5, and the interval between the lower threshold and the higher threshold is fixed at 0.4. The number of hosts is configured

TABLE III: Energy consumption with different CPU utilization in Watts

Server	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
HP ProLiant G4	86	89.4	92.6	96	99.5	102	106	108	112	114	117
HP ProLiant G5	93.7	97	101	105	110	116	121	125	129	133	135

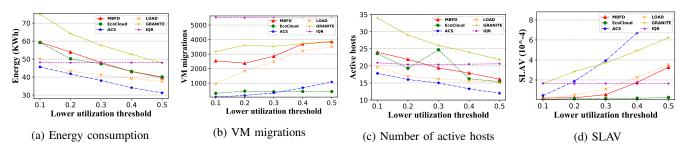| (a) Energy consumption	(b) VM migrations	(c) Number of active hosts	(d) SLAV

Fig. 3: Performance comparison of algorithms under Synthetic workloads (ratio of hosts and VMs number is 50:50)

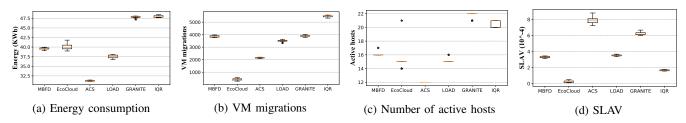| (a) Energy consumption	(b) VM migrations	(c) Number of active hosts	(d) SLAV

Fig. 4: Performance comparison of algorithms under Synthetic workloads with lower utilization threshold as 0.5 (ratio of hosts and VMs number is 50:50)

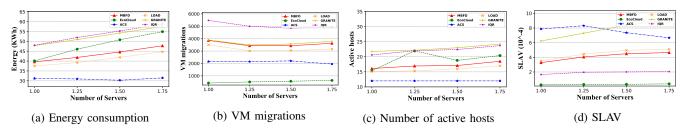| (a) Energy consumption	(b) VM migrations	(c) Number of active hosts	(d) SLAV

Fig. 5: Performance comparison of algorithms under Synthetic workloads (setting hosts number = 50 and varying ratios of hosts and VMs number are 1:1, 1:1.25, 1:1.5, and 1:1.75)

TABLE IV: Host / VM Types and Capacity

Name	CPU (MIPS)	Cores	Memory	Bandwidth	Storage
Host Type 1	1.86 GHz	2	4 GB	1 Gbit/s	1 TB
Host Type 2	2.66 GHz	2	4 GB	1 Gbit/s	1 TB
VM Type 1	2.5 GHz	1	870 MB	100 Mbit/s	1 GB
VM Type 2	2.0 GHz	1	1740 MB	100 Mbit/s	1 GB
VM Type 3	1.0 GHz	1	1740 MB	100 Mbit/s	1 GB
VM Type 4	0.5 GHz	1	613 MB	100 Mbit/s	1 GB

as 800 and the number of VMs is retrieved from PlanetLab traces. The average results are shown in Fig. 6 by running 10 times of experiments, and each experiment is with a one-day PlanetLab trace.

Fig. 6a shows the energy consumption comparison, in which the power decreases when having a larger value of the lower utilization threshold. ACS consumes the least energy compared with other algorithms with 148.7 kWh when the lower utilization threshold is 0.4. EcoCloud consumes more energy consumption than MBFD because it keeps more servers to be active. The numbers of VM migrations are compared in Fig. 6b. EcoCloud achieves an apparent reduction in this metric and only needs 1251 migrations on average when the lower threshold is set to be 0.5. LOAD achieves improvement in the number of migrations and reduces 12.4% migrations compared with MBFD. GRANITE acquires better results when the lower utilization threshold increases. The number of active hosts comparison is shown in Fig. 6c, and ACS can shut down the maximum number of hosts. Fig. 6d shows the SLA violation percentage comparison, and the SLA violation percentage increases as the lower utilization threshold increases. As the figure shows, LOAD and ACS perform worse on this metric compared with MBFD and EcoCloud. EcoCloud maintains a low SLA violation percentage and thus ensures the quality of services.

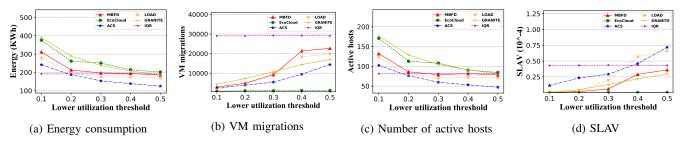We set the lower utilization threshold as 0.5 and the higher

(a) Energy consumption   (b) VM migrations   (c) Number of active hosts   (d) SLAV

Fig. 6: Performance comparison of algorithms under PlanetLab workloads

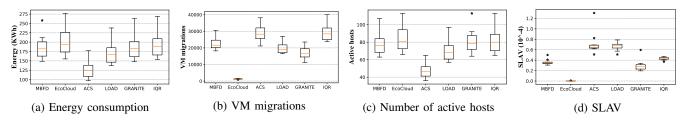(a) Energy consumption   (b) VM migrations   (c) Number of active hosts   (d) SLAV

Fig. 7: Performance comparison of algorithms under PlanetLab workloads with lower utilization threshold as 0.5

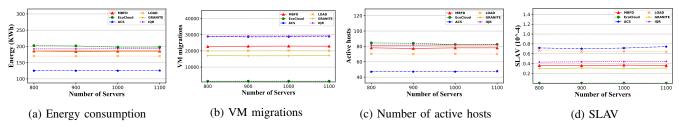(a) Energy consumption   (b) VM migrations   (c) Number of active hosts   (d) SLAV

Fig. 8: Performance comparison of algorithms under PlanetLab workloads (varied number of servers with 800, 900, 1000 and 1100)

utilization threshold as 0.9, and run 10 times experiments repeatedly to show the variance of performance results as shown in Fig. 7. In Fig. 7a, we can notice that ACS and LOAD perform better in energy consumption compared with other baselines. The average energy consumption of ACS is 125.1 KWh while the MBFD, EcoCloud, and GRANITE consume more than 180 KWh. Fig. 7c shows the average results of the number of active hosts, and ACS can achieve the best results with 48 in average. Fig. 7b demonstrates the comparison of VM migrations, and EcoCloud focuses on optimizing this metric and reduces the number of VM migration to be under lower than 2000. Compared with MBFD and GRANITE, LOAD reduces the VM migrations. The SLA violation percentage comparison is presented in Fig 7d, and EcoCloud ensures SLA in a the best manner with 0.01×10^{-4}. Although ACS saves more energy, it performs worst in reducing SLA violations with 0.71×10^{-4}.

To evaluate the performance with a different number of hosts, We also vary the number of hosts from 800 to 1100 in our experiments as shown in Fig. 8. As the number of hosts increases, we can notice that ACS always consumes the minimum power, and the least number of VM migration happens in EcoCloud. ACS only keeps 48 active hosts and leads to the maximum SLAV as 0.72×10^{-4}.

In conclusion, we can notice that under both workloads, ACS can achieve the best energy efficiency in most cases as it also has the least number of active hosts. EcoCloud can reduce more VM migrations compared with other algorithms and trigger the least SLA violations under PlanetLab workloads. We can notice that meta-heuristic algorithms, e.g. ACS, can achieve better performance in energy compared with heuristic algorithms as it searches larger solution space. As the pioneer of VM consolidation-based energy efficient algorithm for cloud data centers, MBFD has been widely accepted as it is easy to implement and efficient. Although the performance of MBFD has been outperformed by recent algorithms, the main idea of MBFD has been referred, such as in GRANITE and LOAD, where the holistic energy and SLA violation are optimized respectively. GRANITE is suggested for the case that aims to optimize more energy-consuming components rather than only the hosts. EcoCloud is suitable for the scenarios that VM migrations should be minimized, such as the network is the bottleneck of the whole system, as it can reduce VM migrations significantly. LOAD can achieve good performance if the future resource usage can be accurately predicted, thus it suits the conditions that the system has adequate history resource usage data or resource usage has a typical pattern, like Wikipedia traces.

VII. Conclusions and Future Work

In this paper, we present an investigation of 5 state-of-the-art energy efficient algorithms based on VM consolidations for cloud data centers. We discuss and compare the algorithms from multiple perspectives, including core ideas, architectures, modelling and algorithm complexity. We also implement these algorithms in CloudSim and configure the parameters as noted in these algorithms. By conducting experiments under both synthetic and real traces, the results show that these algorithms can efficiently reduce the energy consumption while balancing the trade-offs between energy and some other metrics, such as VM migrations and SLA violations.

Based on our investigation, some future research directions are identified as follows:

- Utilization threshold settings have important impacts on energy consumption, thus self-adaptive threshold configuration approach can be further investigated.
- Most VM consolidation-based energy efficient algorithms consider the CPU resource as the bottleneck and the key factor in their power models. Resource usage and energy consumption related to other resources, e.g. network, can be further considered.
- Meta-heuristic algorithms have shown their better performance in improving energy consumption compared with traditional heuristic algorithms, however, the related time cost should be reduced, which can be achieved by reducing the solution space.
- Evaluations can be conducted under recently published workloads, such as Google traces and Alibaba traces.

References

[1] R. Buyya, S. N. Srirama, G. Casale, R. Calheiros, and et al., "A manifesto for future generation cloud computing: Research directions for the next decade," *ACM Computing Surveys*, vol. 51, no. 5, pp. 105:1–105:38, Nov. 2018.

[2] T. Kaur and I. Chana, "Energy efficiency techniques in cloud computing: A survey and taxonomy," *ACM Computing Surveys (CSUR)*, vol. 48, no. 2, pp. 1–46, 2015.

[3] S. S. Gill, S. Tuli, M. Xu, I. Singh, K. V. Singh, D. Lindsay, S. Tuli, D. Smirnova, M. Singh, U. Jain, H. Pervaiz, B. Sehgal, S. S. Kaila, S. Misra, M. S. Aslanpour, H. Mehta, V. Stankovski, and P. Garraghan, "Transformative effects of iot, blockchain and artificial intelligence on cloud computing: Evolution, vision, trends and open challenges," *Internet of Things*, vol. 8, p. 100118, 2019.

[4] (2014) Data center energy: Reducing your carbon footprint — data center knowledge. [Online]. Available: http://www.datacenterknowledge.com/archives/2014/12/17/undertaking-challenge-reduce-data-center-carbon-footprint

[5] S. S. Gill and R. Buyya, "A taxonomy and future directions for sustainable cloud computing: 360 degree view," *ACM Computing Surveys*, vol. 51, no. 5, pp. 104:1–104:33, Dec. 2019.

[6] R. Ghosh, S. P. R. Komma, and Y. Simmhan, "Adaptive energy-aware scheduling of dynamic event analytics across edge and cloud resources," in *2018 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)*, May 2018, pp. 72–82.

[7] O. Soualah, M. Mechtri, C. Ghribi, and D. Zeghlache, "Energy efficient algorithm for vnf placement and chaining," in *Proceedings of the 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, ser. CCGrid '17. Piscataway, NJ, USA: IEEE Press, 2017, pp. 579–588. [Online]. Available: https://doi.org/10.1109/CCGRID.2017.84

[8] M. Xu, W. Tian, and R. Buyya, "A survey on load balancing algorithms for virtual machines placement in cloud computing," *Concurrency and Computation: Practice and Experience*, vol. 29, no. 12, pp. 4123–4138, 2017.

[9] M. Xu and R. Buyya, "Brownout approach for adaptive management of resources and applications in cloud computing systems: A taxonomy and future directions," *ACM Comput. Surv.*, vol. 52, no. 1, pp. 8:1–8:27, Jan. 2019.

[10] W. Tian, M. Xu, A. Chen, G. Li, X. Wang, and Y. Chen, "Open-source simulators for cloud computing: Comparative study and challenging issues," *Simulation Modelling Practice and Theory*, vol. 58, pp. 239 – 254, 2015, special issue on Cloud Simulation.

[11] R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. De Rose, and R. Buyya, "Cloudsim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms," *Software: Practice and Experience*, vol. 41, no. 1, pp. 23–50, 2011.

[12] B. Wickremasinghe, R. N. Calheiros, and R. Buyya, "Cloudanalyst: A cloudsim-based visual modeller for analysing cloud computing environments and applications," in *2010 24th IEEE international conference on advanced information networking and applications*. IEEE, 2010, pp. 446–452.

[13] S. K. Garg and R. Buyya, "Networkcloudsim: Modelling parallel applications in cloud simulations," in *2011 Fourth IEEE International Conference on Utility and Cloud Computing*. IEEE, 2011, pp. 105–113.

[14] Y. Mansouri, A. N. Toosi, and R. Buyya, "Data storage management in cloud environments: Taxonomy, survey, and future directions," *ACM Computing Surveys (CSUR)*, vol. 50, no. 6, pp. 1–51, 2017.

[15] A.-C. Orgerie, M. D. d. Assuncao, and L. Lefevre, "A survey on techniques for improving the energy efficiency of large-scale distributed systems," *ACM Computing Surveys (CSUR)*, vol. 46, no. 4, p. 47, 2014.

[16] Z. Á. Mann, "Allocation of virtual machines in cloud data centers—a survey of problem models and optimization algorithms," *Acm Computing Surveys (CSUR)*, vol. 48, no. 1, p. 11, 2015.

[17] 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*, vol. 52, pp. 11–25, 2015.

[18] A. Beloglazov, J. Abawajy, and R. Buyya, "Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing," *Future Generation Computer Systems*, vol. 28, no. 5, pp. 755–768, 2012.

[19] C. Mastroianni, M. Meo, and G. Papuzzo, "Probabilistic consolidation of virtual machines in self-organizing cloud data centers," *IEEE Transactions on Cloud Computing*, vol. 1, no. 2, pp. 215–228, 2013.

[20] X. Li, P. Garraghan, X. Jiang, Z. Wu, and J. Xu, "Holistic virtual machine scheduling in cloud datacenters towards minimizing total energy," *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, no. 6, pp. 1317–1331, 2018.

[21] M. Ranjbari and J. A. Torkestani, "A learning automata-based algorithm for energy and sla efficient consolidation of virtual machines in cloud data centers," *Journal of Parallel and Distributed Computing*, vol. 113, pp. 55–62, 2018.

[22] 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," *IEEE Transactions on Services Computing*, vol. 8, no. 2, pp. 187–198, March 2015.

[23] M. Yue, "A simple proof of the inequality ffd (l) ¡ 11/9 opt (l) + 1, for all the ffd bin-packing algorithm," *Acta Mathematicae Applicatae Sinica*, vol. 7, no. 4, pp. 321–331, Oct 1991.

[24] F. Farahnakian, P. Liljeberg, and J. Plosila, "Lircup: Linear regression based cpu usage prediction algorithm for live migration of virtual machines in data centers," in *2013 39th Euromicro Conference on Software Engineering and Advanced Applications*. IEEE, 2013, pp. 357–364.

[25] K. Park and V. S. Pai, "Comon: a mostly-scalable monitoring system for planetlab," *ACM SIGOPS Operating Systems Review*, vol. 40, no. 1, pp. 65–74, 2006.