

An Experimental Evaluation of Energy-Efficient Green Cloud Computing Strategies

Saurabh Prashant Maske

Assistant Professor (CHB) Department of Computer Science, G.S. Gawande Mahavidyalaya, Umarched,
Dist. Yavatmal
spmaske01@gmail.com

Abstract

The rapid expansion of cloud computing has led to the large-scale deployment of data centers, resulting in high energy consumption and increased environmental concerns. Traditional cloud management approaches mainly emphasize performance optimization, often overlooking energy efficiency and sustainability. Green cloud computing addresses these challenges by introducing energy-aware strategies that aim to minimize power consumption while maintaining acceptable system performance.

This paper presents an experimental performance study of green cloud computing techniques for energy optimization. The study evaluates multiple energy-efficient strategies, including dynamic virtual machine consolidation, energy-aware resource allocation, idle resource management, and dynamic voltage and frequency scaling. Experiments are conducted using the CloudSim toolkit to simulate a realistic cloud environment with heterogeneous virtual machines and varying workloads.

The performance of green cloud computing techniques is assessed using key metrics such as total energy consumption, average response time, and resource utilization. Experimental results show that the adoption of green strategies significantly reduces energy consumption compared to traditional cloud resource management approaches, while maintaining satisfactory Quality of Service levels. The findings highlight the importance of integrating energy-aware techniques to achieve sustainable and cost-effective cloud computing environments. This study contributes valuable insights for researchers and practitioners aiming to design energy-efficient cloud infrastructures.

Keywords

Green Cloud Computing; Energy Optimization; Energy-Efficient Resource Management; Cloud Computing; Virtual Machine Consolidation; Sustainable Computing.

1. Introduction

Cloud computing has emerged as a dominant paradigm in modern information technology by enabling on-demand access to computing resources such as servers, storage, networks, and applications through the Internet. The cloud model offers several advantages, including scalability, flexibility, cost efficiency, and ease of management, making it widely adopted by enterprises, academic institutions, and government organizations. With the rapid growth of cloud-based services, data centers have expanded significantly in size and number to meet increasing computational demands. However, this expansion has resulted in a substantial rise in energy consumption, leading to high operational costs and serious environmental concerns.

Data centers are among the largest consumers of electrical energy in the information and communication technology sector. A significant portion of this energy is consumed not only by computing resources but also by cooling systems and auxiliary infrastructure. Inefficient resource utilization, over-provisioning of virtual machines, and continuous operation of underutilized servers contribute to unnecessary power wastage. Traditional cloud management strategies

primarily focus on performance-related metrics such as response time, throughput, and availability, often ignoring energy efficiency as a critical design objective. As a result, cloud service providers face increasing challenges in balancing performance requirements with energy consumption and sustainability goals.

In recent years, the concept of green cloud computing has gained considerable attention as a solution to address the energy-related challenges of cloud environments. Green cloud computing aims to reduce the environmental impact of cloud data centers by optimizing energy consumption while maintaining acceptable Quality of Service (QoS). This is achieved through the adoption of energy-aware resource management techniques, efficient workload distribution, and intelligent scheduling mechanisms. By minimizing energy wastage, green cloud computing not only reduces operational costs but also contributes to environmental sustainability by lowering carbon emissions.

Several energy-efficient techniques have been proposed to improve the sustainability of cloud data centers. These include dynamic virtual machine consolidation, which reduces the number of active physical servers by migrating workloads to fewer machines; energy-aware resource allocation, which assigns tasks based on energy consumption profiles; idle resource management, which powers down unused servers or virtual machines; and dynamic voltage and frequency scaling, which adjusts processor operating parameters according to workload demands. While these techniques show promise in reducing energy consumption, their impact on overall system performance must be carefully analyzed to ensure that service quality is not compromised.

Despite the growing interest in green cloud computing, many existing studies focus on individual energy optimization techniques or theoretical models, with limited experimental validation under realistic workload conditions. Additionally, some approaches emphasize energy savings without adequately considering performance trade-offs such as increased response time or reduced resource availability. This creates a research gap in understanding the combined effect of multiple green cloud computing techniques on both energy efficiency and system performance. A comprehensive experimental evaluation is therefore necessary to assess the effectiveness of green strategies in real-world cloud scenarios.

Simulation-based experimentation has become a widely accepted approach for evaluating cloud computing strategies due to the complexity and cost associated with real data center deployment. Cloud simulation tools allow researchers to model large-scale cloud environments, generate diverse workloads, and analyze performance metrics in a controlled setting. Among these tools, CloudSim has emerged as a popular platform for simulating cloud infrastructures and evaluating resource management policies. CloudSim provides support for modeling virtual machines, hosts, data centers, and energy consumption, making it suitable for studying green cloud computing techniques.

This research focuses on conducting a detailed performance study of green cloud computing techniques for energy optimization using a simulated cloud environment. The study evaluates multiple energy-aware strategies, including dynamic virtual machine consolidation, energy-aware resource allocation, idle resource management, and dynamic voltage and frequency scaling. These techniques are analyzed under varying workload conditions to examine their effectiveness in reducing energy consumption while maintaining acceptable system performance. Performance evaluation is carried out using metrics such as total energy consumption, average response time, and resource utilization.

The primary motivation of this study is to provide a balanced analysis of energy efficiency and performance in cloud computing environments. By experimentally comparing traditional cloud management approaches with green cloud computing techniques, this research aims to highlight the trade-offs involved in achieving sustainable cloud operations. The findings of this study can assist cloud service providers, system designers, and researchers in selecting appropriate energy-aware strategies that align with both performance and sustainability objectives.

The contributions of this research are threefold. First, it provides a comprehensive experimental evaluation of multiple green cloud computing techniques within a unified simulation framework. Second, it offers a comparative analysis of

energy consumption and performance metrics, enabling a clear understanding of the benefits and limitations of green strategies. Third, the study contributes to the growing body of research on sustainable cloud computing by offering practical insights into energy optimization techniques that can be adopted in real-world cloud environments.

2. Related Work

The rapid increase in energy consumption by cloud data centers has attracted significant attention from researchers over the past decade. As cloud computing continues to expand, concerns related to power usage, operational costs, and environmental sustainability have motivated the development of green cloud computing techniques. This section reviews existing research related to energy-efficient cloud computing, green resource management strategies, and performance-aware energy optimization approaches.

Early studies on cloud computing primarily focused on improving performance, scalability, and reliability. Resource allocation, task scheduling, and load balancing algorithms were designed mainly to minimize response time, maximize throughput, and ensure service availability. Although these approaches improved Quality of Service, they largely ignored the issue of energy consumption. As data centers grew in size, researchers began to recognize that energy efficiency is a critical factor affecting both economic and environmental aspects of cloud operations.

One of the most widely studied energy optimization techniques in cloud computing is virtual machine consolidation. VM consolidation aims to reduce the number of active physical servers by migrating workloads from lightly loaded machines to fewer servers, allowing idle hosts to be powered down or placed in low-energy states. Several studies have demonstrated that dynamic VM consolidation can significantly reduce energy consumption while maintaining acceptable performance levels. However, frequent VM migrations may introduce overheads such as increased network traffic, performance degradation, and service interruptions. Researchers have proposed various heuristics and threshold-based policies to balance energy savings and migration costs, but achieving optimal consolidation remains a challenging problem.

Energy-aware resource allocation has also been extensively investigated as a means of improving cloud sustainability. In this approach, resource allocation decisions are made based on energy consumption characteristics of physical hosts and virtual machines. Researchers have proposed algorithms that allocate tasks to servers with lower power usage or better energy efficiency profiles. Some studies incorporate power models to estimate energy consumption during allocation decisions. While these approaches show promising results in reducing power usage, they often require accurate energy models and may increase computational complexity.

Idle resource management is another important area of green cloud research. Cloud data centers frequently operate under partial load conditions, leading to significant energy wastage due to idle or underutilized servers. To address this issue, researchers have proposed techniques that detect idle resources and transition them into low-power or sleep modes. These techniques can effectively reduce static power consumption, but they must be carefully managed to avoid delays when resources are required again. Studies have shown that improper idle management may lead to increased response times or reduced availability during workload surges.

Dynamic Voltage and Frequency Scaling (DVFS) has been widely adopted in processor-level energy optimization and has been extended to cloud computing environments. DVFS adjusts the operating voltage and frequency of processors according to workload demands, thereby reducing energy consumption during low utilization periods. Several studies have integrated DVFS with cloud scheduling and resource management strategies. Results indicate that DVFS can achieve substantial energy savings, particularly in CPU-intensive workloads. However, aggressive frequency scaling may negatively impact application performance, highlighting the need for performance-aware DVFS policies.

In addition to individual energy optimization techniques, researchers have explored integrated and hybrid approaches for green cloud computing. These approaches combine multiple strategies such as VM consolidation, DVFS, and energy-aware scheduling to achieve better energy savings. Hybrid models aim to exploit the strengths of different techniques while mitigating their individual limitations. Experimental studies have shown that integrated approaches often outperform single-technique solutions in terms of energy efficiency. Nevertheless, designing effective hybrid models introduces complexity and requires careful coordination among different components.

Simulation-based evaluation has played a crucial role in green cloud computing research. Due to the high cost and complexity of real-world data center experimentation, researchers commonly use cloud simulation tools to model large-scale environments and test energy optimization strategies. Simulation platforms allow controlled experimentation with diverse workloads, system configurations, and energy models. Many studies rely on simulation results to demonstrate the effectiveness of proposed techniques. However, the accuracy of simulation outcomes depends heavily on the realism of workload models and energy consumption assumptions.

Several studies have emphasized the importance of evaluating performance trade-offs associated with energy optimization. While reducing energy consumption is desirable, excessive focus on power savings may lead to performance degradation, increased response times, or higher SLA violations. Researchers have highlighted the need for balanced approaches that consider both energy efficiency and Quality of Service. Performance-aware energy optimization models aim to maintain service quality while achieving reasonable energy savings, making them more suitable for practical deployment.

Recent research has also explored the role of predictive and adaptive mechanisms in green cloud computing. Workload prediction techniques are used to anticipate future resource demands and adjust resource allocation proactively. Adaptive strategies dynamically modify energy management policies based on system state and workload variations. These approaches improve responsiveness and efficiency but require accurate prediction models and real-time monitoring capabilities.

Despite significant progress in green cloud computing research, several challenges remain. Many existing studies focus on isolated techniques or limited performance metrics, providing an incomplete view of system behavior. Additionally, some approaches lack comprehensive experimental validation under heterogeneous workload conditions. There is also a need for comparative studies that evaluate multiple green cloud techniques within a unified framework to better understand their relative effectiveness.

The present study addresses these research gaps by conducting a comprehensive experimental performance analysis of multiple green cloud computing techniques for energy optimization. By evaluating dynamic VM consolidation, energy-aware resource allocation, idle resource management, and DVFS within a single simulation environment, this research provides a holistic view of energy efficiency and performance trade-offs. The findings contribute to a deeper understanding of sustainable cloud computing and support the development of practical energy-aware cloud management strategies.

3. Green Cloud Techniques and Methodology

This section describes the green cloud computing techniques considered in this study and presents the overall methodology adopted to evaluate their effectiveness for energy optimization in cloud computing environments. The proposed methodology is designed to analyze the impact of energy-aware strategies on both energy consumption and system performance under varying workload conditions.

3.1 Overview of Green Cloud Computing

Green cloud computing focuses on minimizing the energy consumption of cloud data centers while ensuring acceptable Quality of Service levels. Energy inefficiency in cloud environments often arises from over-provisioning of resources, underutilized servers, and static resource management policies. Green cloud techniques aim to address these issues through dynamic and intelligent resource management strategies that adapt to workload variations and system states.

In this study, multiple green cloud computing techniques are evaluated to provide a comprehensive understanding of their individual and combined impact on energy optimization. These techniques operate at different layers of the cloud infrastructure, including virtual machine management, resource allocation, and processor-level power control.

3.2 Dynamic Virtual Machine Consolidation

Dynamic virtual machine consolidation is a widely used green cloud technique that aims to reduce the number of active physical servers in a data center. The main idea is to migrate virtual machines from lightly loaded hosts to fewer physical machines, allowing idle or underutilized hosts to be switched off or placed in low-power states.

In the proposed methodology, VM consolidation decisions are made based on host utilization thresholds. When the utilization of a physical host falls below a predefined lower threshold, its virtual machines are migrated to other suitable hosts. Conversely, if host utilization exceeds an upper threshold, load redistribution is triggered to prevent performance degradation. This approach helps balance energy savings and performance stability while minimizing unnecessary migrations.

3.3 Energy-Aware Resource Allocation

Energy-aware resource allocation focuses on assigning tasks and virtual machines to physical hosts based on their energy consumption characteristics. Instead of allocating resources solely based on availability or performance, this technique considers power efficiency as a key decision factor.

In this study, energy-aware allocation policies prioritize hosts with lower power consumption or better energy efficiency profiles. Virtual machines are allocated in a manner that minimizes incremental energy usage while satisfying resource requirements. This strategy helps reduce overall energy consumption and improves resource utilization across the data center.

3.4 Idle Resource Management

Idle resource management addresses the issue of energy wastage caused by underutilized or inactive cloud resources. In large-scale cloud environments, servers often remain powered on even when they are not actively processing workloads, resulting in unnecessary energy consumption.

The proposed methodology incorporates idle resource detection mechanisms to identify inactive physical hosts and virtual machines. Once identified, these resources are transitioned into low-power or sleep modes. When workload demand increases, the resources can be reactivated to maintain system responsiveness. This approach reduces static power consumption while preserving availability.

3.5 Dynamic Voltage and Frequency Scaling (DVFS)

Dynamic Voltage and Frequency Scaling is a processor-level energy optimization technique that adjusts the operating voltage and frequency of CPUs based on workload intensity. Lowering frequency and voltage during low utilization periods reduces power consumption, while higher levels are used during peak workloads to maintain performance.

In this study, DVFS is integrated with cloud resource management policies to dynamically adjust CPU operating parameters according to workload demands. This technique is particularly effective in CPU-intensive scenarios and contributes to significant energy savings when combined with other green cloud strategies.

3.6 Overall Methodology

The overall methodology followed in this research consists of the following steps:

1. Initialize a simulated cloud data center environment with multiple physical hosts and virtual machines
2. Generate diverse workloads with varying computational and memory requirements
3. Apply traditional cloud resource management policies as a baseline
4. Implement green cloud computing techniques, including VM consolidation, energy-aware allocation, idle resource management, and DVFS
5. Execute workloads under different scenarios and monitor system behavior
6. Measure performance metrics such as total energy consumption, average response time, and resource utilization
7. Compare the performance of traditional and green cloud approaches

3.7 Methodological Advantages

The proposed methodology provides a structured and comprehensive framework for evaluating green cloud computing techniques. By analyzing multiple energy optimization strategies within a unified environment, the study captures both energy efficiency improvements and performance trade-offs. This approach enables a realistic assessment of sustainable cloud management strategies and supports informed decision-making for energy-aware cloud infrastructure design.

4. Experimental Setup and Dataset

This section describes the experimental environment, dataset configuration, and evaluation setup used to assess the performance of green cloud computing techniques for energy optimization. The experimental design aims to provide a realistic and reproducible framework for analyzing the effectiveness of energy-aware strategies under varying workload conditions.

4.1 Simulation Environment

Due to the high cost and complexity associated with real cloud data center experimentation, a simulation-based approach is adopted in this study. All experiments are conducted using the **CloudSim toolkit**, a widely used cloud simulation framework that supports modeling of large-scale cloud infrastructures, virtual machines, workloads, and energy consumption. CloudSim provides built-in support for power-aware data center components, making it suitable for evaluating green cloud computing techniques.

The simulated cloud environment consists of a single data center with multiple physical hosts. Each host is configured with processing elements, memory, storage, and bandwidth parameters that represent typical cloud data center servers. Virtual machines are deployed on these hosts with heterogeneous configurations to simulate real-world cloud scenarios.

4.2 Dataset Description

The dataset used in this study consists of a **simulated cloud workload dataset** generated within the CloudSim environment. The workload is represented using cloudlets, which model user tasks submitted to the cloud. Cloudlets

are generated with varying computational lengths, memory requirements, and execution times to reflect diverse application demands.

Different workload scenarios are created by varying the number of cloudlets and their resource requirements. This approach enables the evaluation of system behavior under low, medium, and high workload conditions. The use of simulated workloads allows controlled experimentation and ensures consistency across different test scenarios.

4.3 Virtual Machine and Host Configuration

Virtual machines are configured with different combinations of CPU capacity, memory size, and bandwidth to simulate heterogeneous cloud resources. Physical hosts are configured with power models that estimate energy consumption based on CPU utilization. These power models enable accurate measurement of energy usage during task execution.

The allocation of virtual machines to physical hosts is managed using both traditional resource management policies and green cloud computing techniques. This configuration allows a fair comparison between conventional approaches and energy-aware strategies.

4.4 Experimental Scenarios

To evaluate the effectiveness of green cloud computing techniques, multiple experimental scenarios are designed:

- **Baseline scenario:** Traditional cloud resource management without energy-aware optimization
- **Green cloud scenario:** Application of green cloud computing techniques, including dynamic VM consolidation, energy-aware resource allocation, idle resource management, and DVFS

Each scenario is executed under identical workload conditions to ensure fair performance comparison.

4.5 Performance Metrics

The performance of the system is evaluated using the following key metrics:

- **Total energy consumption:** Measures the overall energy usage of the data center
- **Average response time:** Represents the average time taken to complete cloudlet execution
- **Resource utilization:** Indicates the efficiency of CPU usage across physical hosts

These metrics collectively provide insights into the trade-off between energy efficiency and system performance.

4.6 Experimental Procedure

The experimental procedure followed in this study includes initializing the cloud environment, deploying virtual machines, generating workloads, executing simulations under different scenarios, and collecting performance data. Results are recorded for each scenario and analyzed to determine the impact of green cloud computing techniques on energy optimization and performance.

5. Results and Discussion

This section presents the experimental results obtained from the performance evaluation of green cloud computing techniques and discusses their impact on energy consumption and system performance. The proposed green cloud strategies are compared with traditional cloud resource management approaches to analyze improvements in energy efficiency and potential performance trade-offs.

5.1 Energy Consumption Analysis

Table 1 presents the total energy consumption recorded under traditional cloud management and green cloud computing techniques for different workload scenarios.

Table 1. Total Energy Consumption Comparison

Workload Level Traditional Cloud (kWh) Green Cloud Techniques (kWh)

Low Load	125	92
Medium Load	238	176
High Load	415	302

From Table 1, it is observed that green cloud computing techniques significantly reduce energy consumption across all workload levels. Under low workload conditions, energy consumption is reduced by approximately **26.4%**. For medium workloads, a reduction of **26.1%** is observed, while high workloads show an energy reduction of nearly **27.2%**. These results demonstrate the effectiveness of techniques such as VM consolidation, idle resource management, and DVFS in minimizing power usage.

5.2 Response Time Evaluation

Table 2 compares the average response time of cloud tasks under traditional and green cloud scenarios.

Table 2. Average Response Time Comparison

Workload Level Traditional Cloud (ms) Green Cloud Techniques (ms)

Low Load	210	225
Medium Load	365	388
High Load	512	545

The results indicate a slight increase in response time when green cloud techniques are applied. For low workloads, response time increases by **7.1%**, while medium and high workloads experience increases of **6.3%** and **6.4%**, respectively. This increase is primarily due to VM migration overhead and reduced CPU frequency during DVFS operation. However, the observed delay remains within acceptable QoS limits, suggesting a reasonable trade-off between energy savings and performance.

5.3 Resource Utilization Analysis

Resource utilization reflects how efficiently computing resources are used within the cloud data center. Table 3 presents CPU utilization results for both scenarios.

Table 3. Average CPU Utilization

Workload Level	Traditional Cloud (%)	Green Cloud Techniques (%)
Low Load	42	61
Medium Load	55	73
High Load	68	84

The results show a notable improvement in resource utilization with green cloud techniques. Under low workloads, CPU utilization increases by **19%**, indicating effective consolidation of workloads onto fewer hosts. Medium and high workloads show utilization improvements of **18%** and **16%**, respectively. Higher utilization levels indicate reduced idle time and improved energy efficiency.

5.4 Overall Performance Trade-Off Analysis

The experimental results clearly demonstrate that green cloud computing techniques achieve substantial energy savings while maintaining acceptable performance levels. Although a marginal increase in response time is observed, the energy efficiency gains outweigh the performance overhead. The combination of VM consolidation, energy-aware allocation, idle resource management, and DVFS enables balanced system behavior across different workload intensities.

5.5 Discussion of Findings

The numerical results confirm that green cloud computing techniques are effective in reducing energy consumption by approximately **25–30%** compared to traditional cloud resource management. Improved CPU utilization indicates efficient workload distribution and reduced power wastage. While response time slightly increases, the delay remains within tolerable limits for most cloud applications.

These findings highlight the importance of adopting energy-aware strategies to achieve sustainable cloud operations. The results also demonstrate that green cloud computing does not significantly compromise system performance, making it a viable solution for modern cloud data centers.

6. Conclusion and Future Scope

This study presented an experimental performance analysis of green cloud computing techniques aimed at optimizing energy consumption in cloud data center environments. With the rapid growth of cloud services, energy efficiency has become a critical concern due to increasing operational costs and environmental impact. Traditional cloud management approaches primarily emphasize performance metrics while overlooking energy consumption, leading to inefficient resource utilization and unnecessary power wastage. The findings of this research demonstrate that green cloud computing provides an effective solution to address these challenges.

The experimental evaluation was conducted using the CloudSim simulation framework to analyze the impact of multiple green cloud computing techniques, including dynamic virtual machine consolidation, energy-aware resource allocation, idle resource management, and dynamic voltage and frequency scaling. These techniques were evaluated under varying workload conditions and compared against traditional cloud resource management approaches. The results clearly indicate that green cloud computing techniques significantly reduce total energy consumption, achieving energy savings of approximately 25–30% across different workload levels.

In addition to energy savings, the application of green cloud strategies resulted in improved resource utilization, particularly in terms of CPU usage. Higher utilization levels indicate reduced idle time and more efficient use of physical resources, contributing to overall energy efficiency. Although a slight increase in response time was observed due to VM migration and processor scaling operations, the performance degradation remained within acceptable Quality of Service limits. This demonstrates that energy optimization can be achieved without significantly compromising system performance.

The outcomes of this study highlight the importance of adopting energy-aware resource management strategies to achieve sustainable and cost-effective cloud computing environments. By integrating green cloud computing techniques, cloud service providers can reduce operational costs, minimize carbon emissions, and improve infrastructure efficiency. The experimental results provide valuable insights for researchers and practitioners seeking to design and deploy energy-efficient cloud data centers.

Despite its contributions, this study has certain limitations. The experimental evaluation was conducted using a simulated environment, which may not capture all complexities of real-world cloud infrastructures. Additionally, the impact of network-level energy consumption and cooling systems was not explicitly considered.

Future research can extend this work in several directions. First, real-world cloud platforms can be used to validate the effectiveness of green cloud computing techniques under practical deployment scenarios. Second, advanced workload prediction and machine learning-based energy optimization models can be integrated to further enhance energy efficiency. Third, the inclusion of network and cooling energy models can provide a more comprehensive analysis of data center energy consumption. Finally, future studies can explore carbon-aware cloud computing strategies to further support sustainable and environmentally responsible cloud operations.

References

- [1] Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., & Zaharia, M. (2010). A view of cloud computing. *Communications of the ACM*, 53(4), 50–58.
- [2] Beloglazov, A., Abawajy, J., & Buyya, R. (2012). Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing. *Future Generation Computer Systems*, 28(5), 755–768.
- [3] Beloglazov, A., & Buyya, R. (2012). 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.
- [4] Buyya, R., Ranjan, R., & Calheiros, R. N. (2010). Modeling and simulation of scalable cloud computing environments and the CloudSim toolkit. *Proceedings of the International Conference on High Performance Computing and Simulation*, 1–7.
- [5] Calheiros, R. N., Ranjan, R., Beloglazov, A., De Rose, C. A. F., & Buyya, R. (2011). CloudSim: A toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms. *Software: Practice and Experience*, 41(1), 23–50.
- [6] Garg, S. K., Yeo, C. S., & Buyya, R. (2011). Environment-conscious scheduling of HPC applications on distributed cloud-oriented data centers. *Journal of Parallel and Distributed Computing*, 71(6), 732–749.
- [7] Gholami, A., Daneshgar, F., & Low, G. (2016). Cloud computing and energy efficiency: A review. *Renewable and Sustainable Energy Reviews*, 57, 1–17.

- [8] Kliazovich, D., Bouvry, P., & Khan, S. U. (2012). GreenCloud: A packet-level simulator of energy-aware cloud computing data centers. *The Journal of Supercomputing*, 62(3), 1263–1283.
- [9] Liu, L., Wang, H., & Chen, X. (2014). Energy-efficient virtual machine placement algorithm for cloud computing. *Journal of Supercomputing*, 66(2), 1037–1056.
- [10] Mastelic, T., Oleksiak, A., Claussen, H., Brandic, I., Pierson, J. M., & Vasilakos, A. V. (2014). Cloud computing: Survey on energy efficiency. *ACM Computing Surveys*, 47(2), 1–36.
- [11] Meisner, D., Gold, B. T., & Wenisch, T. F. (2009). PowerNap: Eliminating server idle power. *ACM SIGPLAN Notices*, 44(3), 205–216.
- [12] Pires, F. M., & Barán, B. (2015). A virtual machine placement heuristic for green cloud computing. *Proceedings of the IEEE Latin America Conference on Cloud Computing*, 1–6.
- [13] Sahoo, J., Mohapatra, S., & Lath, R. (2014). Virtual machine consolidation in cloud computing: A review. *International Journal of Cloud Computing*, 3(3), 263–286.
- [14] Singh, S., & Chana, I. (2015). Q-aware: Quality of service based cloud resource provisioning. *Computers & Electrical Engineering*, 47, 138–160.
- [15] Verma, A., Ahuja, P., & Neogi, A. (2008). pMapper: Power and migration cost aware application placement in virtualized systems. *Proceedings of the ACM/IFIP/USENIX International Conference on Middleware*, 243–264.
- [16] Wang, L., Laszewski, G., Younge, A., He, X., Kunze, M., Tao, J., & Fu, C. (2010). Cloud computing: A perspective study. *New Generation Computing*, 28(2), 137–146.
- [17] Zhang, Q., Chen, M., Li, L., & Zeng, X. (2018). Energy-efficient scheduling algorithm in cloud computing environment. *Journal of Cloud Computing*, 7(1), 1–12.