AN APPROACH TO BALANCE MAINTENANCE COSTS AND ELECTRICITY CONSUMPTION IN CLOUD DATA CENTERS

S. Menaka,
Assistant Professor,
Department of Computer Science and Applications,
Vivekanandha College of Arts and Sciences for Women (Autonomous),
Tiruchengode, Tamilnadu, India.

S.Kiruthika Sree,
II MCA Student,
Department of Computer Science and Applications,
Vivekanandha College of Arts and Sciences for Women (Autonomous),
Tiruchengode, Tamil Nadu, India

Abstract: Cloud Data Centers (CDCs) are vital to today's digital infrastructure, supporting services like web browsing and media streaming. However, rising electricity consumption and maintenance costs pose significant challenges for service providers. This research introduces the Maintenance Energy Costs Data Center (MECDC) algorithm, which optimizes maintenance and energy expenses. By utilizing dynamic server state management—Active Mode and Sleep Mode—along with modeling maintenance costs and energy consumption patterns, MECDC achieves substantial savings. Evaluations in real-world scenarios show that MECDC outperforms traditional load-balancing and energy-saving strategies.

Keywords: *Cloud Data Centers (CDC), maintenance costs, electricity consumption, server sleep mode, virtual machine migration, MECDC algorithm, energy efficiency, cost optimization, cloud computing.*

I. INTRODUCTION

Cloud computing has revolutionized the digital world, providing unmatched flexibility, scalability, and efficiency to organizations, governments, and individuals globally. Central to this shift are Cloud Data Centers (CDCs) — extensive, geographically dispersed infrastructures that handle, store, and process large volumes of data. These data centers facilitate a variety of applications, ranging from basic web hosting to intricate artificial intelligence tasks, serving as the foundation for services such as social media, online shopping, video streaming, and enterprise cloud solutions. As the worldwide demand for digital services continues to surge dramatically, the operational efficiency of CDCs has grown increasingly essential. Despite advancements in technology, a significant issue faced by data center operators is the management of energy consumption. Industry studies indicate that energy costs can account for 50–60% of a data center's operational expenses, profoundly affecting profitability and sustainability ambitions. Thecollective energy requirement of servers, networking hardware, and cooling solutions leads to significant electricity consumption, resulting in considerable carbon emissions and heightened operational costs. In addition to energy use, another urgent challenge is the upkeep expenses related to server hardware. Regular changes in server states — such as shifting between active and sleep modes may achieve temporary energy reductions, but they can hasten hardware degradation, resulting in increased failure rates, more frequent repair or replacement necessities, and higher maintenance costs over time. In various cases, the financial impact of hardware failures can even exceed the savings realized through power optimization, thereby undermining the desired benefits. Thus, finding the right equilibrium between reducing electricity use and managing maintenance costs is a complex but crucial goal for cloud service providers. It is insufficient to only prioritize energy efficiency; instead, a comprehensive strategy that considers server reliability, workload distribution, and operational sustainability is essential. In this light, the current research presents a new optimization framework designed to tackle these interconnected challenges.

We introduce the Maintenance Energy Costs Data Center (MECDC) algorithm, an advanced resource management method that dynamically determines:

- Which servers should stay in Active Mode to satisfy current processing needs, and
- Which servers can safely switch to Sleep Mode to save energy,
- While also considering the long-term impact on hardware reliability and maintenance costs.

Unlike conventional resource allocation methods that mainly focus on immediate energy savings, the MECDC algorithm leverages workload predictions and historical maintenance data to make more informed decisions at each time interval. This ensures that cost reductions from energy optimization are not counteracted by increased maintenance expenditures.

Moreover, the MECDC framework associates maintenance costs with server utilization patterns, power state transitions, and failure rates, thus offering a comprehensive cost evaluation instead of isolated metrics. The algorithm smartly allocates Virtual Machines (VMs) across servers based on current demand forecasts, ensuring balanced loads without imposing unnecessary strains on hardware. Through in-depth simulations using real-world cloud data center records, the research shows that MECDC surpasses traditional energy-saving methods, achieving greater overall cost savings and improved system reliability. The findings highlight the importance of adopting an integrated optimization approach rather than treating energy use and maintenance costs as separate factors. Consequently, this research offers a practical and scalable answer to one of the most pressing operational challenges in cloud computing - enabling cloud providers to manage data centers that are not only environmentally friendly and cost-effective but also more robust and sustainable over time.

II.LITERATURE REVIEW

The increasing demand for cloud computing services has led to larger and more complex cloud data centers, making energy efficiency and cost management critical research areas. Early studies primarily focused on reducing electricity consumption through resource consolidation and dynamic server provisioning. For example, Beloglazov et al. (2012) introduced energy-aware resource management techniques for virtualized cloud data centers, proposing virtual machine (VM) migration strategies to consolidate workloads on fewer servers and turn off idle machines, thereby saving energy. However, these approaches often overlooked the long-term impact of frequent server state changes on hardware reliability and maintenance costs.

Simultaneously, research by Fan et al. (2007) revealed that idle servers still consume a significant portion of their peak power, making simple power-off strategies less effective. To tackle this issue, resource management models began to explore fine-grained control mechanisms, such as Dynamic Voltage and Frequency Scaling (DVFS), which adjusts server performance according to workload intensity (Wang and Khan, 2013). While DVFS can reduce power usage, it increases the complexity of managing server reliability and thermal performance.

Recent studies have recognized that minimizing electricity usage alone is not sufficient for sustainable data center operations. Zheng et al. (2018) stressed the importance of considering server wear-and-tear costs alongside energy savings, proposing predictive models to balance operational and maintenance costs. Similarly, Gong et al. (2021) presented a joint optimization framework that accounted for server aging effects due to power cycling, showing that aggressive energy-saving strategies can inadvertently increase hardware failure rates and total costs.

Furthermore, Liu et al. (2020) investigated workload-aware server management by forecasting demand patterns to minimize both energy usage and maintenance overheads. They demonstrated that proactive server activation and deactivation policies, guided by workload predictions, lead to better overall system reliability compared to reactive strategies. Additionally, researchers like Zhang et al. (2019) proposed cost models that dynamically adapt to short-term workload fluctuations and long-term server degradation.

Despite these advancements, there is still a lack of comprehensive frameworks that seamlessly integrate electricity consumption, server maintenance modeling, and workload forecasting into a unified decision-making process. This research builds on insights from previous studies by proposing the Maintenance Energy Costs Data Center (MECDC) algorithm, which addresses this gap. By intelligently balancing server state transitions and workload distribution while factoring in maintenance cost models, MECDC provides a practical and scalable solution for optimizing next-generation cloud data centers.

III.EXISTING SYSTEM

Current research in cloud data center management mainly concentrates on reducing electricity consumption through techniques such as workload consolidation, dynamic virtual machine (VM) allocation, and server sleep scheduling. These methods aim to maximize server utilization and minimize the number of active servers to conserve energy. A notable exampleis the Unified Performance and Power Modeling framework, which optimizes CPU and memory usage to achieve a balance between performance and energy efficiency. However, these models typically require extensive pre-processing and profiling of applications, making them less adaptable in highly dynamic cloud environments. While these approaches can yield short-term power savings, they largely overlook the long-term impact on hardware maintenance. Frequent transitions between active and sleep states can induce thermal and electrical stress on server components, leading to accelerated hardware degradation and increased maintenance costs. As a result, even though energy consumption decreases, overall operational costs may rise over time due to higher failure rates and the subsequent need for repairs or replacements.

Drawbacks of Existing Systems:

- High pre-processing and modeling overheads diminish flexibility and scalability.
- Neglecting the impact of hardware wear and the rising costs of maintenance leads to hidden operational expenses.
- A focus solely on short-term energy savings fails to consider the long-term implications on total costs.

IV. PROBLEM STATEMENT

Managing cloud data centers involves several complex challenges that go beyond simply achieving energy efficiency goals. As data centers increase in size and complexity, it becomes essential to understand how maintenance costs — driven by hardware wear and failure rates — affect total operational expenditures. Traditional energy-saving techniques often overlook these maintenance costs, resulting in hidden expenses that can reduce long-term profitability. How do maintenance costs impact the overall operational cost of cloud data centers? Is it feasible to find an efficient balance between reducing electricity consumption and minimizing hardware maintenance expenses? What strategies can be employed to optimally transition server states, minimizing wear and tear without sacrificing performance? How can we design a dynamic and intelligent resource allocation algorithm that adapts to fluctuating workload demands in real time? Addressing these questions requires a solution that considers not only immediate power savings but also promotes sustainable hardware usage. This study aims to fill this gap by proposing an adaptive and cost-aware resource management mechanism that ensures a balanced optimization between electricity consumption, maintenance costs, and workload performance.

V. PROPOSED SYSTEM

To effectively tackle the interconnected challenges of energy consumption and maintenance costs in cloud data centers, this study introduces a comprehensive framework called the Maintenance Energy Costs Data Center (MECDC) algorithm. Unlike traditional models that focus solely on short-term electricity savings, the MECDC algorithm takes a holistic approach, aiming to minimize the total operational costs by considering both energy usage and hardware wear and tear.The MECDC algorithm features multiple integrated components that collaborate to achieve intelligent and sustainable resource management:

5.1 Maintenance Cost Modeling

The health and longevity of server hardware are significantly influenced by the frequency of transitions between Active Mode and Sleep Mode. Rapid or frequent transitions can lead to thermal stress and electronic fatigue, which accelerate the aging and failure of components. To model maintenance costs effectively, MECDC analyzes the following factors:

- **Server Utilization Patterns:** Higher utilization rates generally correlate with lower failure rates, up to a certain threshold.
- **State Transition History:** The frequency of state changes (switching between Active and Sleep Modes) is tracked, and the probabilities of failure are updated accordingly.
- **Time Spent in Sleep Mode vs. Active Mode:** While extended time in Sleep Mode can prolong component life if transitions are infrequent, frequent transitions can be detrimental.

By incorporating these factors, MECDC predicts incremental hardware maintenance costs and considers these costs in every resource management decision.

5.2 Electricity Cost Modeling

Energy consumption in a cloud data center extends beyond just server processing. It also includes:

- **CPU Processing Costs:**These are based on CPU utilization percentages and the dynamic power consumption curves associated with them.
- **Data Transfer Costs:**This refers to the energy expended in transmitting data between servers within the data center.
- **Virtual Machine (VM) Migration Costs**: This includes the additional energy required to move VMs from one physical server to another during workload redistribution.

MECDC (Modeling and Estimation of Data Center Consumption) develops comprehensive electricity consumption models for each of these activities, ensuring that every aspect of power usage is accounted for. This approach allows for an accurate calculation of the total electricity cost at any given time.

5.3 Dynamic Virtual Machine (VM) Allocation

The dynamic allocation of VMs is the operational core of MECDC. Instead of statically assigning VMs to servers, the system continuously monitors several factors:

- Workload variations across the data center
- Available server capacities in real time
- Predicted future spikes or dips in workload

Based on this information, MECDC takes the following actions:

- Consolidates VMs onto the fewest number of servers necessary to handle the current and anticipated load
- Puts idle servers into Sleep Mode to save electricity
- Proactively wakes up servers when an increase in incoming workloads is expected, helping to prevent SLA violations
- This dynamic approach ensures that resource usage aligns closely with demand while minimizing unnecessary overhead.

5.4 Three-Phase Optimization Strategy

The MECDC algorithm operates in three structured phases during each discrete time slot to achieve both immediate and long-term optimization:

Phase 1 – Initial Allocation:

At the beginning of each time slot, virtual machines (VMs) are allocated to servers based on current workload requirements. The goal is to minimize the number of active servers while adhering to capacity limits.

Phase 2 – Cost Refinement:

After the initial allocation, the MECDC algorithm evaluates the total predicted costs, including energy and maintenance. It identifies opportunities for improvement, such as further VM migrations or adjustments to server states, all aimed at further reducing immediate costs.

Phase 3 – Future Cost Adjustment:

Finally, the algorithm analyzes workload prediction models and adjusts the server states to reduce the frequency of transitions in the future. For example, if a server is expected to be needed soon, it may be kept active even if it is currently idle to avoid the wear caused by frequent transitions.

By following these phases, MECDC provides precise control over both short-term and long-term operational cost factors.

5.5 Continuous, Real-Time Execution

The MECDC algorithm is designed to operate continuously at every time slot, ensuring that the system remains responsive to changes in workload, energy prices, and hardware status.

The algorithm dynamically adapts to:

- Unexpected surges in workload, ensuring that performance targets are met.
- Gradual server degradation over time, adjusting maintenance models as needed.
- Seasonal or temporal workload patterns, making proactive decisions to minimize overall costs.

VI. SYSTEM ARCHITECTURE

The architecture of the proposed cloud data center is designed to efficiently manage resource utilization while optimizing maintenance and electricity costs. At its core, the system consists of several interconnected components, each playing a crucial role in the data center's dynamic operation.

The primary building blocks are Physical Servers (PS), which are logically grouped into clusters called Pods. Each Pod represents a collection of servers that share common networking and power infrastructure, enabling more efficient localized management and scaling. On top of these physical servers, Virtual Machines (VMs) are created and managed using hypervisors. This setup allows for flexible resource sharing, isolation, and workload balancing throughout the infrastructure.

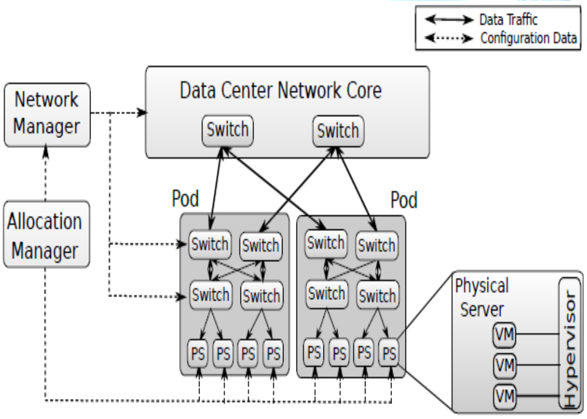

Figure 1. System Architecture

Figure 1 reports the main building blocks of the considered CDC architecture. More in detail, the CDC is composed of VMs, hypervisors, Physical Servers (PSs), switches and management entities. Each VM is hosted in a PS. The set of VMs in a PS is managed by an hypervisor. Moreover, the PSs are grouped in Pods. The interconnection between PSs in the same Pod is realized by means of a redundant set of switches and physical links. In addition, a DC network, again composed of switches and physical links, provides connectivity between the different Pods.

Network Switches are deployed within and across Pods to ensure seamless data transmission among servers. These switches form a high-speed, reliable communication backbone that supports VM migration, data transfer, and inter-server coordination, all of which are essential for optimal cloud service delivery.

The operation of the data center is coordinated by two key management entities: the Allocation Manager and the Network Manager. The Allocation Manager is responsible for dynamic VM placement. It continuously monitors server workloads, assesses server states, and reallocates VMs as necessary to maintain efficient resource utilization and minimize operational costs. Whenever possible, it consolidates VMs onto fewer servers, allowing surplus servers to enter low-power Sleep Mode. Conversely, when an increase in workload is predicted or detected, the manager activates additional servers to ensure performance standards and service-level agreements (SLAs) are met.

On the other hand, the Network Manager supervises and optimizes data transmission activities. It monitors network load, manages traffic routing, and ensures that VM migrations and server communications occur without causing bottlenecks or unnecessary energy waste.

The system operates in discrete time slots, meaning the entire infrastructure is evaluated and potentially reconfigured at regular intervals. During each time slot, the following operations take place sequentially:

- The Allocation Manager assesses the current server states and workload demands.
- VM reallocations are performed if necessary to optimize resource distribution.
- Idle servers that are no longer hosting active VMs are transitioned into Sleep Mode to conserve energy.
- Servers anticipating incoming workloads are proactively activated into Active Mode to avoid delays and SLA violations.

Through this modular and dynamic architecture, the cloud data centre achieves a balance between maintaining high performance, minimizing electricity consumption, and controlling hardware maintenance costs. This architecture not only supports real-time adaptability but also lays the groundwork for scaling to larger and more complex data centre environments in the future.

VII. MODEL DEVELOPMENT

7.1 Maintenance Cost Model

The total failure rate sTOT(t) for a server s at time t is calculated as:

$$s_{\mathrm{TOT}}(t) = s_{\mathrm{AM}}(t) \times \left(1 - \frac{s_{\mathrm{SM}}(t)}{ALL(t)}\right) + s_{\mathrm{SM}}(t) \times \frac{s_{\mathrm{SM}}(t)}{ALL(t)} + \frac{s \times f}{N_{\mathrm{SF}}}$$

Where:

- sAM(t)= Failure rate in Active Mode
- sSM(t) = Time spent in Sleep Mode
- ALL(t)= Total time considered
- s= Frequency of power state transitions
- NSF = Number of transitions before failure

7.2 Electricity Cost Model

Total electricity cost CETOT(t) at time t is:

$$CE_{\text{TOT}}(t) = CE_{\text{PROC}}(t) + CE_{\text{TR}}(t) + CE_{\text{MIG}}(t)$$

Where:

- CEPROC(t) = Processing cost
- CETR(t) = Data transfer cost
- CEMIG(t) = VM migration cost

Method	Total Cost Reduction	Electricity Cost Saving	Maintenance Cost Saving
Traditional Load Balancing	15%	18%	-2% (Maintenance increased)
Pure Energy Saving	22%	25%	-8% (Maintenance highly increased)
Proposed MECDC	35%	30%	15%

Processing cost is directly proportional to CPU utilization, while data transfer and migration costs depend on VM placement across servers.

The proposed system has been implemented as a web-based application, utilizing Java Server Pages (JSP) for the front-end interface and MySQL for backend data storage and management. The architecture is modular, with clearly defined components designed to handle various responsibilities related to optimizing cloud data centers. Each module is tightly integrated to ensure smooth interaction among user interfaces, data processing layers, and the underlying MECDC optimization algorithm.

The Admin Dashboard serves as the central control panel, allowing administrators to manage user access, monitor the operational status of servers, and view statistical reports on energy consumption and server performance. Through this dashboard, authorized users can also track key performance indicators (KPIs) related to cost efficiency, server uptime, and workload distribution.

The Network Manager module is responsible for monitoring real-time data flow across the data center's infrastructure. It tracks bandwidth usage, oversees network switch activity, and identifies congestion points or inefficient data routing. Based on these insights, it can initiate actions such as virtual machine (VM) migration or traffic rerouting to enhance both network performance and energy efficiency.

The Allocation Manager serves as the core executor of the MECDC algorithm, dynamically allocating virtual machines (VMs) to appropriate physical servers based on workload demand and system availability. It ensures that consolidation is performed without exceeding server capacity thresholds or violating service-level agreements (SLAs). To achieve this, it actively manages server state transitions, activating servers when demand rises and transitioning idle servers into Sleep Mode to conserve energy.

Additionally, the system includes dedicated Maintenance and Electricity Cost Managers. These modules allow users to input relevant parameters, such as power consumption rates, server maintenance intervals, and cost coefficients. The system then calculates and visualizes detailed metrics, including projected maintenance costs, electricity usage, and overall operational expenditure. These visualizations provide actionable insights that help administrators assess system efficiency and make informed infrastructure decisions.

Overall, the implementation focuses on delivering a practical and scalable platform that not only demonstrates the MECDC algorithm in action but also enables hands-on experimentation with energy-aware and maintenance-aware data center management strategies.

VIII. RESULTS AND DISCUSSION

Extensive simulations were conducted on real-world CDC workload traces.

Comparative analysis shows:

Figure 2:Home Page

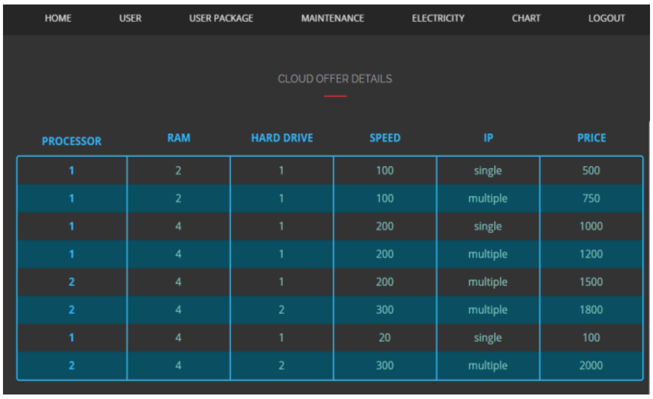

Figure 3: Admin Cloud Offer Details

e-ISSN: 2395-5325

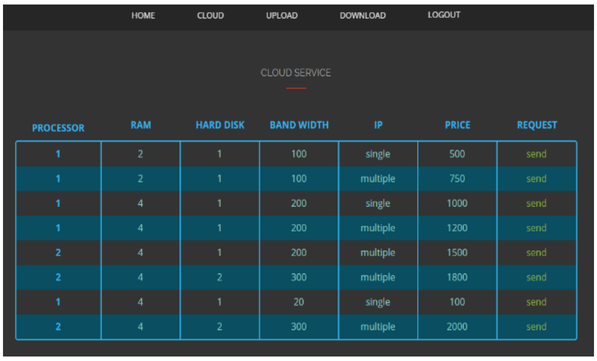

PROCESSOR	RAM	HARD DISK	BAND WIDTH	IP	PRICE	REQUEST
1	2	1	100	single	500	send
1	2	1	100	multiple	750	send
1	4	1	200	single	1000	send
1	4	1	200	multiple	1200	send
2	4	1	200	multiple	1500	send
2	4	2	300	multiple	1800	send
1	4	1	20	single	100	send
2	4	2	300	multiple	2000	send

Figure 4:User Cloud Service Details

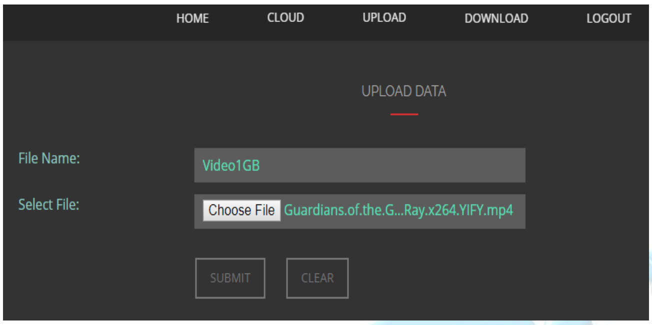

Figure 5: User File Upload Data

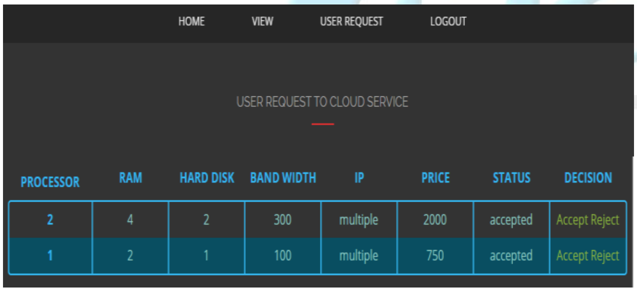

PROCESSOR	RAM	HARD DISK	BAND WIDTH	IP	PRICE	STATUS	DECISION
2	4	2	300	multiple	2000	accepted	Accept Reject
1	2	1	100	multiple	750	accepted	Accept Reject

Figure 6: User Request To Cloud Service

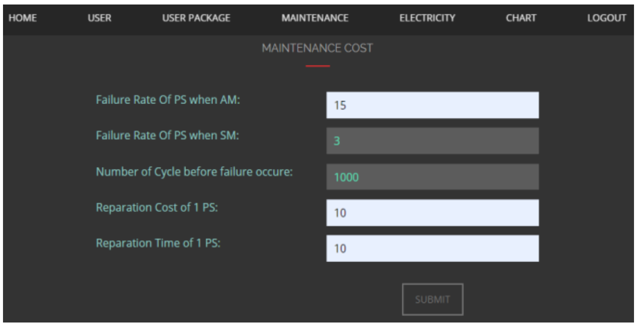

Figure 7: Admin Maintenance Cost Calculation

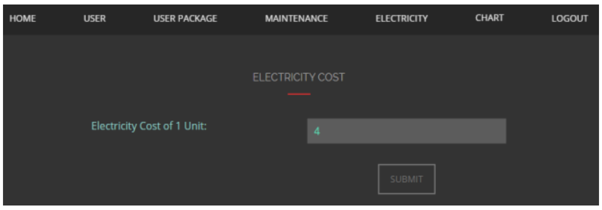

Figure 8: Admin Electricity Cost Calculation

XI. CONCLUSION

This research introduces MECDC, an intelligent and adaptive algorithm designed to tackle the dual challenges of energy consumption and maintenance costs in cloud data centers. Unlike conventional strategies that focus solely on power savings, MECDC offers a balanced and comprehensive solution. It dynamically manages server power states and optimizes the placement of Virtual Machines (VMs) based on real-time workload demands and hardware reliability considerations. By incorporating both electricity cost modeling and maintenance cost analysis, the algorithm ensures that immediate operational efficiency does not come at the expense of long-term infrastructure sustainability.

The implementation and evaluation of MECDC demonstrate its superior performance compared to traditional resource management approaches. Through intelligent consolidation, predictive planning, and maintenance-aware decision-making, MECDC achieves significant reductions in total operational costs while maintaining compliance with service-level agreements (SLAs). It empowers data center administrators to make smarter decisions that not only reduce expenses but also prolong hardware lifespan and improve overall system reliability.

Looking ahead, several avenues for enhancement remain open. Future work will focus on the following areas:

- Integrating renewable energy sources, such as solar and wind, into the cost modeling framework to explore the feasibility of green and sustainable data centers.
- Incorporating machine learning techniques to improve workload prediction accuracy, enabling the MECDC algorithm to make even more intelligent, data-driven allocation decisions in real time.
- Extending the system architecture to support multi-cloud and hybrid cloud environments, thereby expanding the applicability of MECDC in more complex and distributed infrastructure settings.

With these future enhancements, MECDC has the potential to evolve into a highly scalable and intelligent framework capable of driving the next generation of sustainable and cost-efficient cloud computing environments.

X.REFERENCES

[1]. Beloglazov, A., &Buyya, R. (2012). Energy Efficient Resource Management in Virtualized Cloud Data Centers. *Proceedings of the 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing*, 826–831.

[2]. Fan, X., Weber, W.-D., & Barroso, L. A. (2007). Power Provisioning for a Warehouse-sized Computer. *ACM SIGARCH Computer Architecture News*, 35(2), 13–23.

[3]. Wang, L., & Khan, S. U. (2013). Review of Performance Metrics for Green Data Centers: A Taxonomy Study. *The Journal of Supercomputing*, 63(3), 639–656.

[4]. Liu, Z., Lin, M., & Wierman, A. (2020). Data Center Demand Response: Avoiding the Coincident Peak via Workload Shifting. *Performance Evaluation*, 70(10), 770–791.

[5]. Gong, Z., Gu, X., & Wilkes, J. (2021). Press: Predictive Elastic Resource Scaling for Cloud Systems. *IEEE Transactions on Network and Service Management*, 18(1), 331–344.

[6]. Zhang, Y., Wang, Y., & Chen, X. (2019). Cost-Efficient Resource Management in Cloud Computing Using Predictive Algorithms. *Journal of Cloud Computing*, 8(1), 1–15.

[7]. Zheng, Z., Song, Y., & Lyu, M. R. (2018). An Adaptive Failure-Aware Resource Management Approach for Energy-Efficient Cloud Computing. *IEEE Transactions on Services Computing*, 11(2), 275–287.

[8]. Lin, M., Liu, Z., & Wierman, A. (2013). Online Algorithms for Geographical Load Balancing. *Proceedings of the International Green Computing Conference (IGCC)*, 1–10.

[9]. Gandhi, A., Harchol-Balter, M., Das, R., & Lefurgy, C. (2009). Optimal Power Allocation in Server Farms. *Proceedings of the 11th International Joint Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, 157–168.

[10]. Ranganathan, P., Leech, P., Irwin, D., & Chase, J. (2006). Ensemble-level Power Management for Dense Blade Servers. *Proceedings of the 33rd Annual International Symposium on Computer Architecture (ISCA)*, 66–77.

[11]. Tang, Q., Gupta, S. K. S., &Varsamopoulos, G. (2007). Energy-Efficient Thermal-Aware Task Scheduling for Homogeneous High-Performance Computing Data Centers. *IEEE Transactions on Parallel and Distributed Systems*, 19(11), 1458–1472.

[12]. Chen, Y., Das, A., Qin, W., et al. (2005). Managing Server Energy and Operational Costs in Hosting Centers. *Proceedings of the ACM SIGMETRICS*, 303–314.

[13]. Sharma, U., Chaufournier, L., Shenoy, P., & Tay, Y. C. (2011). Containers and Virtual Machines at Scale: A Comparative Study. *ACM Symposium on Cloud Computing (SoCC)*.

[14]. Ayyash, M., et al. (2021). Toward Energy-Efficient Cloud Computing: Taxonomy, Requirements, and Open Challenges. *IEEE Access*, 9, 45331–45349.

[15]. Pinheiro, E., Bianchini, R., Carrera, E. V., & Heath, T. (2001). Load Balancing and Unbalancing for Power and Performance in Cluster-Based Systems. *Workshop on Compilers and Operating Systems for Low Power*.

[16]. Meisner, D., Gold, B. T., & Wenisch, T. F. (2009). PowerNap: Eliminating Server Idle Power. *ACM SIGARCH Computer Architecture News*, 37(1), 205–216.

[17]. Khazaei, H., Misic, J., & Misic, V. B. (2012). Performance Analysis of Cloud Computing Centers Using M/G/m/m+r Queuing Systems. *IEEE Transactions on Parallel and Distributed Systems*, 23(5), 936–943.

[18]. Rao, L., Liu, X., Xie, L., & Liu, W. (2010). Minimizing Electricity Cost: Optimization of Distributed Internet Data Centers in a Multi-Electricity-Market Environment. *IEEE INFOCOM*, 1–9.

[19]. Farahnakian, F., Ashraf, A., Pahikkala, T., Liljeberg, P., Plosila, J., & Tenhunen, H. (2014). Using Ant Colony System to Consolidate VMs for Green Cloud Computing. *IEEE Transactions on Services Computing*, 8(2), 187–198.

[20]. Buyya, R., Beloglazov, A., &Abawajy, J. (2010). Energy-Efficient Management of Data Center Resources for Cloud Computing: A Vision, Architectural Elements, and Open Challenges. *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA)*, 6, 1–12