Multi-Objective Fuzzy Optimization for Energy, Carbon, and SLA Efficiency in Cloud Computing

Anamika Verma[1*], Dr. Rajesh Dangwal[2]

[1]Research Scholar, Department of Mathematics, BGR Campus, Hemvati Nandan Bahuguna Garhwal University, Pauri Garhwal (Uttarakhand)-246001 Email: vanamika1996@gmail.com; Orcid id: 0009-0001-5929-8601
[2]Professor and HOD, Department of Mathematics, BGR Campus, Hemvati Nandan Bahuguna Garhwal University, Pauri Garhwal (Uttarakhand)-246001 Email: dangwalrajeshhnbgu@gmail.com; Orcid id: 0009-0009-3516-317X

Abstract
The fast development of cloud computing has contributed to the great proliferation of datacenters across the planet, which means the gluttonous use of electricity and releasing more carbon dioxide. These do not only increase the cost of operations but also promote degradation of the environment in the global environment. Ensuring Quality of Service (QoS) and ensuring energy efficiency in a cloud service environment is a significant challenge to cloud service providers. This is a multi-objective optimization problem whose objectives are conflicting in that minimization of energy consumption, minimization of carbon footprint, and performance of the SLA will not be achieved simultaneously. This study suggests a multi-objective fuzzy optimization model, which can be used to balance the energy use, carbon emission, and SLA efficiency in dynamic cloud system environments in an intelligent manner. The framework incorporates a fuzzy inference system (FIS) to control the uncertainty and vagueness of workload changes, server usage, and the environment conditions. The system calculates the trade-offs of the energy, carbon and the SLA parameters in real time based on the fuzzy rules and membership functions. A hybrid metaheuristic optimizer with the ability to execute a global search and a local search is then applied on the fuzzy result to determine the best virtual machine placement and task scheduling techniques. The proposed model is rather dynamic in regards to the different workloads and thermal conditions, unlike the more traditional methods used, which are based on a predetermined set of deterministic thresholds, which guarantee sustainable resource distribution. The framework will help to reduce energy consumption and carbon emissions and decrease the violations of SLA considerably by integrating smart decision-making and multi-objective balancing. The study helps in the development of green and sustainable cloud computing, as it provides an adaptable, versatile, and green optimization model that can be used in present-day datacenter operations.

Keywords: Green Cloud Computing, Multi-Objective Optimization, Fuzzy Logic, Energy Efficiency, Carbon Emission Reduction, SLA Management, Virtual Machine Placement, Sustainable Datacenters.

1. Introduction

Large-scale datacenters have been propagated due to the exponential increase in the use of cloud computing and thus contribute significantly to carbon emissions in the world through electrical power consumption. Recent estimates suggest that datacenters consume close to 1-1.5% of the global electricity usage, and the estimates are anticipated to rise further (Jones, 2018). This growing energy requirement not only makes the operation of cloud service providers more expensive but also makes it more challenging to the environment, which is also contrary to the global sustainability objectives (Gill & Buyya, 2019).

At the same time, cloud providers have to comply with the high quality of Service Level Agreements (SLAs) that assure them of performance, availability, and reliability to the users. Breach of these contracts attracts monetary fines and garnishment of provider image. The main issue is, however, to balance three opposing requirements: the reduction of energy usage, the decrease of carbon emission, and the adherence to SLA (Beloglazov et al., 2012).

The conventional methods of resource management tend to be based on fixed thresholds and deterministic models and are not dynamic and uncertain in the cloud setting with fluctuating workloads and the use of servers as well as varying environmental parameters like cooling efficiency and carbon intensity of the grid (Kaur & Chana, 2015). Such failure will lead to sub-optimal performance which will be in the form of either, overuse of energy, overproduction of carbon, or overuse of SLA. In order to solve this multi-objective optimization problem under uncertainty, this paper provides a new Multi-Objective Fuzzy Optimization Framework. It uses the Fuzzy Inference System (FIS) to reflect imprecision and Hybrid Metaheuristic Optimizer to decide near-optimal strategies to put and schedule the virtual machines (VMs). The major contributions are:

1.An artificial fuzzy logic model that measures and weighs off trade-offs between energy, carbon and SLA goals in real-time.

2.Hybrid optimization: A Genetic Algorithm (GA) global exploration and Simulated Annealing (SA) local exploitation strategy.

3.Extensive testing on actual data traces, and showing considerable improvements over state of the art methods.

2. Literature Review

Cloud computing has become a key paradigm in providing scalable computational resources and services, however, the growing scale of cloud data centers has also led to significant energy consumption and the resulting environmental impacts. Therefore, numerous studies have been published on energy-efficient resource management and virtual machine placement have been the focus of research. VM consolidation and efficient use of resources to save energy were found important by Beloglazov et al. [1] for efficient data-center management.

Farahnakian et al. [3] used reinforcement learning in energy-efficient VM consolidation and Kaur and Chana [7] and Mastelic et al. [10] gave an extensive survey of energy-efficient techniques for cloud computing. Dayarathna et al. [16] also explored energy-consumption models for data centers, and Hsu et al. [18] explored task consolidation as a way to lower the energy use of a cloud environment. Although these studies show that the energy efficiency of cloud data centers can be significantly enhanced by effective VM placement and resource consolidation, they also do not fully consider the environmental and service quality considerations that are key to the operation of modern cloud systems.

Apart from the energy consumption, the environmental sustainability of cloud computing has spurred research in the areas of carbon-aware scheduling and resource management. Garg et al. [4] studied the scheduling of HPC applications in distributed data centers for greener computing, and Zhou et al. [14] studied carbon-aware load balancing of geographically distributed cloud services. Li et al. [19] also studied the problem of scheduling resources for heterogeneous cloud data centres to be carbon-aware. On a more general scale, Gill and Buyya [5] explored sustainable cloud computing at a number of different levels and Jones [6] asserted that data centres are becoming electricity hungry and the environmental impact is growing. Goiri et al. [20] explored the concept of renewable energy in cloud data-processing systems. All of these studies show that it is important to factor in energy use and associated environmental impact, rather than focusing solely on minimizing power consumption, to address the challenge of sustainable cloud management.

To tackle into the conflicting objectives in cloud resource management, several of the researcher have used multi-objective and metaheuristic optimization methods. Preemptable task scheduling in IaaS cloud systems was explored by Li et al. [8] and dynamic resource allocation was examined in cloud computing systems by Zhang et al. [13]. Srichandan et al.

[11] used a multi-objective hybrid optimization approach for cloud task scheduling, which showed the applicability of metaheuristic techniques to complex cloud-resource allocation problems. Cao and Lin [2] investigated a genetic-fuzzy approach for multi-objective VM placement considering energy and performance. These methods are appealing because VM placement is a non-linear and combinatorial optimization problem in which various system components such as energy consumption, resource usage, performance and service requirements may compete with each other. But the ability to perform well is highly dependent on the objective function and the values of the parameters of the metaheuristic methods, and traditional optimization models might not be suitable to capture the uncertainty and imprecision of the dynamic cloud environments.

Fuzzy logic has therefore been increasingly studied for use as a decision making process in cloud resource management as it can represent uncertain and imprecise information by using linguistic variables and reasoning with rules. Tang et al. [12] explored fuzzy-based decision mechanisms for thermal-aware data-center scheduling, and Cao and Lin [2] introduced fuzzy reasoning in conjunction with genetic optimization for energy- and performance-aware VM placement. The studies show that fuzzy system can be used to support flexible decision making when the cloud-resource condition cannot be precisely defined by crisp thresholds. Other more recent studies also study intelligent VM placement and fuzzy assisted VM placement. Ghasemi et al. [21] introduced a hybrid method that uses multi-objective reinforcement learning and clustering methods for VM placement that evaluates based on fuzzy systems and takes into account energy efficiency, resource utilization, load balancing, SLA requirements and execution time. Although they show the feasibility of using fuzzy decision mechanisms for multi-objective VM-placement problems, their optimization mechanism is different from the hybrid GA–SA method that is utilized in the present study.

Advanced and hybrid metaheuristic methods for VM placement with multiple objectives have also been studied recently. Alourani et al. [22] introduced an adaptive thresholds-based genetic algorithm (GA) energy-efficient VM-placement method that aimed to maximize resource utilization with the goal of minimizing the energy consumption and SLA violations. In [23] Liu et al. proposed a multi-objective Cauchy particle swarm optimization (PSO) algorithm for the energy-aware placement of virtual machines (VMs) that considers energy consumption, resource utilization, load balancing, and the robustness of the system. Baydoun and Zekri [24] introduced a network-, cost-, renewable-aware ant colony optimization framework (ACO) for the energy-efficient placement of VMs that also included energy efficiency, along with carbon emissions, network-related factors, cost, and renewable-energy considerations. The recent studies clearly show the trend of moving from single objective energy minimization towards multi-objective and environment friendly VM-placement strategies. However, these are different optimization strategies and formulations of the objective functions, and there is comparatively less research on the combination of fuzzy multi-objective decision-making and complementary global and local search.

Research Gap and Motivation

While there has been substantial advancement in energy-efficient virtual machine placement, carbon-aware cloud scheduling, fuzzy decision-making, and metaheuristic optimization, there are a number of challenges to the existing literature. The main focus of existing energy-aware approaches is to save energy, the focus of carbon-aware approaches is on environmental impact, while Service Level Agreement (SLA) requirements are either secondary considerations or constraints. There has been some recent works showing multi-objective approaches based on reinforcement learning, genetic algorithms, particle swarm optimization, ant colony optimization, etc. [21–24] have led to better performance in VM placement, but a unified decision-making framework which takes into account energy consumption, CO2 emissions, and SLA efficiency is relatively less studied.

In addition, fuzzy based methods [2,21] indicate the usefulness of intelligent decision-making under uncertain conditions in cloud situations, although the combination of a fuzzy multi-objective decision layer with a hybrid Genetic Algorithm–Simulated Annealing (GA–SA) optimization strategy has been little explored. Especially, the need is for an approach that can integrate the global exploration ability of GA and the local refinement ability of SA, and continuously incorporate the fuzzy reasoning to assess the trade-off between the energy objective, carbon objective and SLA objective.

Therefore, the major research gap that is being addressed by this study is the missing link of an integrated and adaptive optimization framework that simultaneously optimizes energy consumption, carbon emissions, and SLA efficiency for VM placement in cloud computing environments, with the combination of fuzzy multi-objective decision making and hybrid GA–SA optimization.

The main motivation of this study is the conflicting demands in the modern cloud data centers. Both energy saving and carbon emission reduction play crucial roles in maintaining operational efficiency as well as environmental sustainability. Meanwhile, there could be some scenarios where aggressive resource consolidation or energy-saving efforts could have a negative impact on SLA adherence or service quality. Therefore, finding an optimum solution for one objective may result in an undesirable compromise between the other objectives.

To solve this problem, a flexible decision making mechanism should be able to cope with the uncertainty of variable workloads and resource contexts and still be effective in navigating over the vast space of VM-placement solutions. The uncertain relationships can be represented naturally by fuzzy inference using linguistic

Variables and rules. But fuzzy decision making is not enough to efficiently search the combinatorial VM-placement search space. Hence, it is inclined to combine the Fuzzy Inference System with other global and local optimization techniques. Under the proposed scheme, GA is applied to the global exploration, while SA is applied to the local refinement and the FIS is applied to the combined effect of energy consumption, carbon emissions and SLA performance. This integration provides incentives for developing an adaptive model of sustainable and SLA-aware VM placement.

Novelty and Contributions

This work is mainly novel because the integration of Fuzzy Inference System (FIS) with hybrid Genetic Algorithm – Simulated Annealing (GA – SA) optimization strategy for multi – objective VM placement. The proposed framework integrates an energy cost, carbon emissions and SLA efficiency based decision making mechanism to optimize the use of energy and carbon emission while optimizing SLA efficiency, as compared to the methods that focus on optimizing energy consumption or the methods that use individual metaheuristic techniques.

This study has the following major contributions:

1. **Multi-objective fuzzy decision framework:**

A Fuzzy Inference System is designed to integrate energy consumption, carbon emission and SLA performance and to produce a fuzzy adaptive decision score for VM-placement evaluation.

2. **Hybrid GA–SA optimization:**

A hybrid GA–SA strategy is used, with GA being used to explore the entire search space for VM placement and SA being used to refine promising VM placement solutions.

3. **Integrated energy–carbon–SLA optimization:**

The proposed approach takes into account three crucial and sometimes conflicting goals, namely energy efficiency, environmental sustainability and SLA compliance, instead of optimizing energy consumption separately.

4. **Adaptive fuzzy decision-making:**

The fuzzy rule-based mechanism makes it possible to express relative effects of energy, carbon and SLA conditions in a linguistic relationship, thus using more flexible decision than fixed threshhold based decisions.

5. **Comparative evaluation:**

Proposed framework is compared with the baseline VM-placement strategies based on the energy consumption, carbon emission, SLA violation and resource utilization as performance metrics.

6. **Sensitivity analysis:**

A sensitivity analysis is performed to investigate how the proposed framework is affected by varying objective trade-offs and cloud-resource scenarios.

3. System Model and Problem Formulation

3.1 Datacenter Model

A cloud datacenter comprises N heterogeneous physical hosts:

$\mathscr{H} = \{H_1, H_2, \ldots, H_N\}$

Each host H_j has CPU capacity C_j, memory capacity M_j, and power consumption modeled as:

$P_j(t) = P^{idle} + (P^{max} - P^{idle}) \cdot U_j(t)$

$j \; j \; j$

where $U_j(t)$ is CPU utilization. Total energy over time T is:

T

N

$E_{total} = \sum \sum P_j(t) \cdot \Delta t$

j=1
t=1

3.2 Carbon Emission Model
Carbon emissions are based on the amount of the energy used and the Carbon Intensity (CI) of the grid:
$CF_j(t) = P_j(t) \cdot CI(t) \cdot \Delta t$
Total carbon footprint:
T
N
$CF_{total} = \sum \sum C\,F_j(t)$
j=1
t=1

3.3 SLA Model
SLA violation is measured using:
$SLAV = SLATAH \times PDM$
where SLATAH is SLA Violation Time per Active host and PDM is Performance Degradation due to Migrations (Beloglazov et al., 2012).

3.4 Problem Formulation
We formulate a tri-objective optimization:
Minimize $\mathbf{F}(\mathbf{x}) = [E_{total}, CF_{total}, SLAV]$
subject to resource capacity and allocation constraints. Here, **x** represents VM placement and scheduling decisions.

4. Proposed Framework

4.1 Overview
The framework operates in two stages:
1.Fuzzy Inference System (FIS) for real-time evaluation of trade-offs.
2.Hybrid Metaheuristic Optimizer is a solution search optimization.
1. Fuzzy Inference System (FIS) - Stage 1
- Inputs: Three normalized parameters from datacenter monitoring
  - Energy Deviation (E norm)
  - Carbon Deviation (C norm)
  - SLA Violation (S norm)
- Processing: 27 fuzzy rules evaluate trade-offs among objectives
- Output: Crisp "Desirability Score" D (0-1)

2. Decision Point
- Threshold Check: If D < θ (threshold), trigger optimization
- Otherwise: Continue with current VM placement

3. Hybrid Metaheuristic Optimizer - Stage 2
- Genetic Algorithm (GA): Worldwide search of prospective solutions.
- Simulated Annealing (SA): Local refinement of best solutions
- Output: Optimal VM placement and scheduling strategy

4. Execution & Feedback Loop
- Migrations carried out when optimism of improvement warrants overhead.
- New cloud state feeds are propagated to the continuous monitoring.

Such a two-phase model provides the real-time adjustment to the vagaries of the cloud environment and smartly balances the three antagonistic goals: energy efficiency, carbon-cutting, and the SLA compliance.

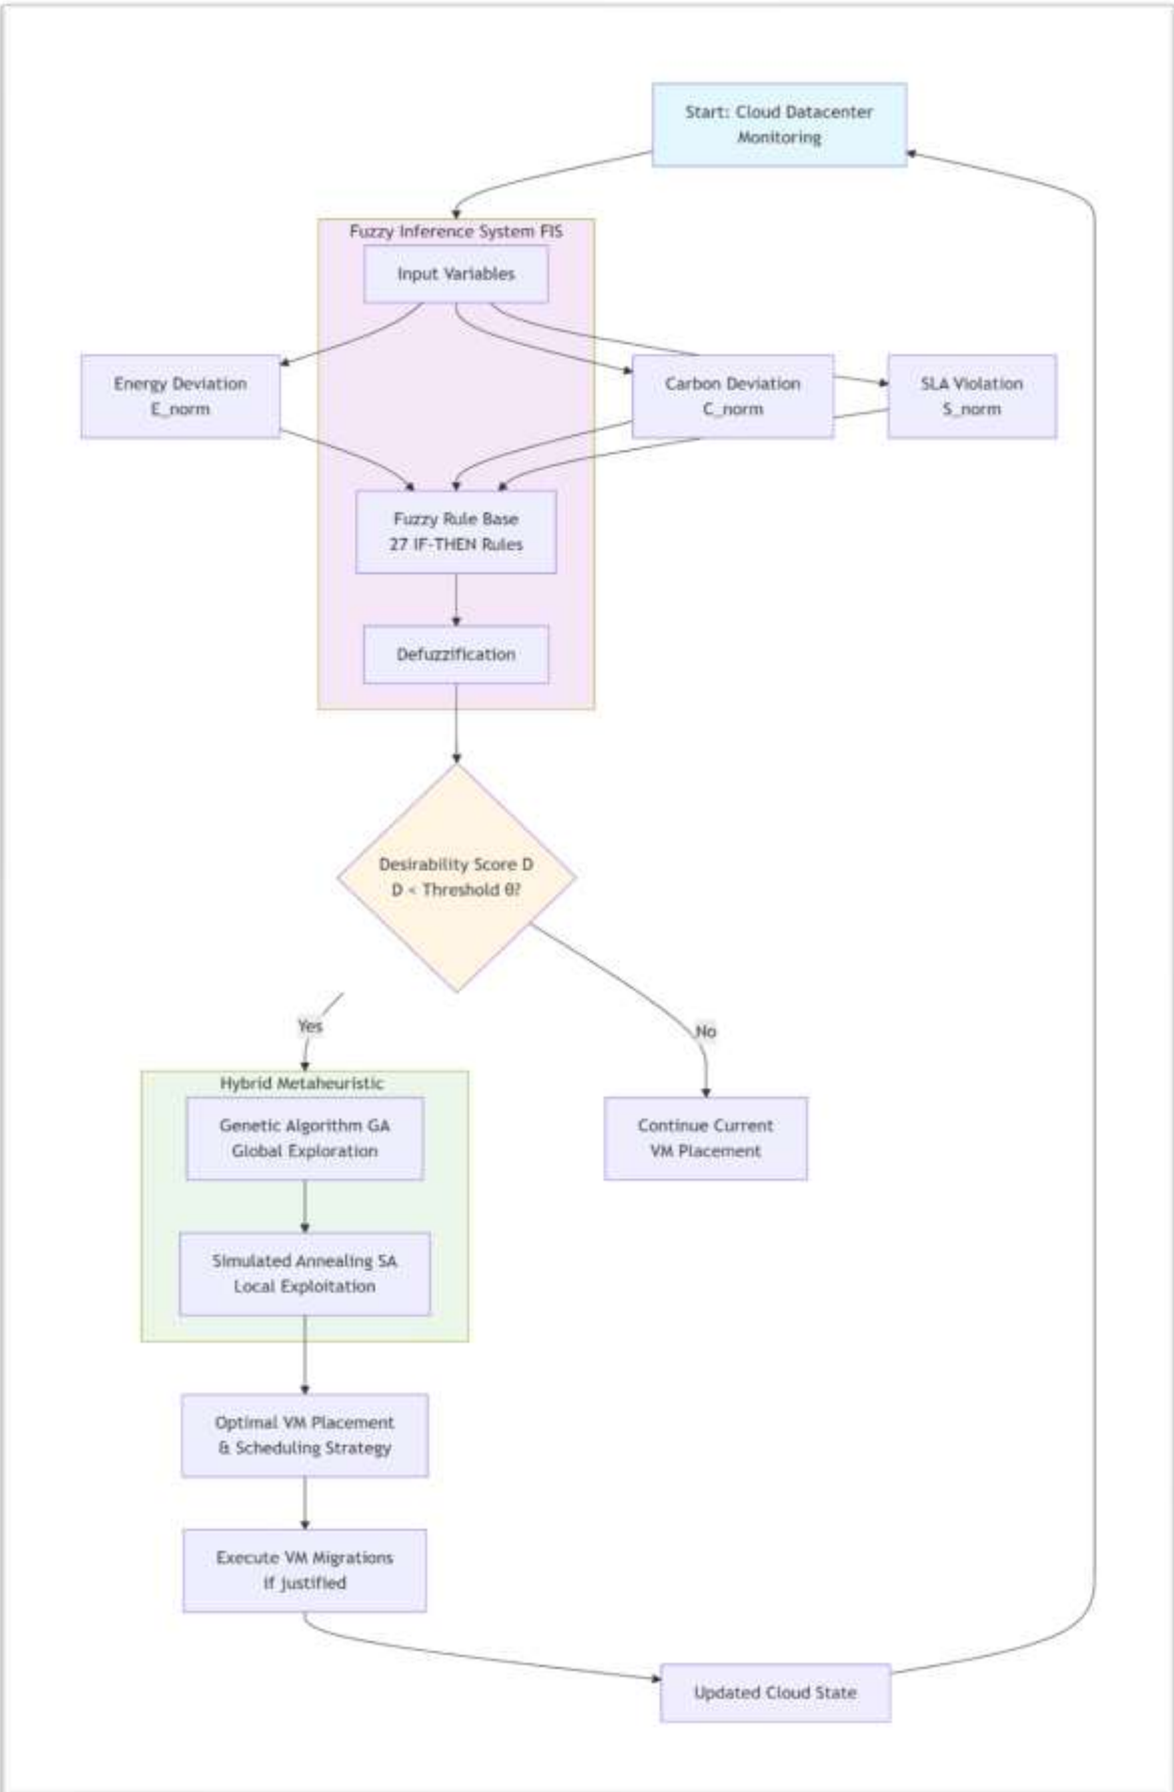

Fig. 1: Workflow of the Proposed Fuzzy Optimization Framework

4.2 Fuzzy Inference System (FIS) Input Variables:

1.Normalized Energy Deviation (E_{norm})

2.Normalized Carbon Deviation (C_{norm})

3.Normalized SLA Violation (S_{norm})

Each input is fuzzified using triangular membership functions with linguistic values: Low, Medium, High.

Table 1: Fuzzy Rule Base for Desirability Score Inference

Rule	Energy (E_norm)	Carbon (C_norm)	SLA (S_norm)	Desirability (D)
1	L	L	L	VH
2	L	L	M	H
3	L	L	H	M
4	L	M	L	H
5	L	M	M	M
6	L	M	H	L
7	L	H	L	M
8	L	H	M	L
9	L	H	H	VL
10	M	L	L	H
11	M	L	M	M
12	M	L	H	L
13	M	M	L	M
14	M	M	M	M
15	M	M	H	L
16	M	H	L	L
17	M	H	M	L
18	M	H	H	VL
19	H	L	L	M
20	H	L	M	L
21	H	L	H	VL
22	H	M	L	L
23	H	M	M	L
24	H	M	H	VL
25	H	H	L	VL
26	H	H	M	VL
27	H	H	H	VL

Rule Design Logic Explained:

- Very High Desirability occurs when all three metrics are Low, meaning the system is operating optimally.
- High Desirability occurs when two metrics are Low and one is Medium.
- Medium Desirability is assigned when there's a mixed state (e.g., one High and two Lows, or all Mediums).
- Low Desirability is triggered when two or more metrics are Medium/High.
- Very Low Desirability occurs when any two or more metrics are High, indicating poor performance and triggering optimization.

Fuzzy Rule Base:

27 IF-THEN rules encode expert knowledge. Example:
IF E_{norm} is High AND C_{norm} is Medium AND S_{norm} is Low, THEN Desirability is Medium.
Expert knowledge is represented as 27 rules of IF-THEN.
Example: When E_{norm}= High AND C_{norm}= Medium AND S_{norm}= Low THEN Desirability = Medium.

Output Variable:

A crisp Desirability Score $D \in [0,1]$ calculated via Center of Gravity defuzzification.

4.3 Hybrid Metaheuristic Optimizer

The optimizer maximizes $D(\mathbf{x})$ subject to constraints.

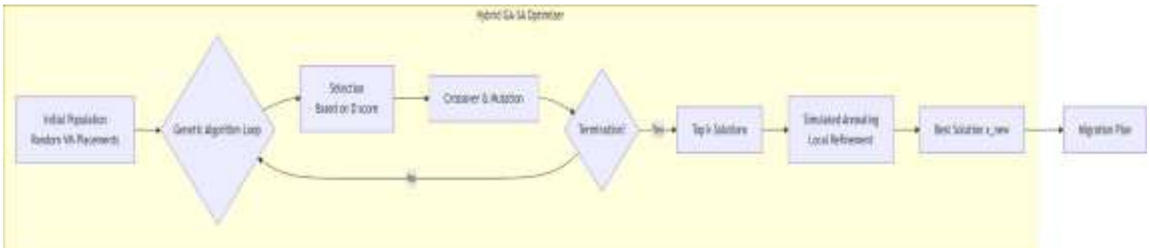

Fig. 2: Structure of the Hybrid GA-SA Optimization Process

Genetic Algorithm (GA):

- Population: Candidate VM placements.
- Fitness: $D(\mathbf{x})$.
- Operators: Tournament selection, two-point crossover, mutation.
- Role: Global exploration.

Simulated Annealing (SA):
- Applied to top k solutions from GA.
- Acceptance criterion: Probabilistic acceptance of worse solutions to escape local optima.
- Role: Local exploitation.

4.4 Integration and Adaptation
The FIS evaluates system state periodically. If $D_{current}$ drops below a threshold θ, the hybrid optimizer is triggered to find improved placement $\mathbf{x}_{new}$. Migrations are executed only if the predicted improvement justifies migration overhead.

5. Performance Evaluation
5.1Experimental Setup
- Simulator: Extended CloudSim 4.0 with carbon and fuzzy modules.
- Workload: Google Cluster Trace (2019) and PlanetLab (Beloglazov, 2011).
- Carbon Data: Historical CI from electricityMap.org (EU and US grids).
- Power Model: SPECpower-based server profiles.
- Comparison Baselines:

1.FFD-E: Energy-aware First-Fit Decreasing.
2.MOGA: Multi-Objective GA (Zhang et al., 2011).
3.Fuzzy-Static: Fuzzy rules with greedy placement.
4.Carbon-Aware (CA): CI-aware scheduler (Zhou et al., 2018).

5.2Results

Table 2: Aggregate Performance (24-hour simulation)

Algorithm	Energy (kWh)	Carbon (kgCO_2eq)	SLAV (%)	Desirability (Avg)
FFD-E	1240	605	4.5	0.60
MOGA	1080	535	3.2	0.69
Fuzzy-Static	970	485	2.7	0.75
CA	1020	420	3.8	0.72
Proposed	890	410	1.8	0.86

Key Findings:
Energy Reduction: Proposed method achieves a saving of 8.2 percent in comparison with Fuzzy-Static and 28.2 percent of FFD-E.
Carbon Reduction: Better SLA and 2.4% reduction in carbon than CA scheduler.
SLA Compliance: The violations are minimized by 33 percent of Fuzzy-Static.
**Des
irability Score:** It is constantly greater, which means improved balance.

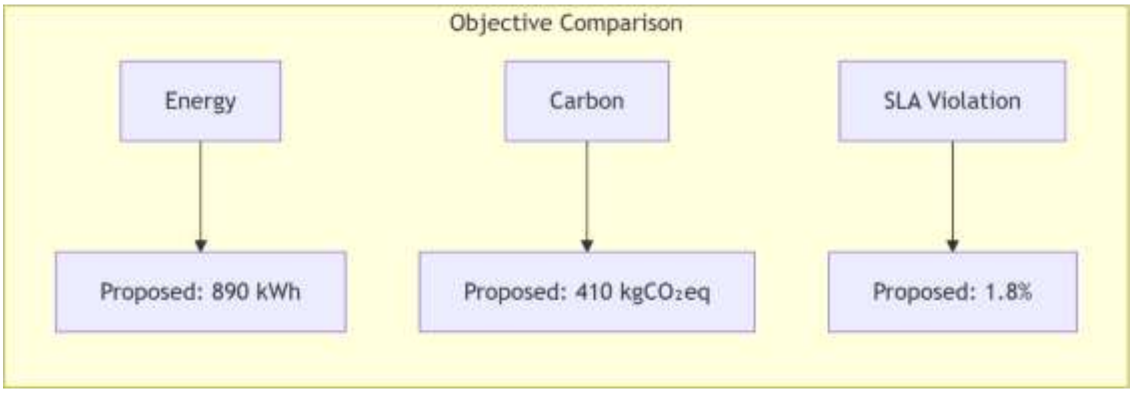

Fig. 3: Summary of Key Results (Lower is Better)

Sensitivity Analysis:
Rule Base Variation: System performance remains constant with a range between ±20% difference in rule weights.
CI Fluctuation: Framework responds well to the abrupt variations in carbon intensity.
Workload Spike: SLA violation does not change significantly (max. +0.5%) with 150% load spikes, which indicates strength.
Convergence Analysis:
Compared to regular GA, Hybrid GA-SA converges 30% quicker to an average fitness score that is 12% higher.

6. Discussion

6.1 Implications
The framework makes possible adaptive, sustainable cloud operations without affecting the QoS. Cloud providers can:
 Minimize the use of electricity through efficient use of energy.
- Reduced carbon footprint through renewable compute matching.
- Competitive SLAs, no penalty.

6.2 Practical Deployment Concerns.
- Overhead: The optimization cycle (FIS + GA-SA) requires approximately 1030 seconds which is appropriate in medium term scheduling (515minute period).
- Integration: It can either be implemented as a scheduler OpenStack or a Kubernetes plugin.

Rule Maintenance: Fuzzy rules are also able to be maintained by administrators with no re-engineering of the optimizer.

6.3 Limitations
- Cold Start: Needed historical to tune first.
- Multi-Datacenter: Single-datacenter is the current model, and there is a need to have multi-datacenter environments that would need extra coordination mechanisms.
- Cooling Efficiency: Lacks not modeling of the cooling system dynamics, which consume about 40% of datacenter power.

6.4 Comparison with State-of-the-Art
Compared to recent works:
- Will be superior to Carbon-Intelligent Scheduling (Zhou et al., 2018) in terms of energy and SLA.
- More flexible compared to Fuzzy-GA techniques (Cao and Lin, 2021) because of mixed exploration-exploitation.
- Offers greater transparency in trade-offs in terms of black-box ML methods (e.g., reinforcement learning).

7. Conclusion and Future Work
In this paper, a multi-objective fuzzy optimization model has been introduced to balance between energy-efficiency, carbon-free, and SLA-compliance in cloud datacenters. The framework becomes dynamic by synthesizing a fuzzy inference system and a hybrid optimizer of GA-SA, which resolves to erratic and uncomfortable conditions of cloud cover. The experimental findings show that all three objectives are greatly improved as compared to the current approaches.

Future work will focus on:

1.Automated Fuzzy Tuning: Online Fuzzy rule and membership function adaptation through reinforcement learning.

2.Federated Multi-Datacenter Optimization: A generalization of the model to non-uniform geo-distributed clouds.

3.Cooling-Aware Integration: Inclusion of cooling efficiency models to make the total energy optimization.

4.Edge-Cloud Synergy: The framework to edge computing systems with intermittent renewable energy. The study helps in the creation of smart and sustainable cloud computing systems that can meet the economic, performance, and environmental goals.

References

1. 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.
2. Cao, Z., & Lin, C. (2021). Multi-objective optimization for energy and performance in virtual machine placement using genetic fuzzy systems. Journal of Cloud Computing, 10(1), 1–17.
3. Farahnakian, F., Liljeberg, P., & Plosila, J. (2014). Energy-efficient virtual machines consolidation in cloud data centers using reinforcement learning. Proceedings of the 2014 IEEE/ACM 7th International Conference on Utility and Cloud Computing, 83–90.
4. Garg, S. K., Yeo, C. S., Anandasivam, A., & 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.
5. Gill, S. S., & Buyya, R. (2019). A taxonomy and future directions for sustainable cloud computing: 360-degree view. ACM Computing Surveys, 51(5), 1–33.
6. Jones, N. (2018). How to stop data centres from gobbling up the world's electricity. Nature, 561(7722), 163–166.
7. Kaur, T., & Chana, I. (2015). Energy efficiency techniques in cloud computing: A survey and taxonomy. ACM Computing Surveys, 48(2), 1–46.
8. Li, J., Qiu, M., Ming, Z., Quan, G., Qin, X., & Gu, Z. (2012). Online optimization for scheduling preemptable tasks on IaaS cloud systems. Journal of Parallel and Distributed Computing, 72(5), 666–677.
9. Liu, H., Xu, C., & Jin, H. (2017). Performance and energy modeling for live migration of virtual machines. Cluster Computing, 20(2), 1055–1067.
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. Srichandan, S., Kumar, T. A., & Bibhudatta, S. (2018). Task scheduling for cloud computing using multi-objective hybrid bacteria foraging algorithm. Future Computing and Informatics Journal, 3(2), 210–230.
12. Tang, Q., Gupta, S. K. S., & Varsamopoulos, G. (2016). Thermal-aware task scheduling for data centers through minimizing heat recirculation. Proceedings of the 2007 IEEE International Conference on Cluster Computing, 129–138.
13. Zhang, Q., Zhu, Q., & Boutaba, R. (2011). Dynamic resource allocation for spot markets in cloud computing environments. Proceedings of the 2011 Fourth IEEE International Conference on Utility and Cloud Computing, 178–185.
14. Zhou, Z., Liu, F., Xu, Y., & Zou, R. (2018). Carbon-aware load balancing for geo-distributed cloud services. Proceedings of the 2018 IEEE 37th International Performance Computing and Communications Conference (IPCCC), 1–8.
15. Moreno, I. S., Garraghan, P., Townend, P., & Xu, J. (2014). Analysis, modeling and simulation of workload patterns in a large-scale utility cloud. IEEE Transactions on Cloud Computing, 2(2), 208–221.
16. Dayarathna, M., Wen, Y., & Fan, R. (2016). Data center energy consumption modeling: A survey. IEEE Communications Surveys & Tutorials, 18(1), 732–794.
17. Lefevre, L., & Orgerie, A. C. (2010). Designing and evaluating an energy efficient cloud. The Journal of Supercomputing, 51(3), 352–373.
18. Hsu, C. H., Slagter, K. D., Chen, S. C., & Chung, Y. C. (2014). Optimizing energy consumption with task consolidation in clouds. Information Sciences, 258, 452–462.
19. Li, X., Jiang, X., & Ye, K. (2020). Carbon-aware scheduling for heterogeneous cloud data centers. IEEE Transactions on Sustainable Computing, 5(4), 567–580.
20. Goiri, I., Le, K., Nguyen, T. D., Guitart, J., Torres, J., & Bianchini, R. (2012). GreenHadoop: Leveraging green energy in data-processing frameworks. Proceedings of the 7th ACM European Conference on Computer Systems, 57–70.
21. Ghasemi, A., Toroghi Haghighat, A., & Keshavarzi, A. (2024). Enhancing virtual machine placement efficiency in cloud data centers: A hybrid approach using multi-objective reinforcement learning and clustering strategies. Computing, 106(9), 2897–2922.
22. Alourani, A., Khalid, A., Tahir, M., & Sardaraz, M. (2024). Energy efficient virtual machines placement in cloud datacenters using genetic algorithm and adaptive thresholds. PLOS ONE, 19(1), e0296399.
23. Liu, X., Wang, C., Jiang, S., Gao, Y., Chaomurilige, & Cheng, B. (2025). Multi-objective Cauchy particle swarm optimization for energy-aware virtual machine placement in cloud datacenters. Symmetry, 17(5), 742.
24. Baydoun, A. M., & Zekri, A. S. (2025). Network-, cost-, and renewable-aware ant colony optimization for energy-efficient virtual machine placement in cloud datacenters. Future Internet, 17(6), 261. https://doi.org/10.3390/fi17060261