MAIZX: A Carbon-Aware Framework for Optimizing Cloud Computing Emissions

Federico Ruilova Alfaro
fedra@kth.se
KTH Royal Institute of Technology
Stockholm, Stockholm, Sweeden;
Technische Universität Berlin
Berlin, Berlin, Germany

Ernst Gunnar Gran
ernst.g.gran@ntnu.no
Norwegian University of Science and Technology (NTNU)
Gjøvik, Norway

Sven-Arne Reinemo
svenar@simula.no
Simula Metropolitan Centre for Digital Engineering
Oslo, Norway

Abstract

Cloud computing drives innovation but also poses significant environmental challenges due to its high energy consumption and carbon emissions. Data centers account for 2-4% of global energy usage, and the ICT sector's share of electricity consumption is projected to reach 40% by 2040.As the goal of achieving net-zero emissions by 2050 becomes increasingly urgent, there is a growing need for more efficient and transparent solutions, particularly for private cloud infrastructures, which are utilized by 87% of organizations, despite the dominance of public-cloud systems.

This study evaluates the MAIZX framework, designed to optimize cloud operations and reduce carbon footprint by dynamically ranking resources, including data centers, edge computing nodes, and multi-cloud environments, based on real-time and forecasted carbon intensity, Power Usage Effectiveness (PUE), and energy consumption. Leveraging a flexible ranking algorithm, MAIZX achieved an 85.68% reduction in CO_2 emissions compared to baseline hypervisor operations. Tested across geographically distributed data centers, the framework demonstrates scalability and effectiveness, directly interfacing with hypervisors to optimize workloads in private, hybrid, and multi-cloud environments. MAIZX integrates real-time data on carbon intensity, power consumption, and carbon footprint, as well as forecasted values, into cloud management, providing a robust tool for enhancing climate performance potential while maintaining operational efficiency.

Keywords

carbon reduction in cloud, carbon-aware computing, energy-aware clouds, private cloud optimization, sustainable cloud computing,carbon performance potential

ACM Reference Format:
Federico Ruilova Alfaro, Ernst Gunnar Gran, and Sven-Arne Reinemo. 2024. MAIZX: A Carbon-Aware Framework for Optimizing Cloud Computing Emissions. In *Proceedings of 1st International Workshop on Low Carbon Computing (Conference acronym 'LOCO 24')*. ACM, New York, NY, USA, 3 pages.

1 Introduction

Cloud computing drives innovation across sectors but raises concerns about its environmental impact, particularly due to high energy consumption and carbon emissions.Currently Data centers account for 2-4% of the global energy usage, with the broader ICT sector consuming 6% [20],[13],[2]. This is projected to rise to nearly 40% of total global electricity consumption by 2040 [1], underscoring the urgent need for sustainable solutions, carbon-aware computing and energy-aware frameworks[15], [24],[19], [16],[14], [11]. Achieving net-zero emissions by 2050 is critical for limiting global warming, and while cloud providers focus on carbon neutrality through grid carbon intensity awareness and renewable energy integration, more transparent and effective methods are needed, especially for private cloud setups, used by 87% of organizations[9],[4],[3].

This research explores the MAIZX framework to revisit its potential and evaluate its climate performance[21]; the framework uses a ranking algorithm that allocates resources based on scores of computing nodes. The framework's scalability and effectiveness were empirically tested by implementations across geographically distributed data centers and validated via simulations. By integrating these metrics into cloud management, MAIZX offers a robust tool for enhancing climate performance and assessing environmental impact in private, hybrid and multi-cloud approaches.

2 MAIZX Ranking Algorithm

The MAIZX framework uses a hybrid architecture that centralizes control while distributed agents collect power consumption and carbon intensity data, both at each distributed node and at the core [21].

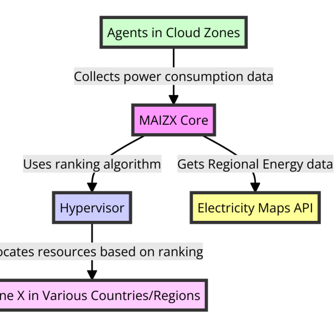

Figure 1: MAIZX Framework Architecture

The ranking algorithm dynamically allocates workloads to nodes with the lowest carbon intensity, prioritizing environmental impact without compromising performance. Centralized components coordinate with the hypervisor using carbon efficiency and power data as it could be observed in Figure 1.

Conference acronym 'LOCO 24', December 03,2024, Glasgow, Scotland, United Kingdom
2024. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

2.1 Key Functionalities of Carbon-Aware allocation

Agents gather real-time energy and carbon intensity data, supporting carbon footprint calculations and forecasting.

The algorithm evaluates nodes on carbon footprint and efficiency, optimizing workload distribution to reduce emissions. It integrates with hypervisors like OpenNebula [18] for efficient scheduling.

The ranking system calculates a node's efficiency by considering key environmental and operational parameters. The algorithm, *MAIZ_RANKING*, is defined as:

$$MAIZ\_RANKING = w_1 \cdot \text{CFP} + w_2 \cdot \text{FCFP} \quad + w_3 \cdot \text{CP\_RATIO} + w_4 \cdot \text{SCHEDULE\_WEIGHT} \tag{1}$$

In this equation, CFP refers to the node's Carbon Footprint, while FCFP denotes the Forecasted Carbon Footprint based on historical data. The Computing Power Ratio (CP_RATIO) reflects the node's energy efficiency, and the Scheduling Weight (SCHEDULE_WEIGHT) accounts for workload priorities and deadlines. Adjustable weights (w_1, w_2, w_3, w_4) are assigned to each factor, enabling the framework to balance environmental impact, performance, and operational needs. This flexible approach ensures that MAIZX can maintain a balance between sustainability and efficiency across various cloud environments.

3 Methodology and Experimental Design

This study evaluates the MAIZX framework's climate performance in private and multi-cloud oriented environments by optimizing operations based on data center locations, hardware usage, and regional energy profiles. The assessment builds on prior MAIZX research[21], using 2022 carbon intensity data[8] to simulate various scenarios: the Baseline Scenario evenly distributes loads without any consideration of carbon intensity or footprint data, serving as a comparison for carbon footprint analysis.; Scenario A directs all computing power to the node with the lowest carbon intensity; Scenario B concentrates tasks on a single node while powering off others to measure energy savings; and Scenario C dynamically shifts loads based on daily carbon intensity fluctuations, highlighting MAIZX's adaptability for reducing emissions.

Data collection consists of power consumption measured every 20 seconds, while carbon intensity is recorded hourly across three regions: Spain, the Netherlands, and Germany. The carbon footprint for each node across the scenarios is calculated using a standard methodology [7, 10], applying the formula:

$$CF = EC \times PUE \times CI \tag{2}$$

where **CF** is the carbon footprint, **EC** is energy consumption, **PUE** is Power Usage Effectiveness, and **CI** is carbon intensity

In order to calculate the climate performance potential (CPP), the impact forecast tool was used, together with the corresponding EU taxonomy group for ICT. The calculation model uses functional unit or (FU) to calculate life cicle analys [12], [6],[23], [22].

4 Results and Analysis

In Scenario C (active load-shifting over one year), the MAIZX framework reduces CO_2 emissions by 85.68% compared to the baseline, optimizing workloads using real-time carbon intensity data, (Figure 2). Each unit, consisting of 60 servers in a 3-node private cloud, reduces emissions by 713.5 kg of CO_2 annually. The main difference between Scenario B and Scenario C is the use of real-time carbon data: Scenario B evenly distributes workloads without considering carbon intensity, whereas Scenario C actively shifts workloads to the nodes with the lowest carbon intensity once, and scenario A as well but leaving the other nodes available. While both scenarios B and C achieve similar reductions, Scenario C is more sustainable long-term due to its dynamic response to fluctuations in carbon intensity, consistently maintaining lower emissions when variations occur.

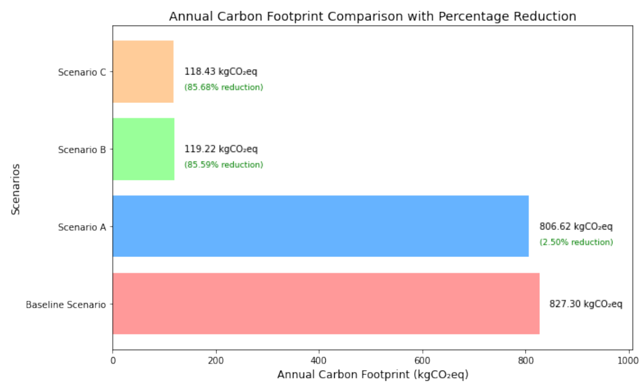

Figure 2: MAIZX Framework Architecture

To evaluate the broader impact, if 1% of the EU Taxonomy target for data-driven climate change monitoring and ICT data processing is considered, it totals 19.754 Mt CO_2eq [5],[12]. Over a 10-year period, implementing the MAIZX framework with its 85% reduction capability would yield the following:

- Total reduction target: 19.754 Mt CO_2eq (19,754,000,000 kg).
- Annual CO_2 reduction per unit: 713.5 kg.
- Units required: 27,686,054.

This showcases the scalability of the MAIZX framework for reducing emissions in large-scale cloud operations, potentially reaching 19.754 Mt CO_2eq in 10 years targeting shifted units. The results highlight MAIZX's potential for substantial environmental and cost savings, particularly in private and multi-cloud environments with optimized power consumption and associated carbon footprint.

5 Conclusion

The framework demonstrates significant potential to reduce CO_2 emissions, particularly in private or multi-cloud setups. Despite being tested in regions with non-renewable energy matrices, MAIZX achieved considerable emission reductions. Empirical data validated the framework's carbon footprint calculations, affirming its methodology and accuracy. Conservative 10-year projections suggest that MAIZX could reduce emissions by 20 Mt CO_2eq—equivalent to planting 90 million trees or removing 2.44 million cars from the road annually. Additionally, MAIZX provides significant eco-cost savings, including €3 billion in human health impacts, €4.65 billion in eco-toxicity, and €2.63 billion in carbon footprint-related costs[12]. These findings underscore MAIZX's potential to support sustainability goals [17] in the ICT sector by optimizing cloud infrastructure for sustainability.

6 Bibliography

References

[1] International Energy Agency. 2021. Net Zero by 2050 - A Roadmap for the Global Energy Sector. (2021).

[2] Ehsan Ahvar, Anne-Cecile Orgerie, and Adrien Lebre. 2022. Estimating Energy Consumption of Cloud, Fog, and Edge Computing Infrastructures. *IEEE Transactions on Sustainable Computing* 7, 2 (April 2022), 277–288. https://doi.org/10.1109/TSUSC.2019.2905900

[3] Rohan Arora, Umamaheswari Devi, Tamar Eilam, Aanchal Goyal, Chandra Narayanaswami, and Pritish Parida. 2023. Towards Carbon Footprint Management in Hybrid Multicloud. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems*. ACM, Boston MA USA, 1–7. https://doi.org/10.1145/3604930.3605721

[4] Cisco. 2022. Global Hybrid Cloud Trends Report. https://www.cisco.com/c/en/us/solutions/hybrid-cloud/2022-trends.html

[5] European Commission European Commission. 2024. EU Taxonomy Navigator. https://ec.europa.eu/sustainable-finance-taxonomy/activities

[6] ecocostsvalue.com. 2024. Life Cycle Assessment. https://www.ecocostsvalue.com/lca/

[7] Tamar Eilam. 2021. Towards transparent and trustworthy cloud carbon accounting. In *Proceedings of the 22nd International Middleware Conference: Extended Abstracts*. ACM, Virtual Event Canada, 1–5. https://doi.org/10.1145/3501255.3501408

[8] electricitymaps.com. 2024. Electricity Maps | Reduce carbon emissions with actionable electricity data. https://www.electricitymaps.com/

[9] Flexera. 2024. Flexera 2023 State of the Cloud | Report. https://info.flexera.com/CM-REPORT-State-of-the-Cloud

[10] GESI. 2024. GeSI. https://www.gesi.org/research/ict-sector-guidance-built-on-the-ghg-protocol-product-life-cycle-accounting-and-reporting-standard

[11] Walid A. Hanafy, Roozbeh Bostandoost, Noman Bashir, David Irwin, Mohammad Hajiesmaili, and Prashant Shenoy. 2023. The War of the Efficiencies: Understanding the Tension between Carbon and Energy Optimization. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems*. ACM, Boston MA USA, 1–7. https://doi.org/10.1145/3604930.3605709

[12] impact forecast.com. 2024. Impact Forecast: Calculate, improve and validate your climate impact. https://impact-forecast.com/

[13] The Independent. 2016. Global warming: Data centres to consume three times as much energy in next decade, experts warn. https://www.independent.co.uk/climate-change/news/global-warming-data-centres-to-consume-three-times-as-much-energy-in-next-decade-experts-warn-a6830086.html Section: Climate.

[14] Loïc Lannelongue and Michael Inouye. 2023. Carbon footprint estimation for computational research. *Nature Reviews Methods Primers* 3, 1 (Feb. 2023), 9. https://doi.org/10.1038/s43586-023-00202-5

[15] Phil Laplante and Jeffrey Voas. 2023. "Frameworking" Carbon-Aware Computing Research. *Computer* 56, 5 (May 2023), 105–108. https://doi.org/10.1109/mc.2023.3240482 Publisher: Institute of Electrical and Electronics Engineers (IEEE).

[16] Diptyaroop Maji, Ben Pfaff, Vipin P R, Rajagopal Sreenivasan, Victor Firoiu, Sreeram Iyer, Colleen Josephson, Zhelong Pan, and Ramesh K Sitaraman. 2023. Bringing Carbon Awareness to Multi-cloud Application Delivery. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems*. ACM, Boston MA USA, 1–6. https://doi.org/10.1145/3604930.3605711

[17] United Nations. 2023. The Sustainable Development Goals Report 2023: Special Edition. https://unstats.un.org/sdgs/report/2023/

[18] OpenNebula. 2023. Scheduler Configuration — OpenNebula 6.6.3 documentation. https://docs.opennebula.io/6.6/installation_and_configuration/opennebula_services/scheduler.html#schg

[19] Ana Radovanović, Ross Koningstein, Ian Schneider, Bokan Chen, Alexandre Duarte, Binz Roy, Diyue Xiao, Maya Haridasan, Patrick Hung, Nick Care, Saurav Talukdar, Eric Mullen, Kendal Smith, MariEllen Cottman, and Walfredo Cirne. 2023. Carbon-Aware Computing for Datacenters. *IEEE Transactions on Power Systems* 38, 2 (March 2023), 1270–1280. https://doi.org/10.1109/TPWRS.2022.3173250

[20] Aimee Ross and Lorna Christie. 2023. Energy consumption of ICT. (Feb. 2023). https://post.parliament.uk/research-briefings/post-pn-0677/

[21] Federico Ruilova Alfaro. 2024. Towards Sustainable Cloud Computing: A Systemic Framework for Leveraging Regional Energy Data to Empower Carbon-Aware Computing: Optimizing Cloud Ecosystems: The MAIZX Framework and Its Ranking Algorithm for Lowering Carbon Footprint. *Dissertation* (2024).

[22] Daniel Valenzuela and Danijel Višević. 2021. How to measure the climate performance potential | World Fund. https://www.worldfund.vc/knowledge/how-to-measure-the-climate-performance-potential-of-startups

[23] J.G. Vogtländer. 2010. *A Practical Guide to LCA for Students, Designers and Business Managers: Cradle-to-grave and Cradle-to-cradle*. VSSD. https://books.google.de/books?id=EkcDaAEACAAJ

[24] Jackson Woodruff, David Schall, Michael F.P. O'Boyle, and Christopher Woodruff. 2023. When Does Saving Power Save the Planet?. In *Proceedings of the 2nd Workshop on Sustainable Computer Systems*. ACM, Boston MA USA, 1–6. https://doi.org/10.1145/3604930.3605719

Received 24 September 2024