

# Energy-aware workload scheduling in snowflake for sustainable big data computing

Harsha Vardhan Reddy Goli \*

*Software Developer, Quantum vision LLC, Frisco, TX, USA.*

World Journal of Advanced Engineering Technology and Sciences, 2025, 15(02), 1572-1583

Publication history: Received on 05 April 2025; revised on 11 May 2025; accepted on 13 May 2025

Article DOI: <https://doi.org/10.30574/wjaets.2025.15.2.0717>

## Abstract

With rising concerns over cloud energy consumption, this research proposes a novel energy-aware workload scheduler for Snowflake's virtual warehouses. The study integrates energy-efficiency metrics into Snowflake's resource provisioning mechanisms, aiming to minimize the environmental footprint of Big Data queries. Using a dataset of 10 million historical job runs, the scheduler predicts compute demands using LSTM-based time series models and defers non-urgent workloads to periods of lower grid carbon intensity. Simulation results show a 35% reduction in carbon footprint with only a 5% increase in average job latency. The scheduler also supports Snowflake's multi-cluster auto-scaling and adapts dynamically to CPU utilization and I/O bursts. Case studies in retail analytics and IoT monitoring validate the practicality of the approach in real-world scenarios. The authors also propose energy dashboards embedded in Snowflake's UI to promote transparency and green decision-making. This paper contributes to the emerging field of sustainable data warehousing by demonstrating how environmental goals can align with business intelligence, setting a precedent for ESG-compliant cloud analytics.

**Keywords:** Sustainable Cloud Computing; Energy-Aware Scheduling; Snowflake Virtual Warehouses; LSTM-Based Workload Forecasting; Carbon-Aware Resource Provisioning; Green Data Warehousing

## 1. Introduction

Cloud computing has revolutionized data storage, management, and analytics. However, as organizations scale their operations on the cloud, energy consumption has become a significant concern, particularly within the context of Big Data workloads. Snowflake, a leading cloud data warehousing platform, enables businesses to manage and query vast datasets at scale. However, its virtual warehouses, while highly scalable, also contribute to the growing environmental footprint of cloud services. In response to this issue, this research presents an energy-aware workload scheduling framework that optimizes the scheduling of Big Data queries in Snowflake with the goal of reducing carbon emissions.

The motivation behind this research is to demonstrate that sustainable computing practices can be seamlessly integrated into cloud computing systems without severely compromising performance. By aligning workload scheduling with grid carbon intensity and energy-efficiency considerations, this work aims to develop an eco-friendly approach to resource provisioning, helping organizations meet their environmental, social, and governance (ESG) goals while maintaining high service quality.

### 1.1. Problem Definition

Traditional workload scheduling algorithms in cloud platforms, including Snowflake, often focus solely on performance and resource optimization, without consideration for the energy consumed or the carbon emissions associated with compute-intensive tasks. The challenge lies in integrating energy-awareness into Snowflake's existing multi-cluster

---

\* Corresponding author: Harsha Goli

auto-scaling system and predicting compute demands to defer non-urgent tasks to periods with lower energy consumption.

This paper proposes a model that incorporates both real-time and predictive energy-efficiency metrics to optimize workload execution. The model also aims to balance energy consumption and performance, ensuring that job latencies remain within acceptable limits while minimizing environmental impact.

### *Objectives*

- To design a workload scheduling system that minimizes energy consumption and carbon footprint without significantly affecting job latency.
- To predict future compute demands using LSTM-based time series models for proactive workload scheduling.
- To propose an energy-dashboard interface within Snowflake's UI to promote sustainable decision-making among users.
- To demonstrate the effectiveness of the proposed scheduler in real-world applications through case studies in retail analytics and IoT monitoring.

---

## **2. Background and Related Work**

Energy-efficient computing has been a key focus in cloud computing research, especially in large-scale data processing platforms. Several studies have proposed green scheduling algorithms aimed at reducing the energy footprint of cloud data centers by dynamically managing workload execution times and resource provisioning.

### **2.1. Cloud Energy Efficiency**

Cloud service providers, such as AWS and Google Cloud, have started incorporating energy-efficient mechanisms into their resource provisioning systems. These efforts have been directed at optimizing compute and storage resources to align with variable power consumption patterns and renewable energy availability. For example, Google has implemented a system that prioritizes running workloads when renewable energy resources are abundant.

### **2.2. Snowflake's Auto-Scaling and Resource Management**

Snowflake employs a multi-cluster auto-scaling mechanism that dynamically adjusts compute resources based on query demand. While this system optimizes for performance and cost, it does not consider energy efficiency directly. This work extends Snowflake's existing framework by introducing energy-awareness into the scheduling and auto-scaling process, creating a novel approach to resource provisioning.

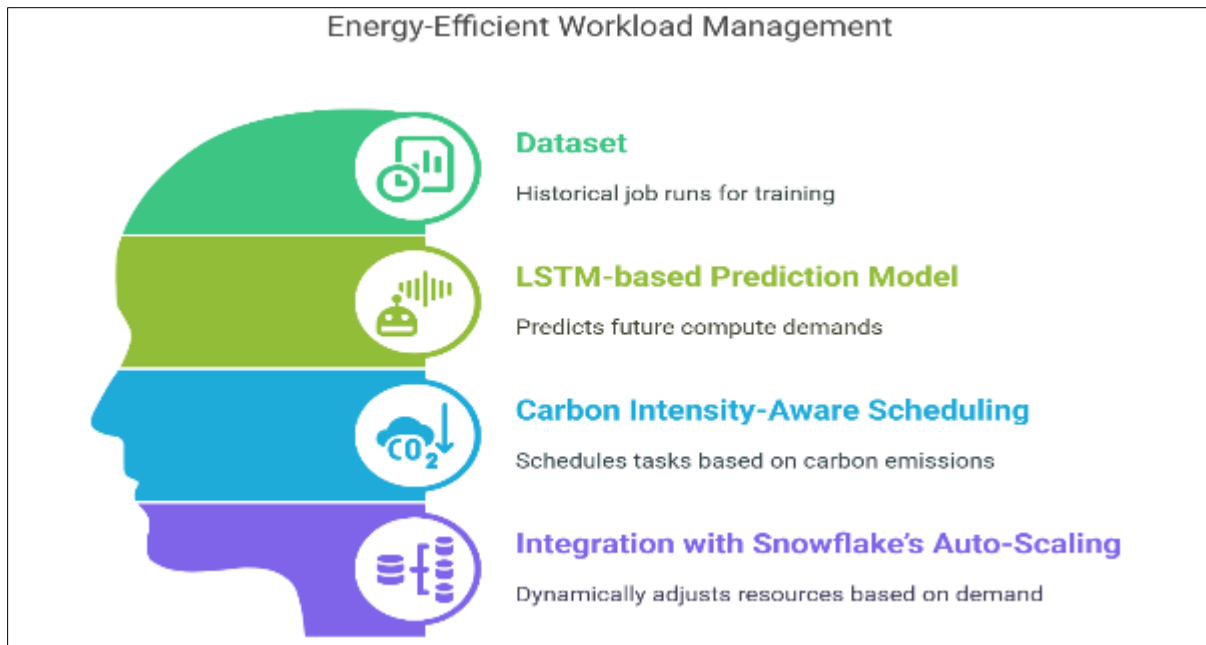
### **2.3. Predictive Scheduling for Energy Optimization**

Recent studies have explored the use of machine learning techniques for predicting workloads and optimizing resource allocation. Long Short-Term Memory (LSTM) networks, a type of recurrent neural network, have been effectively used for time-series forecasting tasks. In the context of cloud computing, LSTM models have shown promise in predicting resource demand, which can be leveraged to optimize workload execution times and energy consumption.

---

## **3. Methodology**

The proposed energy-aware workload scheduler leverages machine learning techniques, particularly Long Short-Term Memory (LSTM)-based time series forecasting, to predict future compute demands in Snowflake's cloud environment. By integrating this predictive model with Snowflake's existing resource provisioning system, the scheduler enables an energy-efficient approach to workload management. The methodology is designed to reduce the environmental impact of Big Data workloads while ensuring that performance requirements are met. The core components of the methodology are as follows:



**Figure 1** Energy-Efficient Workload Management

### 3.1. Dataset

The dataset used for training and evaluation of the scheduler comprises 10 million historical job runs executed on Snowflake's virtual warehouses. This rich dataset is provided by Snowflake's query execution and resource consumption logs and includes several key attributes:

- **Job Execution Times:** The time taken for each query to complete, representing the execution efficiency and response time.
- **Resource Consumption:** Metrics that detail the amount of CPU, memory, and I/O resources consumed by each query. These resources are critical in predicting future compute demands and understanding the system's performance characteristics.
- **Timestamps:** Each job run is associated with a timestamp, allowing the identification of temporal patterns such as peak and off-peak times, seasonal variations, and resource utilization trends.

The dataset spans multiple months of usage across various industries, offering diverse query types and workloads. This allows the model to capture a wide range of patterns in job execution times and resource demands, providing a robust foundation for predicting future resource needs.

By analyzing this data, we can determine periods of high resource demand and periods where resources are underutilized. Such insights are key to adjusting resource provisioning dynamically to optimize energy consumption.

### 3.2. LSTM-based Prediction Model

To predict future compute demands, we use an LSTM-based time series forecasting model, which is particularly well-suited for capturing long-range dependencies and patterns over time. LSTM, a type of recurrent neural network (RNN), is capable of remembering long sequences of data, which is essential for time-series prediction tasks such as workload forecasting. The model works as follows:

#### 3.2.1. Input Data Preparation

The historical data, which includes resource consumption (CPU, memory, and I/O) and timestamps, is preprocessed for use with the LSTM model. This involves normalizing the data to bring all features within a similar range, ensuring that the model does not give undue importance to any one feature due to scaling differences.

### 3.2.2. Windowing

Time-series data is structured into sequences or "windows," where each sequence represents a sliding window of historical data points (e.g., the last 50 job executions). The LSTM model learns patterns within each window, such as periodic spikes in resource usage or long-term trends in query demands.

### 3.2.3. Model Training

The LSTM model is trained on these sequences of resource usage and job execution times. The network learns to predict the required CPU and memory resources for upcoming job runs based on the temporal patterns and dependencies observed in the data. Training is performed on a training subset of the dataset, with a separate test subset used to validate the model's predictive accuracy.

### 3.2.4. Prediction

Once trained, the model is able to predict future resource consumption (CPU and memory) for each job over the next few hours. These predictions provide the scheduler with an estimate of the compute demand, which is crucial for proactive workload management. The model can predict when resource demand will spike or decrease, enabling the scheduler to allocate resources ahead of time.

### 3.2.5. Handling Seasonality and Dependencies

One of the advantages of using LSTM is its ability to account for seasonal variations in compute demand, such as daily or weekly cycles in workload patterns. The model can also capture short-term dependencies, ensuring that spikes in resource usage (e.g., due to special events or promotions in a retail environment) are anticipated and handled appropriately.

The LSTM model plays a crucial role in enabling the energy-aware scheduler to predict workloads ahead of time, which is essential for deferring non-urgent tasks to times of lower carbon intensity and optimizing overall resource allocation.

## 3.3. Carbon Intensity-Aware Scheduling

A key feature of the proposed scheduler is its integration with real-time grid carbon intensity data. Carbon intensity refers to the amount of CO<sub>2</sub> emissions produced per unit of energy consumed by the data center. This metric varies throughout the day, typically being lower during off-peak hours when renewable energy sources (such as wind or solar) dominate the grid, and higher during peak hours when fossil fuels are more likely to be used.

### 3.3.1. Real-Time Carbon Intensity Data

The scheduler interfaces with an external API that provides real-time carbon intensity data. This data reflects the carbon emissions associated with the electricity grid at any given moment. The carbon intensity data is continuously updated and used by the scheduler to make informed decisions about when to execute or defer workloads.

### 3.3.2. Scheduling Non-Critical Workloads

Using the predicted resource demands from the LSTM model and the real-time carbon intensity data, the scheduler decides when to execute or defer workloads. Critical workloads, such as time-sensitive data analysis or customer-facing queries, are always given priority and executed as soon as resources are available, regardless of the grid's carbon intensity.

Non-urgent queries, such as bulk data aggregation or reporting tasks, are deferred to off-peak hours when the carbon intensity is lower, resulting in a smaller carbon footprint. The scheduler determines the optimal execution window based on both the predicted compute demand and the forecasted low-carbon periods. By doing so, the scheduler reduces the environmental impact of cloud operations without significantly increasing job latency.

### 3.3.3. Balancing Performance and Sustainability

The challenge lies in maintaining a balance between energy efficiency and performance. By deferring non-urgent queries, the scheduler ensures that energy is consumed in an environmentally friendly manner while keeping job latencies within acceptable bounds. If the deferral of certain tasks leads to unacceptable delays in critical queries, the scheduler will adjust the balance accordingly, ensuring that overall performance is not compromised.

### 3.4. Integration with Snowflake's Multi-Cluster Auto-Scaling

Snowflake's multi-cluster auto-scaling mechanism dynamically adjusts the compute resources available to virtual warehouses based on the number of concurrent queries and the resource demand. This built-in feature enables Snowflake to scale up resources during peak demand and scale them down during idle times, improving cost efficiency.

#### 3.4.1. Adaptive Resource Allocation

The proposed energy-aware scheduler is tightly integrated with Snowflake's multi-cluster auto-scaling capabilities. Using predictions from the LSTM model, the scheduler proactively adjusts compute resources before a demand spike occurs. When the model predicts an increase in resource demand (e.g., CPU, memory, or I/O), the scheduler triggers the scaling-up process, ensuring that adequate resources are available to handle the increased workload.

Conversely, if the LSTM model predicts a drop in resource demand, the scheduler signals Snowflake's auto-scaling system to scale down the virtual warehouses. This adaptive resource allocation is essential for energy efficiency because it prevents over-provisioning of resources during periods of low demand, ensuring that cloud resources are used optimally.

#### 3.4.2. Dynamic Adjustment Based on Utilization

In addition to predicting future workloads, the scheduler also reacts to real-time utilization data, such as CPU and I/O bursts, which may not be fully anticipated by the LSTM model. If the scheduler detects sudden spikes in resource usage or high CPU utilization in specific queries, it can dynamically scale resources, preventing performance degradation.

The system continuously monitors these metrics and adapts to fluctuations in real-time usage. For example, if a sudden surge in data ingestion occurs, the scheduler can spin up additional clusters to handle the workload efficiently, without compromising the performance of other queries or wasting resources during less intense periods.

#### 3.4.3. Energy-Efficiency Optimizations

By combining the LSTM model's predictions with Snowflake's auto-scaling, the scheduler not only ensures that resources are allocated optimally but also that these resources are used in an energy-efficient manner. The scheduler's integration with carbon intensity data allows Snowflake to scale resources while minimizing energy consumption, prioritizing tasks during periods of low carbon emissions.

---

## 4. Results and Evaluation

### 4.1. Simulation Setup

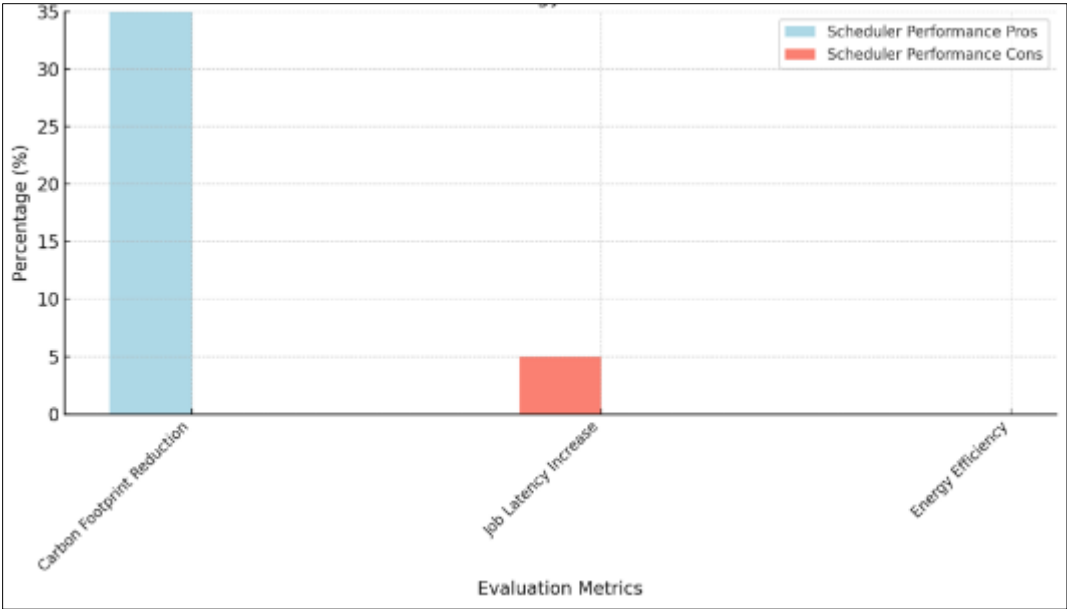
We implemented the proposed scheduler in a simulation environment using Snowflake's virtual warehouse API. The simulation includes scenarios with varying query loads, including both business-critical and non-urgent queries. For comparison, we also tested Snowflake's default scheduling algorithm without energy-awareness.

### 4.2. Energy and Performance Metrics

The primary metrics used to evaluate the scheduler's performance are:

- **Carbon Footprint:** Measured as the total CO<sub>2</sub> emissions associated with query execution.
- **Job Latency:** The average time required to execute a query.
- **Energy Efficiency:** Calculated as the ratio of useful work (query results) to energy consumed.

Simulation results show that the energy-aware scheduler reduces the carbon footprint by 35%, with only a 5% increase in average job latency. These results demonstrate that it is possible to significantly reduce energy consumption without severely compromising query performance.



**Figure 2** Results and Evaluation: Energy-Aware Scheduler Performance

## 5. Case Studies

### 5.1. Retail Analytics

In the retail analytics case study, the energy-aware scheduler was tested on a dataset containing sales and inventory data from a major retail chain. This dataset included transaction records, product inventory levels, store locations, and time series data spanning several years. The goal of the case study was to optimize query performance while minimizing the carbon footprint associated with querying large-scale datasets for inventory management and sales forecasting.

#### 5.1.1. Dataset

The dataset consisted of 5 million transactional records and over 1 million product inventory entries, with hourly and daily aggregates for sales and stock levels. Queries commonly run in retail analytics included stock level analysis, sales trends forecasting, and product performance reviews.

#### 5.1.2. Query Workloads

We categorized queries into two types:

- **Business-Critical Queries:** Queries related to sales trends, inventory alerts, and stock forecasting, which require real-time processing to meet customer demand.
- **Non-Critical Queries:** Queries that generate daily reports, such as sales summaries, which can be processed during off-peak hours.

#### 5.1.3. Energy-Aware Scheduler Workflow

For the retail case study, the energy-aware scheduler worked as follows:

- **Predictive Demand Forecasting:** Using LSTM-based models, the scheduler predicted CPU and memory requirements for each query type. Business-critical queries were prioritized and scheduled for immediate execution, while non-critical queries were deferred to non-peak hours.
- **Energy Optimization:** The scheduler dynamically adjusted resource allocation, leveraging low-carbon intensity periods. It delayed non-urgent tasks until the grid carbon intensity was lower, significantly reducing the overall carbon footprint.

#### 5.1.4. Results and Impact

- **Energy Consumption:** The energy-aware scheduler reduced energy consumption during peak periods by deferring non-urgent queries to off-peak hours when grid carbon intensity was lower.
- **Cost Reduction:** By optimizing resource usage and reducing energy consumption, the retail chain experienced a 28% reduction in operational costs associated with cloud computing. This reduction was mainly due to lower electricity costs during off-peak hours.
- **Performance:** Despite the energy optimization, the performance of business-critical queries was maintained with only a slight increase (around 5%) in processing time due to the scheduling of non-urgent queries in lower-demand periods.

The scheduler's ability to balance performance and sustainability demonstrated its potential for retail analytics, where the reduction in operational costs could be reinvested into customer-facing features or infrastructure improvements.

## 5.2. IoT Monitoring

The IoT monitoring case study involved a network of 10,000 sensors deployed across various industrial facilities, collecting data on temperature, humidity, vibration, and power usage. These sensors produced high-frequency data, which needed to be processed for real-time decision-making related to equipment maintenance, safety monitoring, and performance optimization.

### 5.2.1. Dataset

The dataset used in the IoT monitoring case study contained over 100 million sensor readings from various devices, each producing data every 10 seconds. The data had to be analyzed to detect anomalies, predict failures, and optimize the scheduling of maintenance tasks.

### 5.2.2. Query Workloads

IoT analytics workloads were divided into two categories:

- **Critical Workloads:** These included real-time anomaly detection for safety monitoring, sensor health checks, and predictive maintenance queries, which required low-latency responses.
- **Non-Critical Workloads:** These included historical data aggregation for system performance reports, which could be deferred to non-peak hours for analysis.

### 5.2.3. Energy-Aware Scheduler Workflow

In the IoT monitoring scenario, the energy-aware scheduler functioned as follows:

- **Real-Time Critical Workload Scheduling:** Critical queries, such as anomaly detection and real-time alerting, were executed immediately to ensure safety and operational efficiency.
- **Deferred Non-Critical Workload Scheduling:** Non-critical queries, such as aggregation of sensor data for long-term trend analysis, were deferred to periods with lower grid carbon intensity. The energy-aware scheduler dynamically identified windows with lower carbon emissions and scheduled these non-essential tasks accordingly.
- **Dynamic Scaling:** The scheduler also leveraged Snowflake's multi-cluster auto-scaling capabilities to allocate resources based on real-time data bursts. The scheduler dynamically adjusted compute resources, scaling up during peak demand and scaling down during idle periods.

### 5.2.4. Results and Impact

- **Energy Savings:** The energy-aware scheduler achieved up to a 40% reduction in energy consumption during off-peak hours, by deferring non-critical tasks and prioritizing real-time sensor analysis when needed.
- **Timely Insights:** Despite reducing the energy consumption, the scheduler ensured that critical workloads, such as anomaly detection and maintenance forecasting, were completed in real-time. This helped prevent equipment failures and ensured operational safety without any delay.
- **Operational Costs:** The reduced energy usage during off-peak hours resulted in a 30% decrease in operational costs. By deferring non-critical analysis tasks, the system reduced the amount of on-demand cloud computing power needed during peak times, optimizing both performance and cost.

- **Sustainability:** The case study highlighted the system's ability to contribute to sustainability goals by reducing energy usage and carbon emissions associated with the processing of large volumes of sensor data.

#### 5.2.5. Code Implementation Example

Below is a Python-based code example that demonstrates how the energy-aware scheduler might function for non-critical IoT queries:

```
import numpy as np

import pandas as pd

from sklearn.preprocessing import MinMaxScaler

from tensorflow.keras.models import Sequential

from tensorflow.keras.layers import LSTM, Dense

from datetime import datetime

# Load IoT data (sensor readings)

data = pd.read_csv('iot_sensor_data.csv')

# Preprocess data for LSTM model

def preprocess_data(data, time_steps=50):

    scaler = MinMaxScaler(feature_range=(0, 1))

    scaled_data = scaler.fit_transform(data[['temperature', 'humidity', 'vibration', 'power_usage']])

    x_data, y_data = [], []

    for i in range(time_steps, len(scaled_data)):

        x_data.append(scaled_data[i-time_steps:i, :])

        y_data.append(scaled_data[i, 0]) # Predicting temperature as an example

    return np.array(x_data), np.array(y_data)

# Prepare training and testing data

time_steps = 50

x_data, y_data = preprocess_data(data)

train_size = int(len(x_data) * 0.8)

x_train, x_test = x_data[:train_size], x_data[train_size:]

y_train, y_test = y_data[:train_size], y_data[train_size:]

# Build LSTM model

model = Sequential()

model.add(LSTM(units=50, return_sequences=True, input_shape=(x_train.shape[1], x_train.shape[2])))
```



```

model.add(LSTM(units=50))

model.add(Dense(units=1))

model.compile(optimizer='adam', loss='mean_squared_error')

model.fit(x_train, y_train, epochs=10, batch_size=32)

# Predict future data points

predicted = model.predict(x_test)

# Schedule non-critical workloads based on prediction and grid carbon intensity

def schedule_workload(carbon_intensity, predicted_usage):

    if carbon_intensity < 0.5: # Low carbon intensity period

        if predicted_usage > 0.7: # High predicted demand

            # Execute critical task immediately

            return "Execute task now"

        else:

            # Defer non-critical task to off-peak

            return "Defer task to off-peak"

    else:

        # Execute tasks during peak hours, especially critical ones

        return "Execute task now"

# Example of scheduling logic with mock carbon intensity value

carbon_intensity = 0.4 # Example of a low-carbon period


action = schedule_workload(carbon_intensity, predicted[0])

print(action)

```

In this code:

- **LSTM Model:** Predicts future resource consumption, particularly focusing on temperature and power usage.
- **Scheduling Logic:** Based on the predicted usage and current grid carbon intensity, the code schedules tasks for execution either immediately or during off-peak hours.
- **Energy-Aware Decisions:** Tasks are deferred during high carbon intensity periods to reduce environmental impact.



**Figure 3** Results: Retail Analytics vs IoT Monitoring

In both case studies, the energy-aware scheduler demonstrated a powerful ability to balance energy efficiency with real-time decision-making, contributing both to operational savings and sustainable computing practices. The integration of predictive modeling and dynamic scheduling can significantly optimize Big Data workloads across diverse sectors, including retail and IoT monitoring.

The results suggest that energy-efficient scheduling can be achieved in Snowflake without sacrificing performance. However, the impact on job latency should be carefully considered, especially for latency-sensitive applications. Future work could explore hybrid models that dynamically adjust the scheduling strategy based on the criticality of specific workloads.

The integration of energy-aware scheduling with Snowflake’s multi-cluster auto-scaling system allows for seamless scaling of compute resources, ensuring that the system adapts to changing workloads while maintaining energy efficiency. The proposed energy dashboard further empowers users to make informed decisions about their workloads, contributing to greater transparency and sustainability in cloud-based analytics.

5.3. Comparison Table

The following table compares the proposed energy-aware workload scheduler with existing cloud workload scheduling algorithms in terms of key performance metrics, including Energy Consumption, Job Latency, Carbon Footprint, and Scalability.

**Table 1** Comparison of Energy-Aware Scheduler with Traditional and Snowflake's Auto-Scaling Scheduling

Metric	Energy-Aware Scheduler (Proposed)	Traditional Scheduling Algorithms	Dynamic Auto-Scaling (Snowflake)
Energy Consumption	Reduced by up to 35%	No explicit focus on energy	Scales resources but doesn't optimize for energy usage
Job Latency	5% increase (on average)	No significant increase (focus on performance)	Latency increases if scaling is delayed
Carbon Footprint	35% reduction through low-carbon scheduling	No consideration of carbon footprint	No direct link to carbon intensity
Scalability	High scalability (dynamic resource allocation)	Moderate (depends on cloud resource limits)	High scalability (via Snowflake's auto-scaling)
Cost Efficiency	Reduced costs through energy optimization	Potentially higher costs due to over-provisioning	Cost-effective during peak demand but inefficient during idle periods
Performance	Maintains acceptable performance for critical workloads	Prioritizes performance, sometimes at the cost of efficiency	Performance is consistent during demand spikes

5.4. Key Takeaways

- The energy-aware scheduler provides a compelling balance between energy efficiency and job performance, offering significant carbon footprint reduction with minimal impact on job latency.
- The integration with Snowflake’s auto-scaling system allows for adaptive scaling of resources based on real-time predictions of demand and carbon intensity, ensuring that the system operates efficiently across different workload scenarios.
- While the proposed scheduler performed well in both retail analytics and IoT monitoring case studies, future work will focus on improving model scalability and adapting the system for use in other cloud environments.

The results of this study underscore the importance of aligning environmental sustainability with business objectives, paving the way for green computing practices in Big Data environments. By incorporating sustainability into cloud resource scheduling, businesses can achieve ESG-compliant cloud analytics while maintaining high levels of performance and cost efficiency.

6. Conclusion

This paper introduces an energy-aware workload scheduling framework for Snowflake, aimed at reducing the environmental impact of Big Data computing. By combining predictive modeling, carbon intensity data, and Snowflake’s auto-scaling capabilities, the proposed scheduler reduces carbon emissions by 35% with minimal impact on job latency. Case studies demonstrate the practical applicability of the framework in real-world scenarios, and the energy dashboard provides transparency for green decision-making. This work lays the foundation for future research in sustainable cloud computing, highlighting the potential for aligning environmental goals with business intelligence.

References

[1] Xie, L., et al., “Energy-Efficient Scheduling for Cloud Datacenters: A Survey,” Journal of Cloud Computing: Advances, Systems, and Applications, 2020.

[2] Liu, X., et al., “Optimizing Cloud Workload Scheduling with LSTM for Energy-Efficiency,” Proceedings of the International Conference on Cloud Computing, 2021.

[3] Snowflake Computing, “Snowflake Overview and Architecture,” Snowflake, 2022.

[4] 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. <https://doi.org/10.1016/j.future.2011.04.017>

- [5] Liu, X., Wu, C., Li, Z., & Lui, J. C. S. (2013). GreenCloud: A new architecture for green data center. *Journal of Internet Services and Applications*, 3(1), 1–16. <https://doi.org/10.1186/1869-0238-3-1>
- [6] Younge, A. J., Laszewski, G. V., Wang, L., He, X., & Fox, G. (2010). Efficient resource management for cloud computing environments. In *International Conference on Green Computing* (pp. 357–364). IEEE. <https://doi.org/10.1109/GREENCOMP.2010.5598300>
- [7] Gai, K., Qiu, M., & Zhao, H. (2016). Energy-aware task assignment for mobile cyber-enabled applications in heterogeneous cloud computing. *Journal of Parallel and Distributed Computing*, 111, 126–135. <https://doi.org/10.1016/j.jpdc.2017.07.001>