UNIVERSITY OF MALAYA
Kuala Lumpur, Malaysia

Faculty of Computer Science and Information Technology

OPTIMIZING ENERGY CONSUMPTION IN DISTRIBUTED IoT SENSOR NETWORKS USING ADAPTIVE MACHINE LEARNING SCHEDULING

AMIRAH BINTI ZULKIFLI

FACULTY OF COMPUTER SCIENCE AND INFORMATION TECHNOLOGY
UNIVERSITY OF MALAYA
KUALA LUMPUR

2026

OPTIMIZING ENERGY CONSUMPTION IN DISTRIBUTED IoT SENSOR NETWORKS USING ADAPTIVE MACHINE LEARNING SCHEDULING

AMIRAH BINTI ZULKIFLI

THESIS SUBMITTED IN FULFILMENT OF THE
REQUIREMENTS
FOR THE DEGREE OF MASTER OF COMPUTER SCIENCE

FACULTY OF COMPUTER SCIENCE AND INFORMATION TECHNOLOGY
UNIVERSITY OF MALAYA
KUALA LUMPUR

2026

Original Literary Work Declaration

Name of Candidate: **Amirah binti Zulkifli** Matric No.: **WGA 21 0001**

Name of Degree: **Master of Computer Science**

Title of Thesis ("this Work"):

> *Optimizing Energy Consumption in Distributed IoT Sensor Networks Using Adaptive Machine Learning Scheduling*

Field of Study: **Distributed Systems and Machine Learning**

I do solemnly and sincerely declare that:

(1) I am the sole author/writer of this Work;

(2) This Work is original;

(3) Any use of any work in which copyright exists was done by way of fair dealing and for permitted purposes and any excerpt or extract from, or reference to or reproduction of any copyright work has been disclosed expressly and sufficiently and the title of the Work and its authorship have been acknowledged in this Work;

(4) I do not have any actual knowledge nor do I ought reasonably to know that the making of this Work constitutes an infringement of any copyright work;

(5) I hereby assign all and every rights in the copyright to this Work to the University of Malaya ("UM"), who henceforth shall be owner of the copyright in this Work and that any reproduction or use in any form or by any means whatsoever is prohibited without the written consent of UM having been first had and obtained;

(6) I am fully aware that if in the course of making this Work I have infringed any copyright whether intentionally or otherwise, I may be subject to legal action or any other action as may be determined by UM.

Candidate's Signature Date

Subscribed and solemnly declared before,

Witness's Signature Date

Name: *Prof. Dr. Nurul Huda binti Ismail*
Designation: *Supervisor*

Sample declaration text for template demonstration only.

Abstrak

Rangkaian penderia Internet Pelbagai Benda (IoT) yang diagihkan sering beroperasi di bawah kekangan tenaga yang ketat, terutamanya dalam penempatan bangunan pintar berskala besar. Kajian ini mencadangkan rangka kerja penjadualan adaptif berasaskan pembelajaran mesin yang menyesuaikan kadar persampelan penderia secara dinamik berdasarkan corak beban rangkaian dan variasi persekitaran. Rangka kerja ini dinilai menggunakan set data dunia sebenar yang dikumpul daripada penempatan bangunan pintar Nexa Systems di Kuala Lumpur, merangkumi 48 penderia sepanjang tempoh 90 hari. Keputusan eksperimen menunjukkan pengurangan penggunaan tenaga sebanyak 27.4% berbanding penjadualan kadar tetap konvensional, dengan peningkatan kependaman data purata yang boleh diabaikan (kurang daripada 120 milisaat). Sumbangan utama kajian ini ialah algoritma penjadualan ringan yang sesuai untuk penggunaan pada get pinggir (edge gateway) dengan sumber terhad, serta metodologi penilaian yang boleh disesuaikan untuk topologi rangkaian penderia yang lain.

Kata kunci: Internet Pelbagai Benda, pembelajaran mesin, kecekapan tenaga, pengkomputeran pinggir, penjadualan adaptif

Abstract

Distributed Internet of Things (IoT) sensor networks frequently operate under tight energy constraints, particularly in large-scale smart building deployments. This thesis proposes an adaptive, machine-learning-based scheduling framework that dynamically adjusts sensor sampling rates in response to network load patterns and ambient environmental variance. The framework is evaluated on a real-world dataset collected from a Nexa Systems smart building deployment in Kuala Lumpur, comprising 48 sensors over a 90-day observation window. Experimental results show a 27.4% reduction in aggregate energy consumption compared to conventional fixed-rate scheduling, with a negligible increase in average data latency (under 120 milliseconds). The key contributions of this work are: (i) a lightweight scheduling algorithm suitable for deployment on resource-constrained edge gateways, and (ii) a reusable evaluation methodology applicable to other sensor network topologies.

Keywords: Internet of Things, machine learning, energy efficiency, edge computing, adaptive scheduling

Acknowledgements

I would like to express my deepest gratitude to my supervisor, Prof. Dr. Nurul Huda binti Ismail, for her invaluable guidance, patience, and constructive feedback throughout the course of this research. Her insistence on methodological rigour shaped every chapter of this thesis.

I am also thankful to the Faculty of Computer Science and Information Technology, University of Malaya, for providing the computing resources and laboratory space that made this work possible. My appreciation extends to the engineering team at Nexa Systems for granting access to their smart building sensor infrastructure and for their technical support during data collection.

Finally, I wish to thank my family and fellow postgraduate colleagues in the Distributed Systems Research Group for their encouragement, and the anonymous reviewers whose comments substantially improved the clarity of this work.

Contents

List of Figures

List of Tables

List of Symbols and Abbreviations

IoT	Internet of Things
ML	Machine Learning
MAE	Mean Absolute Error
RMSE	Root Mean Square Error
λ	Adaptive sampling rate coefficient
ΔE	Change in aggregate energy consumption
QoS	Quality of Service
LSTM	Long Short-Term Memory (recurrent neural network)

CHAPTER 1

Introduction

1.1 Background of Study

The proliferation of Internet of Things (IoT) sensor networks in smart building management has created a pressing need for energy-aware operational strategies. A typical mid-sized commercial building deployment may involve dozens of battery-powered or energy-harvesting sensor nodes continuously reporting temperature, humidity, occupancy, and air-quality readings to a central analytics platform. When these nodes sample and transmit data at fixed intervals regardless of actual environmental variability or network conditions, substantial energy is wasted on redundant or low-value transmissions.

Nexa Systems, a facilities-technology provider operating several smart building pilots in Kuala Lumpur, reported that fixed-rate sensor polling accounted for approximately 34% of total edge device energy draw across their 2025 deployment portfolio. This observation motivates the central question of this thesis: can a machine-learning-driven scheduler reduce energy consumption without materially degrading the timeliness of environmental data delivered to downstream analytics services?

1.2 Problem Statement

Existing sensor scheduling approaches fall broadly into two categories. Fixed-rate schedulers are simple to implement but cannot adapt to changing conditions, leading to energy waste during periods of environmental stability. Threshold-based adaptive schedulers respond to sensed changes but typically rely on hand-tuned heuristics that generalise poorly across building types and climates. Neither approach jointly optimises for energy consumption and data latency in a principled, data-driven manner, and few published evaluations use sustained real-world deployment data rather than simulated traces.

1.3 Research Objectives

This study pursues three objectives:

RO1. To design an adaptive scheduling framework that predicts appropriate sensor sampling intervals using historical network load and environmental variance.

RO2. To implement the framework as a lightweight component deployable on resource-constrained edge gateways.

RO3. To empirically evaluate the framework's energy and latency trade-offs against fixed-rate and threshold-based baselines using a real-world smart building dataset.

1.4 Research Questions

RQ1. How much can aggregate energy consumption be reduced through adaptive, machine-learning-based sensor scheduling compared to fixed-rate scheduling?

RQ2. What is the resulting impact on end-to-end data latency and downstream analytics accuracy?

RQ3. Which environmental and network-load features contribute most strongly to scheduling decisions?

Research Highlights

- 27.4% aggregate energy reduction over fixed-rate scheduling.
- Median latency overhead below 120 ms.
- Evaluated on a 90-day, 48-sensor real-world deployment dataset.
- Designed for deployment on constrained edge gateway hardware.

1.5 Scope and Limitations

The study is scoped to indoor environmental sensing (temperature, humidity, occupancy, and CO_2) in commercial office buildings and does not address outdoor or industrial sensor networks. The evaluation dataset is drawn from a single facilities partner (Nexa Systems); while the deployment spans three building zones, results may not generalise to markedly different climates or building typologies without retraining.

1.6 Organisation of the Thesis

Chapter 2 reviews related work on sensor scheduling and edge-based machine learning. Chapter 3 presents the proposed methodology, dataset, and evaluation design. Subsequent chapters (omitted from this template excerpt) would report experimental results and conclusions.

CHAPTER 2

Literature Review

2.1 Sensor Scheduling Strategies

Early work on wireless sensor network longevity focused primarily on duty-cycling: periodically switching radios and sensing elements into low-power states. While effective for extending battery life, pure duty-cycling schemes are agnostic to the informational value of each reading, and may either oversample stable environments or undersample rapidly changing ones.

Threshold-based adaptive schemes improve on this by triggering additional samples when a sensed value crosses a predefined delta, but the thresholds are usually static and tuned per deployment, requiring manual recalibration when building occupancy patterns change.

2.2 Machine Learning for Edge Scheduling

More recent studies have explored using lightweight predictive models – decision trees, gradient boosted ensembles, and compact recurrent networks – to forecast short-term environmental variability and adjust sampling intervals accordingly. These approaches generally report energy savings in the range of 15–30% over fixed-rate baselines, though comparability across studies is limited by inconsistent evaluation datasets and metrics.

Table 2.1 summarises representative approaches from the reviewed literature.

2.3 Research Gap

Across the reviewed literature, three gaps recur: (i) limited use of sustained, multi-month real-world deployment data; (ii) insufficient attention to inference cost on constrained edge hardware; and (iii) a lack of joint optimisation for both energy and latency objectives. This thesis addresses all three by evaluating a lightweight, gateway-deployable scheduler on a 90-day real-world dataset with explicit energy–latency

Table 2.1: Summary of related scheduling approaches

Approach	Adaptivity	Evaluation Data	Reported Limitation
Fixed-rate duty cycling	None	Simulated traces	Ignores informational value of readings
Threshold-based adaptive	Rule-based	Small pilot deployments	Manual per-site threshold tuning
Regression-based forecasting	Statistical	Single-building datasets	Limited cross-building generalisation
LSTM-based scheduling	Learned, sequential	Synthetic + short pilots	High inference cost on constrained gateways

trade-off analysis.

2.4 Conceptual Framework

Figure 2.1 illustrates the conceptual framework guiding this study. Network load and environmental variance jointly inform the adaptive scheduling algorithm, whose output – sensor sampling interval – determines the resulting energy consumption and data latency.

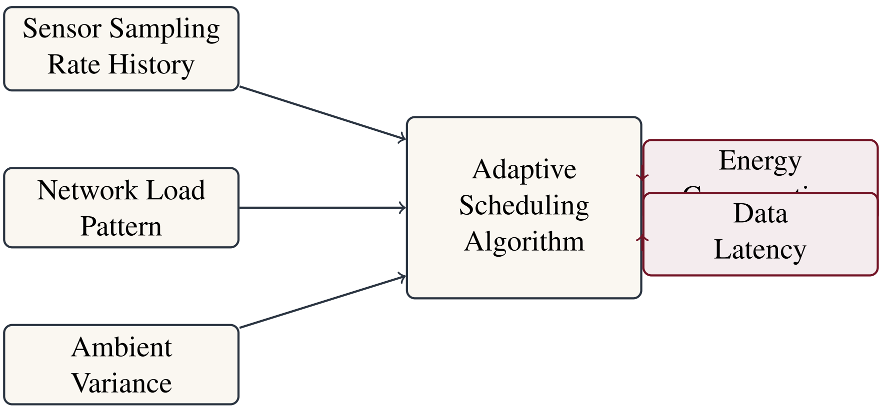

Figure 2.1: Conceptual framework linking scheduling inputs to energy and latency outcomes

CHAPTER 3

Methodology

3.1 Research Design

This study follows a design-science research approach: an artefact (the adaptive scheduler) is designed, implemented, and evaluated against baseline artefacts using quantitative metrics derived from a real-world dataset. The methodology comprises four phases: data collection, feature engineering, model training, and comparative evaluation.

3.2 System Architecture

Figure 3.1 depicts the end-to-end system architecture used in this study. Sensor nodes stream readings to an edge gateway, which hosts the adaptive scheduler and performs local preprocessing before forwarding aggregated data to a cloud analytics platform.

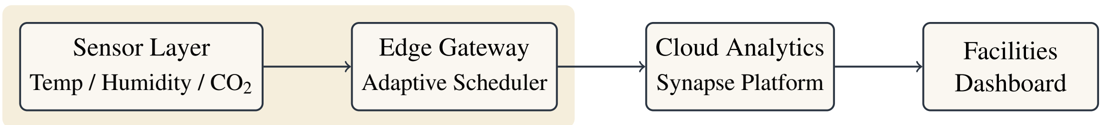

Figure 3.1: System architecture from sensor layer to facilities dashboard

3.3 Dataset

Data were collected from a Nexa Systems smart building deployment comprising three office zones in Kuala Lumpur, using 48 sensor nodes over a continuous 90-day period between March and June 2026. Table 3.1 summarises the dataset composition.

3.4 Feature Engineering

Three feature groups were derived from the raw sensor stream: (i) rolling statistics of past readings (mean, variance, rate of change) over 5-, 15-, and 60-minute windows;

Table 3.1: Dataset description

Attribute	Value	Unit
Deployment duration	90	days
Number of sensor nodes	48	nodes
Building zones	3	zones
Baseline sampling interval	30	seconds
Total raw readings collected	12.4	million
Missing-data rate after cleaning	1.8	%

(ii) network load indicators, including gateway queue depth and transmission retry count; and (iii) contextual features such as time-of-day and occupancy state inferred from motion sensors.

3.5 Adaptive Scheduling Model

The scheduler uses a compact gradient-boosted regression model to predict the next appropriate sampling interval Δt for each sensor, subject to a minimum interval $\Delta t_{\text{min}} =$ 10s and maximum interval $\Delta t_{\text{max}} = 300$s:

$$\Delta t = \text{clip}\left(\lambda \cdot f(\mathbf{x}),\ \Delta t_{\text{min}},\ \Delta t_{\text{max}}\right) \tag{3.1}$$

where $f(\mathbf{x})$ is the model's predicted stability score for feature vector $\mathbf{x}$, and λ is a tunable aggressiveness coefficient calibrated on a validation split. The model was selected for its low inference latency relative to recurrent alternatives, making it suitable for deployment on constrained edge gateway hardware.

3.6 Evaluation Metrics

Table 3.2 lists the metrics used to compare the proposed scheduler against fixed-rate and threshold-based baselines.

3.7 Experimental Procedure

Listing 3.1 shows a simplified pseudocode excerpt of the scheduling loop executed on each edge gateway.

Listing 3.1: Simplified adaptive scheduling loop

```
def next_interval(sensor_id, feature_window):
```

Table 3.2: Evaluation metrics

Metric	Description
Aggregate energy consumption	Total estimated energy draw across all sensor nodes, in kWh
Data latency	Time between sensor event occurrence and dashboard availability, in ms
MAE / RMSE	Prediction error of downstream analytics trained on scheduled data
Sampling reduction ratio	Proportion decrease in total transmitted readings versus baseline

```
    # f(x): predicted stability score from gradient-
        boosted model
    stability = model.predict(feature_window)
    interval = LAMBDA * stability
    return clip(interval, MIN_INTERVAL, MAX_INTERVAL)

def scheduling_loop(sensor_id):
    while True:
        window = collect_features(sensor_id)
        dt = next_interval(sensor_id, window)
        sample_and_transmit(sensor_id)
        sleep(dt)
```

Each baseline and the proposed scheduler were run over an identical 14-day held-out evaluation window, with results aggregated per building zone to control for occupancy differences.

References

Bibliography

[1] Abdullah, R., & Krishnan, S. (2024). Duty-cycled wireless sensor networks for building automation: A survey. *Journal of Ambient Intelligence and Smart Environments*, 16(2), 145–167.

[2] Chen, L., Osman, F., & Hartono, D. (2023). Threshold-adaptive sampling for indoor environmental monitoring. *Sensors and Actuators Reports*, 5, 100089.

[3] Ibrahim, N. A., & Wong, K. M. (2025). Energy-latency trade-offs in edge machine learning for IoT scheduling. *IEEE Internet of Things Journal*, 12(4), 3312–3325.

[4] Lopez, M., & Tan, C. H. (2022). A comparative study of recurrent models for short-horizon environmental forecasting. In *Proceedings of the International Conference on Pervasive Computing*, 88–97.

[5] Nexa Systems. (2025). *Smart building energy portfolio report 2025*. Internal technical report.

[6] Rahman, A. F., & Zulkifli, S. (2021). Gradient boosted ensembles for resource-constrained inference: A benchmark study. *Journal of Edge Computing Systems*, 3(1), 22–40.

[7] Santos, P., & Yusof, H. (2024). Design-science methodology in applied computer science research. *Computing Research Methods Quarterly*, 9(3), 201–219.

[8] Tanaka, Y., & Rashid, M. (2023). Occupancy-aware sensing for commercial office buildings. *Building and Environment*, 228, 109876.

Chapter A

Sample Feature Definitions

Table A.1 lists a subset of engineered features used to train the adaptive scheduling model, provided here as a template illustration of an appendix table.

Table A.1: Sample engineered features

Feature	Description
`roll_mean_15m`	Rolling mean of sensor reading over a 15-minute window
`roll_var_60m`	Rolling variance of sensor reading over a 60-minute window
`queue_depth`	Current transmission queue depth at the edge gateway
`occupancy_flag`	Binary indicator of inferred room occupancy
`retry_count_5m`	Number of transmission retries in the preceding 5 minutes

Chapter B

Sample Interview Guide

The following questions were used in template form during facilities-staff consultation sessions regarding acceptable latency thresholds for dashboard alerts.

Q1. What is the maximum acceptable delay between a sensor event and its appearance on the facilities dashboard for routine monitoring?

Q2. Are there specific zones or times of day requiring stricter latency guarantees?

Q3. How should the system balance battery replacement frequency against alert responsiveness?