INTERNATIONAL JOURNAL of RENEWABLE ENERGY RESEARCH
L. Liu, Vol.15, No.4, December, 2025

Balancing Quality of Service and Lifetime in Solar-Powered IoT via Hyper-Adaptive Duty Cycling

Lei Liu*‡

*School of Intelligent Manufacturing, Weifang University of Science and Technology, Weifang City, Shandong Province, 262700, China

‡ Corresponding Author; Lei Liu, jizhi08_1234@126.com

Received: 05.10.2025 Accepted:09.11.2025

Abstract-Energy-harvesting IoT (EH-IoT) networks enable long-term, autonomous operation for applications such as infrastructure surveillance, smart agriculture, and environmental monitoring. However, the variability of harvested energy—particularly from solar sources—makes it difficult to sustain continuous operation and ensure timely data transmission. Existing studies often lack a comprehensive strategy for managing device activity, especially scheduling transmission and sleep periods to conserve energy, prevent data backlog, and adapt to environmental changes. To address these limitations, this study proposes a Hyper-Adaptive Duty Cycling method that dynamically adjusts device transmission and rest periods based on harvested energy, data generation rates, and remaining battery levels. The method prioritizes high Quality of Service (QoS) through adaptive scheduling rules and short, high-power transmission bursts that increase data throughput. Three years of simulations using real solar panel data demonstrate the approach's effectiveness. The proposed method achieved a high data delivery rate of 78.29% and maintained near-empty data buffers, ensuring timely transmission. However, the increased energy expenditure shortened network lifetime to 24.6 days. In contrast, a fixed-schedule method extended network lifetime to 28.1 days, but delivered only 32% of generated data. These findings illustrate the trade-off between maximizing data delivery and prolonging network lifespan. Overall, the Hyper-Adaptive Duty Cycling approach provides a flexible and responsive energy-management strategy for EH-IoT networks, enabling dynamic adjustment to environmental and operational conditions. It offers system designers a tunable framework for balancing QoS and durability according to application-specific requirements.

Keywords: Hyper-adaptive duty cycling, quality of service, solar-powered IoT, energy-harvesting sensor networks, network lifetime optimization.

1. Introduction

EH-IoT sensor networks are becoming increasingly essential for applications such as environmental monitoring, smart agriculture, and infrastructure surveillance [1]. Their capability to operate autonomously through the collection of ambient energy, primarily derived from solar, thermal, or kinetic sources, enables deployment in remote or inaccessible locations without the necessity for frequent battery replacements [2]. Nevertheless, the intrinsic variability and intermittency of these energy sources present considerable challenges in maintaining network operation, delivering data promptly [3], and ensuring QoS [4], [5]. Consequently, the development of effective energy management strategies—particularly duty cycling algorithms—is crucial for optimizing operational longevity while meeting the requirements of various applications [6].

Traditional duty cycling strategies, including fixed wake-sleep intervals, are straightforward to implement [7]. However, encounter inefficiencies in a dynamic environment [8]. They frequently fail to adapt to fluctuations in harvested energy or workload [9], leading to either energy wastage or missed data transmission opportunities [10]. Consequently, this has spurred the development of adaptive, context-aware techniques that dynamically modify sensor activity based on energy availability and system requirements [11]. Across several technical fields, including robotics and signal processing, the idea of employing adaptive control to manage intricate, dynamic systems is widely accepted [12]. Nonetheless, many of these strategies depend on static thresholds, neglecting workload variations [13], or lack resilience when faced with unpredictable energy patterns.

A promising approach involves integrating backlog-aware and energy-aware decisions [14], wherein duty cycles are adapted to both energy status and data transmission requirements. [15]. Such strategies are capable of harnessing excess energy for aggressive data delivery while conserving energy in times of scarcity [16]. Recent studies have further examined predictive models, employing weather forecasts or machine learning techniques [17], to forecast energy availability [18]. Although these approaches demonstrate potential, their evaluation has primarily been confined to simplified scenarios, such as single-node configurations, idealized energy models, or limited application domains. Consequently, uncertainties persist regarding their effectiveness and applicability in more realistic, long-term contexts.

Within the broader domain of energy-harvesting Internet of Things, previous scholarly investigations have examined techniques ranging from conventional adaptive heuristics to advanced machine learning-based schedulers, including reinforcement learning and correlation-based wake-up schemes. Although these methods have shown improvements in specific contexts, they often overlook critical design considerations such as amortizing wake-up costs through burst transmissions, dynamically balancing QoS and device longevity, and validating performance using real-world energy datasets over extended periods. The absence of thorough evaluation in varied and changing conditions highlights a significant gap in the current literature.

A Hyper-Adaptive Duty Cycling algorithm is proposed in this paper to intelligently manage transmission bursts and sleep intervals based on energy harvesting conditions, backlog size, and real-time battery levels. By incorporating critical reserve thresholds, high-energy burst modes, and an urgency metric that balances energy state with backlog pressure, the approach deliberately prioritizes data throughput to examine the resulting trade-offs against long-term sustainability explicitly. The algorithm's performance is thoroughly assessed against a fixed-duty-cycle baseline through simulation using multi-year solar irradiance data, thereby providing a realistic evaluation of its functionality under both seasonal and short-term variability.

The main contributions of this work are as follows:

- Design of a hyper-adaptive, context-aware duty cycling strategy that integrates battery state, backlog size, and energy availability into transmission and sleep scheduling.
- Comprehensive performance evaluation using multi-year real-world solar data to capture both predictable seasonal cycles and unpredictable short-term fluctuations.
- Quantitative analysis of QoS–lifetime trade-offs that provide actionable insights for tailoring duty cycling strategies to application-specific requirements.

The remainder of this paper is organized as follows: Section 2 reviews related work in adaptive duty cycling and EH-IoT scheduling. Section 3 outlines the proposed methodology, which includes system modeling, duty cycling algorithms, and the experimental setup. Section 4 presents and analyzes the experimental results, while Section 5 discusses the implications, limitations, and potential areas for improvement. Section 6 concludes the paper and outlines directions for future research.

2. Related Work

This section examines contemporary energy management strategies for EH-IoT, including adaptive duty cycling, predictive forecasting, and reinforcement learning. It identifies a gap: the lack of comprehensive solutions that concurrently address data backlogs and energy states, supported by long-term, real-world data.

The work [10] proposes an energy-aware duty-cycle manager for solar-powered IoT nodes. Their method integrates online weather-based sunshine forecasts with device-level measurements, a solar panel model, and a battery model to dynamically adjust each node's sleep interval. Implemented on an ESP32 platform with a custom solar sensor board within a serverless IoT framework, the algorithm proactively extends sleep duration during cloudy days. Field testing has demonstrated the effectiveness of this approach. For example, increasing sleep time before multi-day energy deficits helped the battery maintain a charge above 20%, preventing complete depletion. These findings suggest that proactively adjusting the duty cycle can maintain the necessary charge levels using a smaller battery. However, there are limitations, such as reliance on accurate weather forecasts and specialized hardware. Unpredictable shading or inaccuracies in forecasting can still cause system outages, and this method is specifically designed for solar energy harvesting. In summary, this research contributes to the development of predictive duty cycling by addressing the challenges posed by multi-day energy variability. Nonetheless, its success depends on forecast accuracy, and it is mainly suited to single-node solar energy systems.

Giordano et al. [19] introduce an adaptive sampling algorithm designed for resource-constrained, battery-powered Internet of Things sensors. Drawing inspiration from the TCP Reno AIMD (additive-increase/multiplicative-decrease) algorithm, their finite-state machine (FSM)-based controller modulates the sensor's sampling rate according to the remaining battery capacity and the harvested solar energy. The authors calibrated a solar-cell model using data from 48 days of measurements and utilized open dataset inputs to simulate energy availability. Validation was performed using a prototype EcoTrack device equipped with GNSS and LTE connectivity. Results from three European cities show that the algorithm maintains a minimum activity level (24 localizations per day) while utilizing energy to handle up to approximately 3,000 samples per day when sufficient energy is available. These results demonstrate that the algorithm can maximize the amount of data it collects while still operating independently without requiring additional assistance. However, this method has some limits. It only works with one device and one sampling task, and it doesn't take into account factors such as network communication, delays, or changes in workload. Additionally, it assumes the battery is always 3000 mAh and utilizes historical solar data, which may not accurately reflect real-world conditions or the interactions between multiple devices. In short, this approach demonstrates that smart sampling can continue to run even when energy is being collected; however, it doesn't address scheduling across a network or handling unexpected changes in workload.

In [20], the authors introduce DRDC, a deep reinforcement learning framework for managing the sleep and wake times of energy-harvesting body sensors. A Deep Q-Network (DQN) is trained outside the sensor to decide when each sensor should be awake or asleep. The model uses the battery level and the rate of data change to make more informed decisions. This helps prevent the loss of important data and unnecessary sensor sleep. The reinforcement learning system operates on a local server with a simple neural network, as the sensors have limited resources. The simulation results show that DRDC reduces the time the sensors are awake by about 28% and the amount of data sent by about 50% compared to earlier methods. This makes the system more efficient and better at using energy. Limitations of the approach include its focus on a specific wearable body-area scenario with an assumed network infrastructure, relying on centralized training and known data-rate dynamics; consequently, its performance in highly unpredictable or multi-hop network conditions remains uncertain. Like other reinforcement learning methodologies, it necessitates models derived from simulation and may not directly generalize to all Internet of Things applications. In summary, DRDC demonstrates that deep reinforcement learning can optimize duty cycling under uncertain energy conditions; however, further research is needed to evaluate its practicality beyond body sensor applications.

Ruiz-Guirola et al. [21] propose an intelligent duty-cycling and wake-up strategy for EH-IoT networks with correlated events. The energy and data usage of each device are modeled as Markov chains, where energy is supplied through a modulated Poisson process. The base station uses a K-nearest neighbors (KNN) method to set each device's duty cycle. It uses both the location and timing of device activity to make these decisions and can activate specific devices when additional data is required. In tests with dense networks, this approach reduces missed detections by up to 11 times and cuts energy use by about 50% compared to random duty cycles. This demonstrates that adjusting duty cycles based on learned patterns conserves energy without compromising sensing accuracy. However, there are some limitations, such as the need to know the spatial patterns and perform KNN calculations at the base station. The methodology assumes predictable event patterns and centralized control, which may not apply to all IoT deployments. Moreover, validation has been confined to simulation environments. In conclusion, this research utilizes machine learning techniques and wake-up radios to optimize duty cycles in energy-harvesting IoT networks; however, its reliance on scenario-specific assumptions limits its general applicability.

Researchers of [22] present a comprehensive review of energy harvesting within wireless sensor networks. They examine available renewable sources and various battery-charging techniques, including wireless charging. The paper underscores that, in principle, if the energy harvested were consistent and plentiful, sensor nodes could operate indefinitely. They acknowledge recent scholarly interest in employing reinforcement learning to adapt node parameters to harvesting patterns, such as adjusting duty cycles or transmission power based on predicted energy availability. The review highlights the advantages of energy harvesting, including reduced dependence on batteries, and outlines the challenges, such as irregular energy supply. However, as a survey, it does not introduce novel algorithms or experimental results. Its focus remains on energy harvesting technologies rather than adaptive scheduling strategies. The review does not evaluate specific duty-cycle schemes nor test protocols; consequently, it leaves unresolved how energy harvesting could be operationalized in real Internet of Things systems. In summary, Tran et al. emphasize the potential of energy harvesting and advocate for intelligent adaptation, yet they do not bridge the gap regarding the development of concrete adaptive algorithms for duty cycling in IoT contexts.

Mushtaq et al. [3] provide a comprehensive review of recent advances in sustainable, energy-harvesting wireless sensor networks. Their survey encompasses energy harvesting methods (solar, thermal, vibration, and RF), energy management strategies, cognitive radio integration, secure communication protocols, and routing protocols. They observe that although many protocols now incorporate energy-awareness, significant challenges remain, including unpredictable ambient energy availability, the need to balance network lifetime with QoS, and security vulnerabilities. The authors highlight emerging trends such as AI-driven management systems and innovative storage technologies as future enablers. The contribution of this paper is extensive: it identifies future research directions and advocates for intelligent energy scheduling. The limitations of this review include its broad scope and generality. It does not evaluate or compare specific algorithms, nor does it examine particular duty-cycling strategies in detail. Consequently, while emphasizing the importance of adaptive energy control, it leaves the question of concrete scheduling solutions unresolved.

Study [23] examines throughput-optimal scheduling in RF energy harvesting links subject to time variability. For a single harvest-then-transmit device, they initially derive the optimal offline schedule, delineating the periods designated for harvesting versus transmitting, to maximize data throughput under conditions of a fluctuating RF source. Notably, their findings reveal that the optimal transmit power is dependent solely on the RF harvest rate, independent of the transmission duration. Subsequently, they develop an efficient online heuristic that attains over 90% of the offline optimal performance in simulated environments. This study demonstrates that meticulous scheduling can substantially enhance throughput within variable energy contexts. Nevertheless, its scope remains limited: it considers only a single node with a known energy harvesting process. The approach does not extend to multi-node networks or unpredictable energy losses, nor does it incorporate duty-cycle management or sensing tasks. In conclusion, Shan et al. provide a robust analytical benchmark for energy harvesting scheduling; however, applying this single-link policy to comprehensive IoT networks, particularly those with latency or QoS constraints, remains an unresolved challenge.

Table 1. Summary of methodology, findings, and limitations in recent EH-IoT duty-cycle studies

Ref	Methodology	Findings	Limitations
Gerndt et al. [10]	Sunshine forecasting and device energy model to adapt duty cycle (serverless IoT).	Guarantees continuous operation; adapts sleep to predicted cloud conditions; features a smaller battery.	Depends on accurate weather forecasts and special hardware; limited to solar EH.
Giordano et al. [19]	TCP-Reno–inspired FSM to adapt sensor sampling rate using solar-cell model.	Maintains self-sustainability (≥24 samples/day) and maximizes sampling (up to ~3000).	Single-device focus; assumes fixed battery; ignores network traffic/latency.
Mohammadi et al. [20]	Deep RL (DQN) for body-node duty-cycle; considers data-change rate in reward.	~28% duty-cycle reduction, >50% data overhead saving vs previous RL schemes.	Targets body-area networks; requires centralized training; not tested on real EH variability.
Ruiz-Guirola et al. [21]	Markov model + KNN duty-cycling using spatio-temporal event correlations; wake-ups.	11× lower event misses and 50% less energy use vs. random baseline.	Assumes known event correlations; computation-heavy (KNN at BS); only simulation results.
Tran et al. [22]	Review of EH sources and wireless charging methods for WSNs.	Surveys renewable sources; notes EH can (theoretically) enable indefinite operation; suggests RL for adaptation.	Descriptive survey; no specific duty-cycle algorithms; does not address the adaptive scheduling gap.
Mushtaq et al. [3]	Review of EH-WSN protocols (energy management, routing, CR, security).	Highlights the need for AI-driven energy scheduling; synthesizes the state of the art.	Broad scope; no experiments; lacks detail on specific duty-cycling or scheduling techniques.
Shan et al. [23]	Offline/Online Optimal Harvest-Then-Transmit Scheduling for RF-Powered Links.	Derives a provably optimal schedule; the online heuristic achieves ≈approximately 90% of the maximum throughput.	Single-link model; assumes known energy trace; not applied to network-level duty cycling.

The critical analysis presented in Table 1 indicates that, although recent studies have achieved notable progress, specific gaps persist. For example, as noted by Gerndt et al. [10] and Giordano et al. [19], self-sustainability was attained through forecasting or adaptive sampling; however, their approaches are limited to particular scenarios, such as solar power or single nodes. Mohammadi et al. [20] and Ruiz-Guirola et al. [21] employ DL techniques for duty cycle management, thereby achieving significant improvements; however, these methods rely on centralized training or assumptions of correlation. Moreover, reviews conducted by Tran [22] and Mushtaq [3] emphasize the significance of adaptive energy management, but lack the implementation of concrete scheduling algorithms. Shan et al. [23] propose optimal policies for a single radio-frequency-powered link, yet they do not extend their methodologies to comprehensive IoT networks. None of these studies sufficiently addresses dynamic, multi-node IoT deployments that incorporate long-term predictions and real-time constraints, leaving the identified gap—namely, the necessity for a hyper-adaptive, prediction-aware duty-cycling solution in energy-harvesting IoT—unresolved. This scenario underscores the motivation for our current research to develop such an adaptive scheme.

3. Methodology

The primary objective is to design and compare an advanced adaptive duty cycle algorithm with a traditional fixed-cycle baseline. A simulation approach was employed due to its controlled, repeatable, and cost-effective characteristics, facilitating the analysis of complex system behaviors over time. As illustrated in Figure 1, the process comprises five steps: firstly, describing the simulation environment and tools; secondly, conducting Exploratory Data Analysis (EDA) on solar irradiance data to comprehend patterns; thirdly, developing a system model based on EDA results that encompasses energy harvesting, sensor node traits, and data workload; fourthly, delineating duty cycling strategies—namely, the baseline Fixed-Cycle and the proposed Hyper-Adaptive; and finally, outlining the experimental setup, including parameters and metrics. This framework ensures the derivation of verifiable, realistic insights into a solar-powered IoT sensor network.

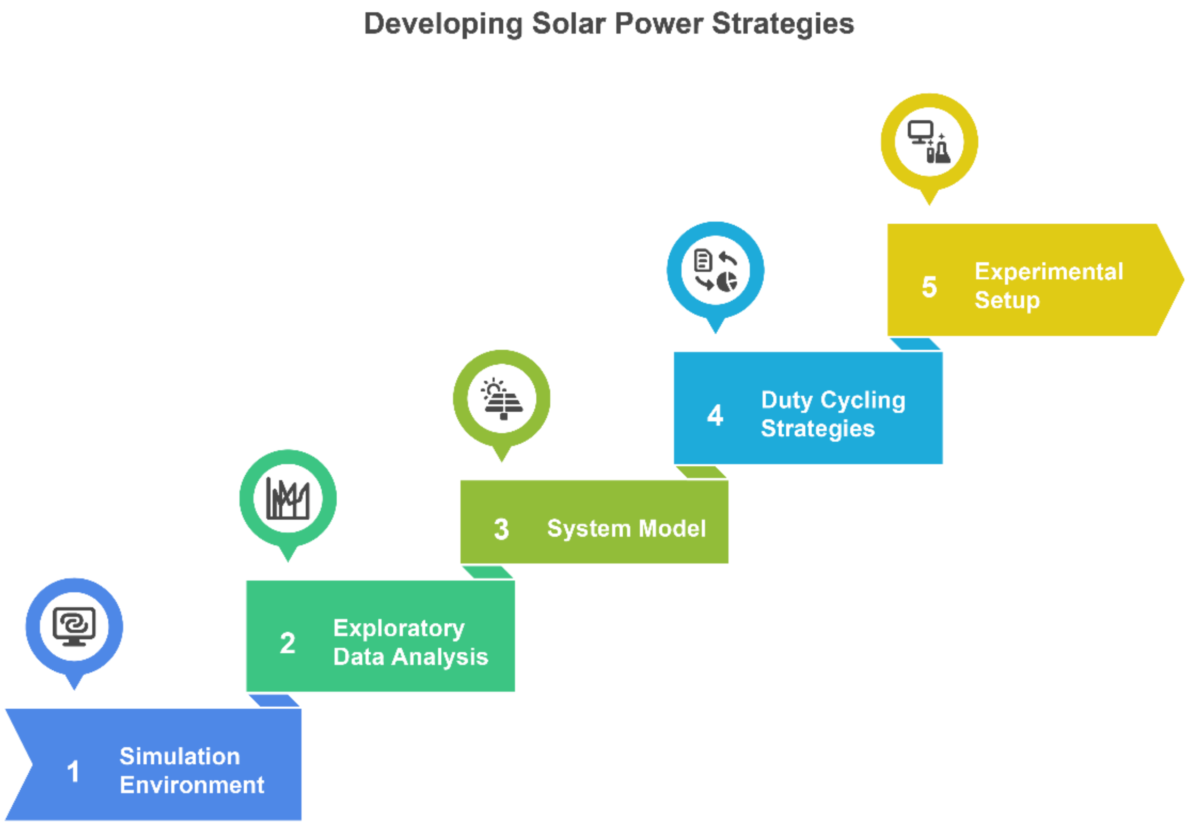

Fig. 1. An overview of the five-step research methodology.

3.1. Simulation Environment and Tools

Developed in Python (version 3.11), the investigation utilized a custom-built discrete-event simulation framework. The selection of Python was driven by its extensive scientific computing ecosystem and robust data analysis libraries.

The foundation of the simulation was established using SimPy (Version 4.0), a process-based discrete-event simulation library. SimPy is particularly well-suited for this research due to its ability to model concurrent processes, such as energy harvesting, workload generation, and sensor state transitions, as independent yet interacting Python generators. This process-based paradigm allows for a natural and intuitive representation of the asynchronous events that characterize an IoT network.

Data manipulation, preprocessing, and analysis were performed using the Pandas library. The three-year solar irradiance dataset was loaded[†], cleaned, and resampled into a time-series DataFrame, which served as the primary input for the energy harvesting model. For visualization, the Matplotlib and Seaborn libraries were employed to generate all plots and figures presented in this paper, enabling a clear visual comparison of the performance of different strategies. The entire simulation and analysis were conducted within a Kaggle Notebook environment, ensuring a reproducible research workflow.

3.2. Exploratory Data Analysis

3.2.1. Temporal patterns of solar irradiance and power

Understanding the temporal patterns of solar irradiance and the resulting harvested power is fundamental to this analysis, as it reveals the dual nature of the energy source that the IoT sensor network relies upon. The short-term daily volatility is clearly illustrated in Figure 2, which contrasts the smooth, predictable energy profile of clear days with the erratic, spiky patterns of cloudy days over a sample week. Figure 3 confirms that this volatility directly translates into the power harvested by the sensor, establishing a proportional relationship between incoming solar irradiance and available energy.

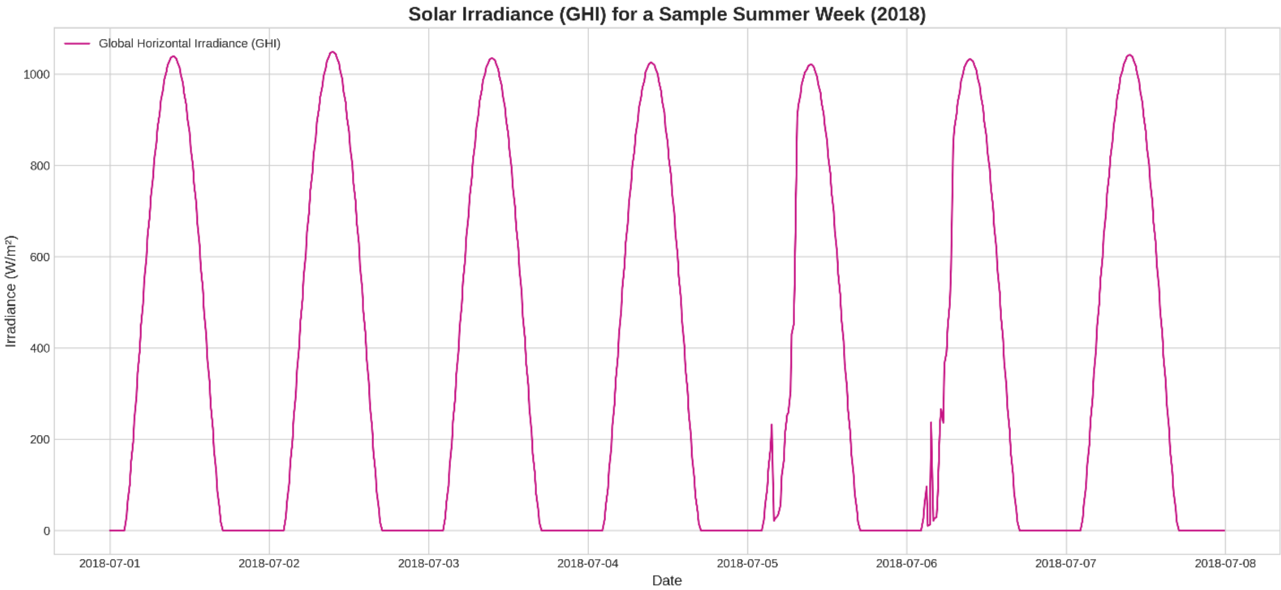

Fig. 2. Solar irradiance (GHI) for a sample summer week.

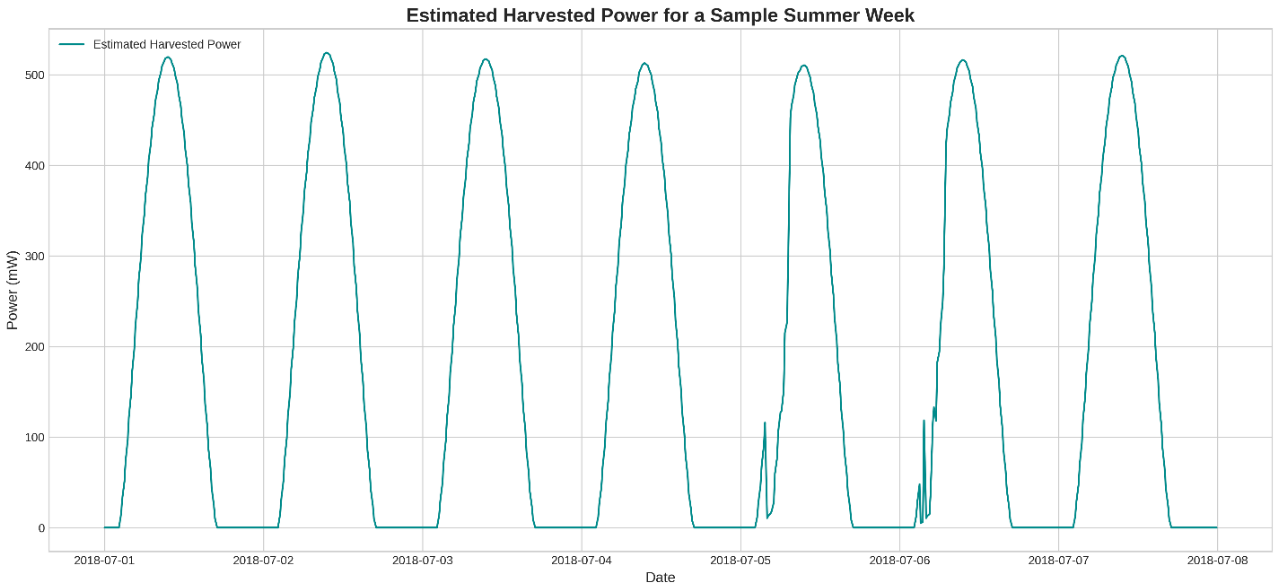

Fig. 3. Estimated harvested power for a sample summer week.

Expanding the scope to a multi-year perspective, Figure 4 illustrates the predictable long-term seasonal cycle. This cycle manifests as a consistent sinusoidal wave that recurs annually, with peak energy availability during the summer and a minimum in winter. Nonetheless, superimposed upon this predictable seasonal trend is a considerable degree of daily volatility, as initially observed in Figure 2. This coexistence of a dependable long-term cycle and short-term intermittency, which appears unpredictable, precisely constitutes the challenge that an adaptive duty cycling strategy must address.

A practical algorithm cannot depend solely on simple averages; it must possess sufficient sophistication to manage both the energy-scarce winter months and abrupt cloudy spells through energy conservation, while also capitalizing on energy-abundant sunny periods. Consequently, this analysis of temporal patterns serves as an essential foundation for designing and justifying a Hyper-Adaptive strategy, as it delineates the dynamic environment within which the system must operate.

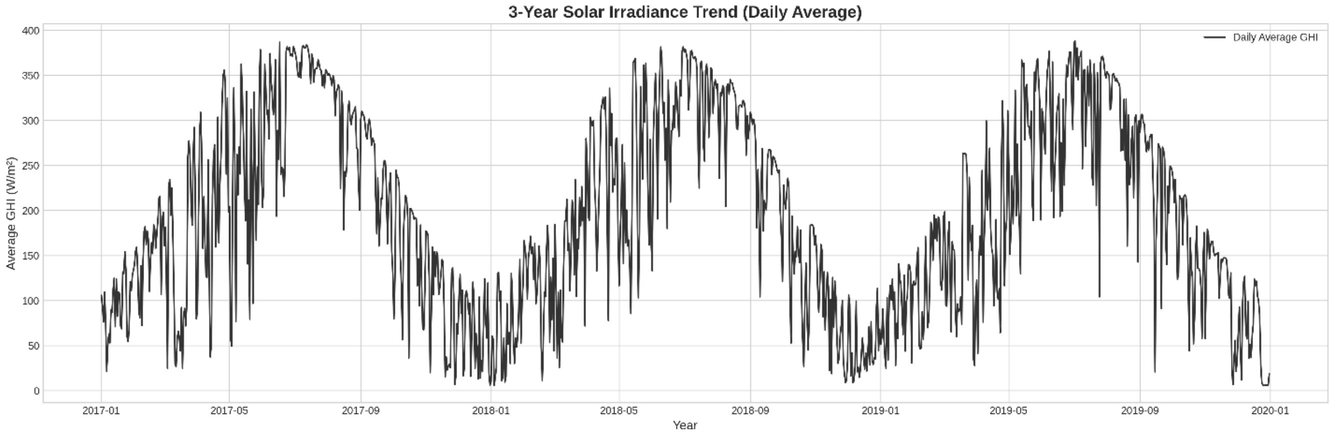

Fig. 4. Three-year trend of daily average solar irradiance (GHI).

3.2.2. Power distribution analysis

Beyond analyzing temporal trends, the power distribution analysis examines the frequency and variability of harvested power levels, offering critical insights into the resource's statistical nature. As shown in the histogram in Figure 5, the distribution of harvested power during daylight hours is strongly right-skewed. This indicates that the most frequent state is one of minimal power generation, corresponding to periods of low light, such as early morning, late evening, or heavy cloud cover. Following this dominant peak at near-zero levels, the distribution settles into a broad, tapering plateau, signifying that while moderate to high power generation is achievable, it occurs with significantly less frequency.

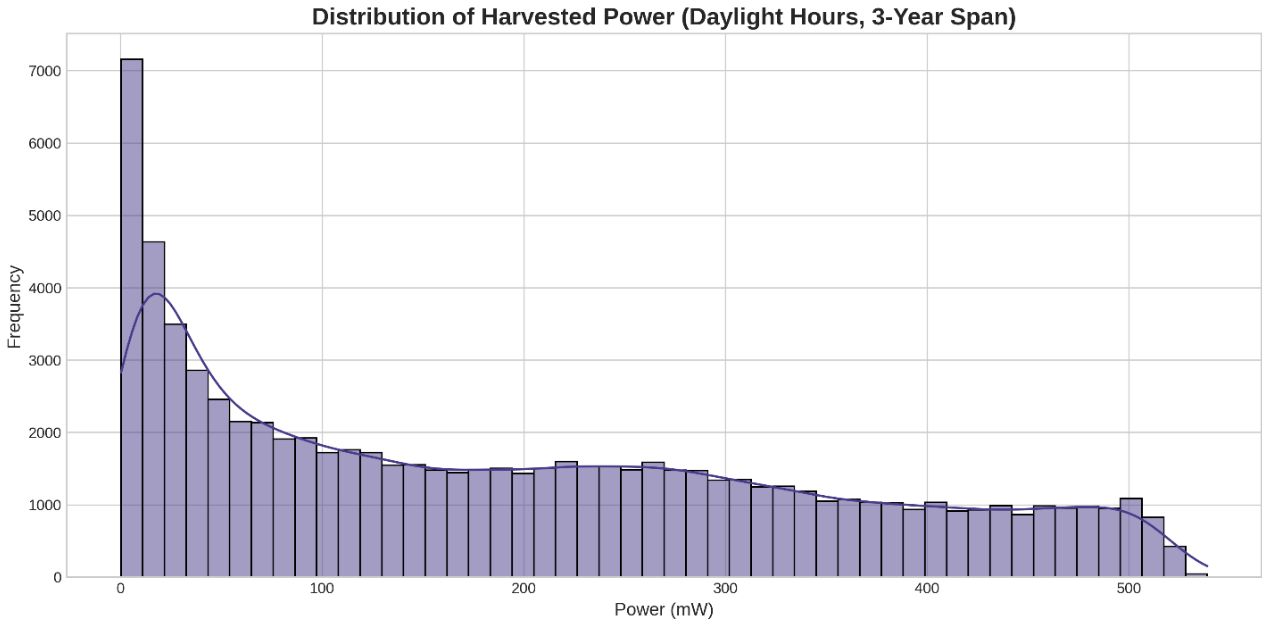

Fig. 5. Frequency distribution of harvested power over three years.

The hourly box plots in Figure 6 further illuminate this by illustrating the typical diurnal generation curve. Power is consistently zero overnight, begins to rise around hours 3-4, peaks between hours 9 and 12, and ceases by hours 16-17. Crucially, this figure highlights that the most significant variance in power output, as represented by the largest box plots and whiskers, occurs during these peak midday hours. This demonstrates that periods of highest potential solar energy are also periods of most significant uncertainty due to fluctuating weather. Taken together, this analysis underscores a fundamental challenge: the system must be designed to survive frequent and prolonged periods of low energy availability while also being capable of responding to the highly variable and unpredictable nature of power generation during peak daylight hours.

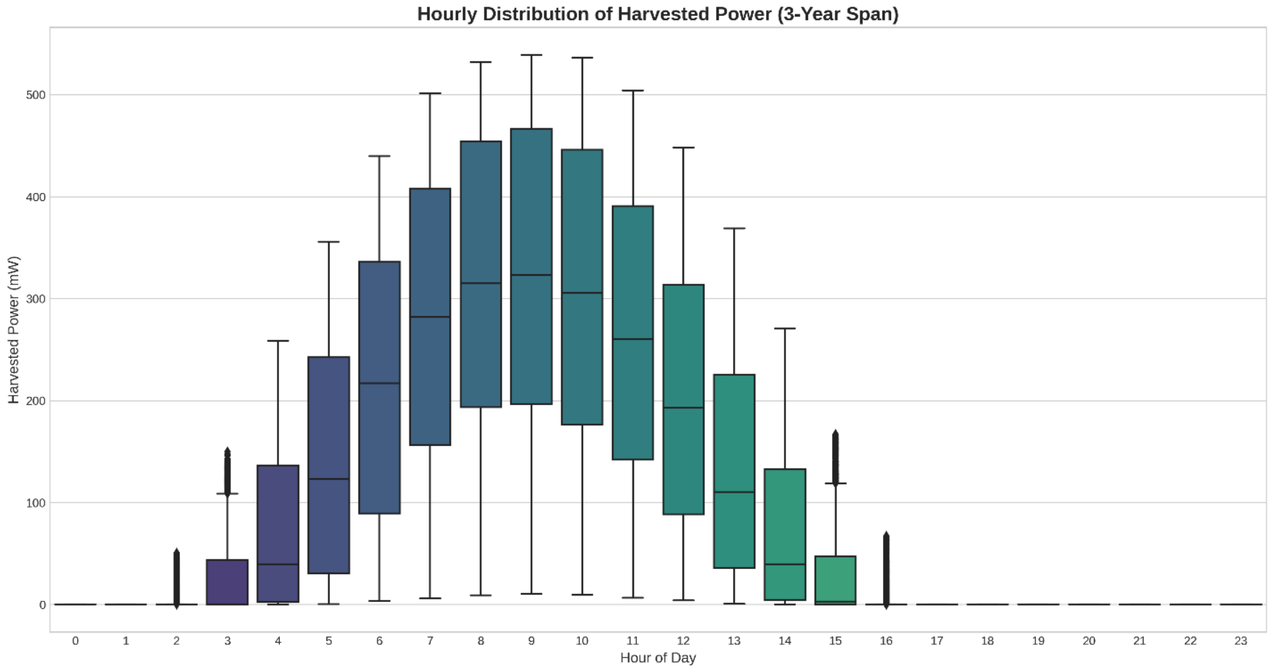

Fig. 6. Hourly distribution of harvested power over three years.

3.3. System Model

To accurately evaluate the duty cycling strategies, a detailed system model was developed, comprising three main components: the energy source, the sensor node, and the data workload.

3.3.1. Energy harvesting model

The energy source constitutes a fundamental element for any investigation involving energy-harvesting systems. To ensure considerable realism, this study employed a publicly accessible dataset sourced from the National Solar Radiation Database (NSRDB), supplied by the National Renewable Energy Laboratory (NREL). The dataset encompasses three complete years (2017-2019) of solar irradiance data for a specific geographic location (39.01°N, 37.26°E). The raw data were provided at a 15-minute temporal resolution. Subsequently, the dataset was preprocessed to establish a continuous time-series index. To balance the precision of the simulation with computational efficiency, the data were resampled to 10-minute intervals through mean aggregation, supplemented by linear interpolation to address any missing values. The primary input for computing the energy harvested by the sensor's solar panel was the Global Horizontal Irradiance (GHI), measured in watts per square meter (W/m²). The calculation of the harvested power (expressed in milliwatts, mW) is delineated in Eq. 1:

$$P_{\text{harvested}} = \text{GHI} \times A_{\text{panel}} \times \eta_{\text{panel}} \times 1000 \qquad (1)$$

where A_{panel} is the area of the solar panel, assumed to be 0.0025 m² (5 cm × 5 cm), and η_{panel} is the panel's energy conversion efficiency, assumed to be 20%, a typical value for small-scale photovoltaic cells. The resulting time series of harvested power served as the energy input for each simulated sensor.

3.3.2. IoT sensor node model

An IoT sensor node is modeled as an entity with limited energy storage and various energy consumption states. Each sensor has a rechargeable battery, represented as a finite server with a maximum capacity of 100,000 mJ. This large capacity allows the sensor to accumulate excess energy during sunny periods, ensuring it can operate throughout the night and on cloudy days. The simulation begins with an initial battery charge of 50,000 mJ, equivalent to 50% of the total capacity. The energy consumption model is event-driven and captures basic activities of a low-power sensor node. Energy costs, based on typical values for low-power microcontrollers and radios, are provided in Table 2.

Table 2. Energy consumption parameters for the simulated IoT sensor node

Parameter	Value	Unit	Description
Wake-up Cost	1.0	mJ	One-time energy cost to transition from sleep to active state.
Packet Transmission Cost	5.0	mJ	Energy cost to transmit a single data packet.
Sleep Power Draw	0.05	mW	Continuous power consumption while in the deep sleep state.

3.3.3. *Workload generation model*

The data generation process, or workload, at each sensor was modeled as a Poisson process. This model is a standard and broadly accepted method for simulating the occurrence of random, independent events, such as sensor readings or event detections. In the simulation, packet generation events were executed by scheduling them with an exponentially distributed inter-arrival time. The average arrival rate (λ) was established at 6 packets per hour, representing a moderate workload suitable for a monitoring application. Each generated packet is timestamped upon creation to facilitate the calculation of data latency.

3.4. *Duty Cycling Strategies*

3.4.1. *Baseline strategy: fixed duty cycle*

The Fixed Duty Cycle strategy exemplifies a straightforward, non-adaptive methodology frequently employed in fundamental IoT applications. Its operational logic is simple and inflexible: the sensor remains in a low-power sleep mode for a predetermined duration of 30 minutes (1800 seconds). Upon the completion of this period, the sensor awakens, incurring an energy expenditure associated with the wake-up process. Subsequently, it assesses its data buffer; if the buffer is not empty and the sensor possesses sufficient energy to transmit a single packet, transmission occurs. Irrespective of the transmission outcome, the sensor promptly re-enters the sleep state for another 1800-second interval. This approach is anticipated to be inefficient, as it may waste energy through unnecessary wake-ups when the buffer is empty or fail to clear its backlog effectively, thereby resulting in increased data latency. It serves as a fundamental baseline for evaluating the advantages of more adaptive strategies.

3.4.2. *Proposed strategy: hyper-adaptive duty cycle*

The proposed Hyper-Adaptive strategy is an intelligent control algorithm designed to maximize data throughput while ensuring sustained long-term energy neutrality in a sensor node. Its operation, conceptually illustrated in Figure 7, follows a sequence of interdependent calculations that jointly determine how many packets to transmit at once and how long the device should subsequently sleep, based on the real-time energy and buffer state of the system.

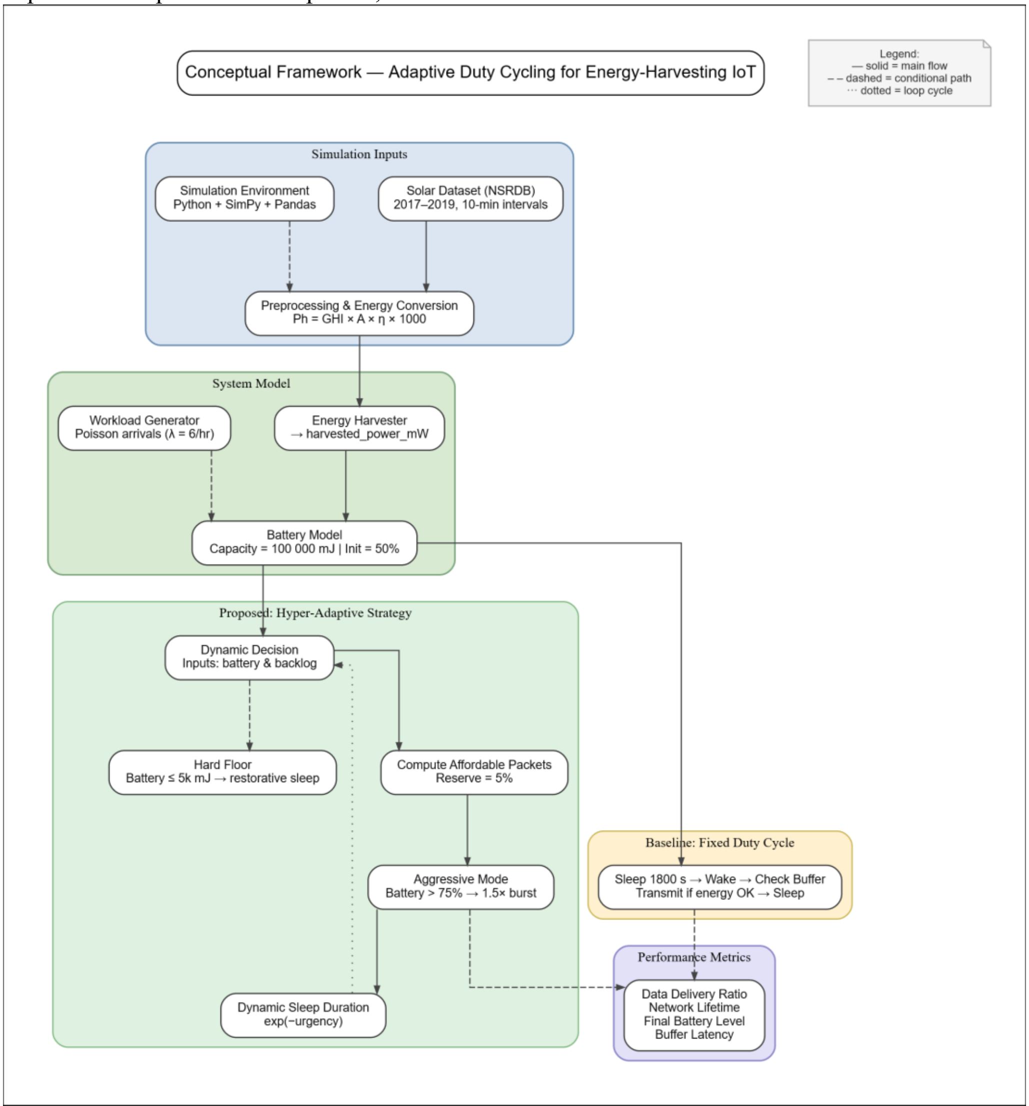

Fig. 7. Conceptual diagram of the hyper-adaptive algorithm's decision logic.

The first stage involves determining the maximum number of packets that can be transmitted in a single burst. This value, (N_{max}), is derived from the available battery energy (E_{batt}) in excess of a reserved safety margin (E_{res}), after accounting for the one-time wake-up cost (E_{wake}). The remaining usable energy is then divided by the per-packet transmission cost (E_{tx}) as provided in Eq.2:

$$N_{max} = \left\lfloor \frac{E_{batt} - E_{res} - E_{wake}}{E_{tx}} \right\rfloor \quad (2)$$

If the battery level surpasses 75% of the total storage capacity, the algorithm increases this burst size by an additional 50% to opportunistically empty the data buffer during periods of energy abundance. This adaptive amplification prevents backlog accumulation and enhances responsiveness when resources are plentiful.

After scheduling the transmission burst, the algorithm evaluates the urgency of the system state to determine the appropriate sleep duration. The total urgency metric (U_{total}) incorporates both battery level urgency and backlog urgency. The battery urgency component (U_{batt}) is modeled as the ratio of the current battery level to the full capacity (E_{cap}), scaled by a tuning factor (α). The backlog urgency ($U_{backlog}$) grows logarithmically with the number of queued packets (N_{buf}), ensuring diminishing influence as buffer size increases. The combined urgency is expressed as Eq.3:

$$U_{total} = \left(\frac{E_{batt}}{E_{cap}}\right)\alpha + \ln\left(1 + N_{buf}\right) \quad (3)$$

Finally, the sleep duration (T_{sleep}) is dynamically selected based on an exponential decay mapping of urgency between a minimum allowable interval (T_{min}) and a maximum interval (T_{max}). A decay constant (k) controls the sensitivity of this response as defined in Eq.4:

$$T_{sleep} = T_{min} + (T_{max} - T_{min}) \cdot e^{-k \cdot U_{total}} \quad (4)$$

This formulation ensures short sleep cycles when urgency is high, either due to low energy or a large backlog, and longer, conservative sleep cycles when urgency is low. However, if the battery falls below a hard safety threshold, the algorithm overrides all calculations and enforces a maximal sleep period to enable energy recovery, thereby maintaining system survivability.

3.5. Experimental Setup and Performance Metrics

For an equitable comparison, the Fixed and Hyper-Adaptive sensors were simulated under identical conditions. Each simulation lasted 30 days, equivalent to 720 hours, allowing for the examination of both the daily fluctuations in energy use and the longer-term trends over time. The strategies were checked using necessary measures: Data Delivery Ratio as a way to check the QoS, Network Lifetime to see how long the system can run before the batteries run out, Final Battery Level to show how sustainable the strategy is, and Data Buffer Length to track how much delay there is in data. Examining these numbers helps to understand the various ways duty cycling strategies impact energy-harvesting networks.

4. Results

4.1. Battery Performance Evaluation

The primary indicator of a strategy's sustainability is its impact on the sensor node's battery level over time. This evaluation assesses the ability of each duty cycling approach to maintain energy neutrality and prolong the operational life of the system.

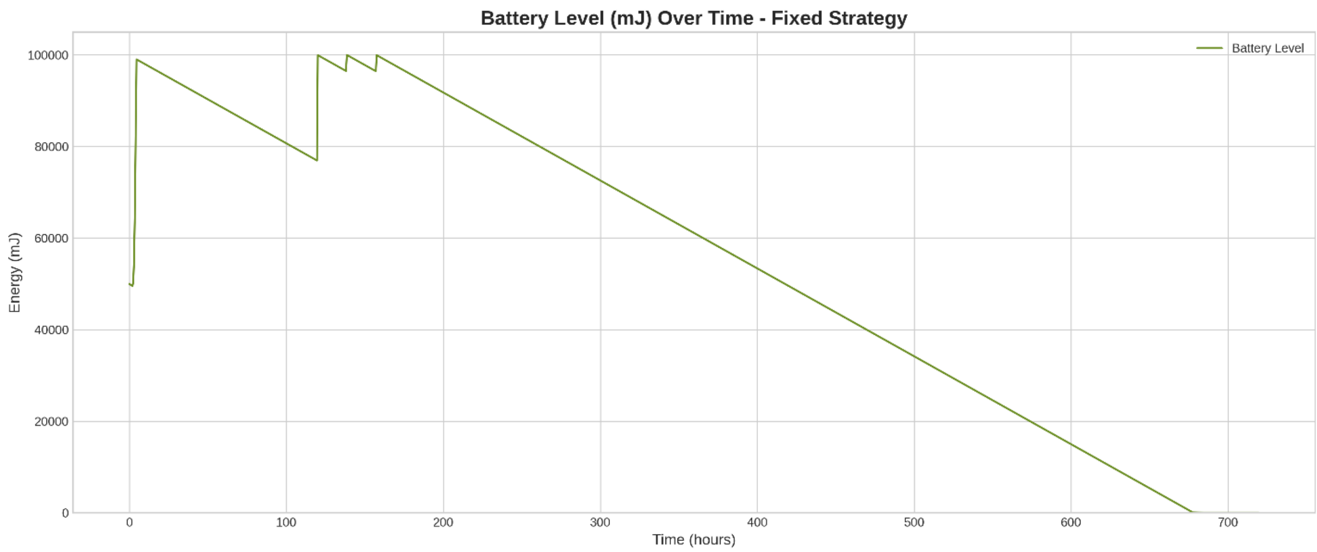

Fig. 8. Battery level over time: Fixed strategy.

Figure 8 depicts the energy level of a battery, measured in millijoules (mJ), over approximately 700 hours of operation under a Fixed Strategy. The x-axis represents time in hours, while the y-axis indicates the stored energy. The simulation commences with an initial charge of 50,000 mJ. Around the 120-hour mark, a significant recharge event restores the battery to its maximum capacity of 100,000 mJ. However, the most significant feature of the plot is the prolonged and steady linear decline that follows. This consistent downward slope signifies a net energy deficit, with the battery depleting at an average rate of approximately 200 mJ per hour. Ultimately, this unsustainable energy expenditure results in the complete depletion of the battery, which reaches zero energy at approximately 675 hours (28.1 days). The overall trajectory confirms that the Fixed Strategy is unviable for long-term operation, as its energy expenditure consistently exceeds energy intake.

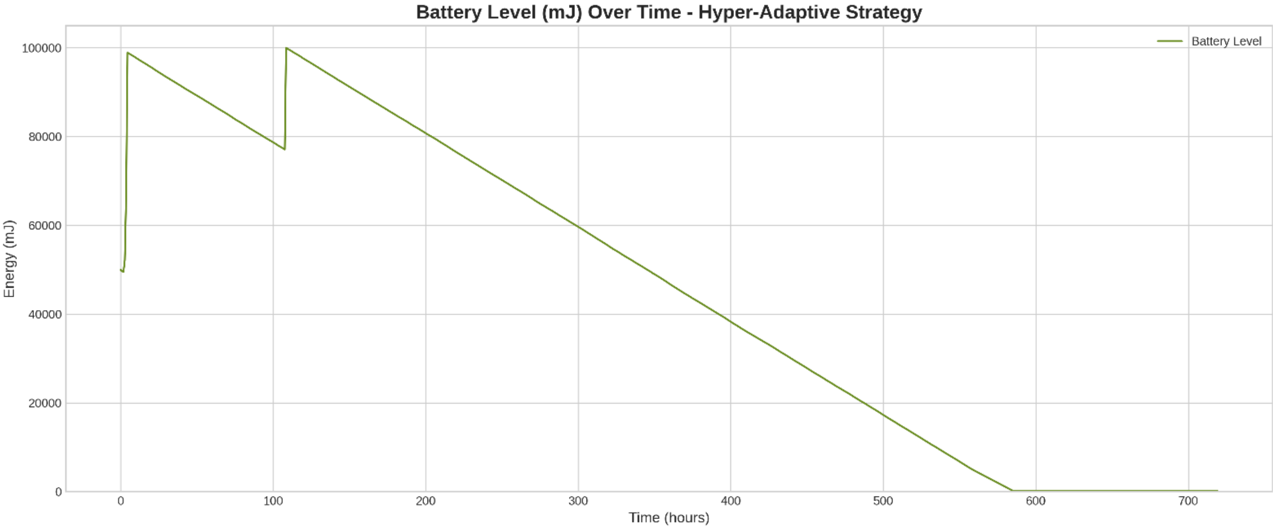

Fig. 9. Battery level over time: Hyper-adaptive strategy.

Figure 9 demonstrates the performance of a Hyper-Adaptive Strategy, illustrating the battery's energy level in millijoules over the simulation period. Similar to the previous strategy, the battery is fully recharged to its 100,000 mJ capacity around the 120-hour mark. However, immediately after this recharge, the system enters a phase of continuous and rapid linear energy decline. This steep, consistent downward slope signifies a critical imbalance, where energy consumption significantly exceeds replenishment —a direct result of the algorithm's aggressive focus on maximizing data throughput. The battery depletes at an average rate of approximately 213 mJ per hour during this phase. As a result, the stored energy is exhausted much earlier, at approximately 590 hours (24.6 days). These findings indicate that, despite prioritizing data transmission, the implementation of the Hyper-Adaptive Strategy was unable to achieve long-term energy neutrality, resulting in a shorter operational lifetime compared to the Fixed Strategy.

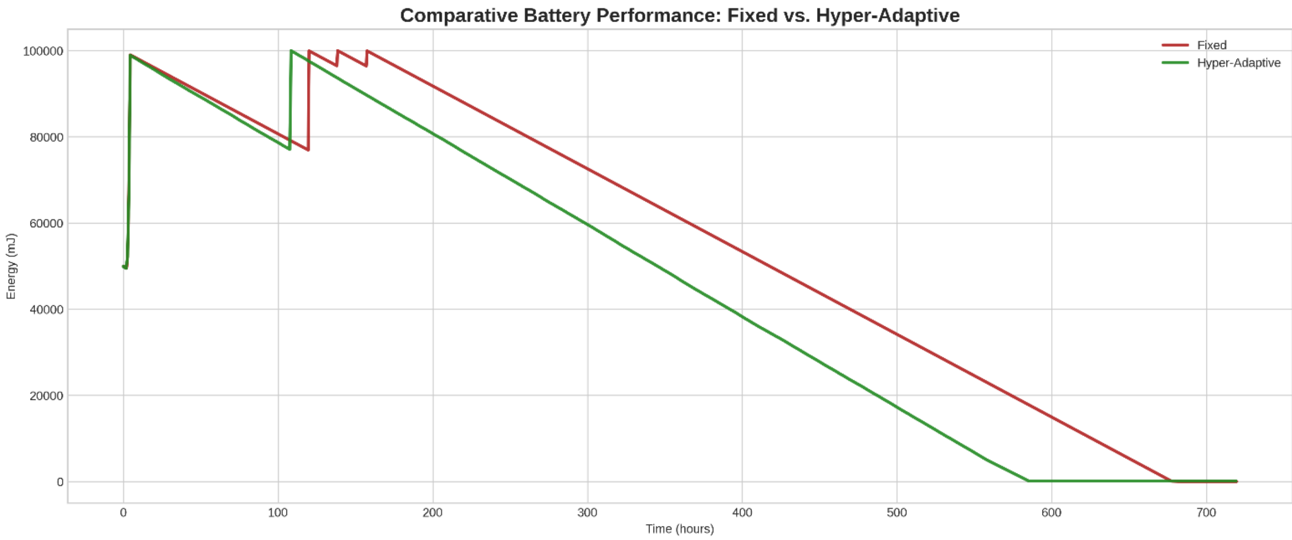

Fig. 10. Comparative battery performance: Fixed vs. hyper-adaptive.

Figure 10 presents a direct comparative analysis of the battery performance for the Fixed and Hyper-Adaptive strategies. The graph visually confirms the trade-off between energy consumption and operational longevity. The defining characteristic is the slope of each line after the final recharge event, which represents the average rate of energy depletion for each strategy.

The Hyper-Adaptive strategy's aggressive focus on data throughput results in a steep depletion rate of approximately 213 mJ per hour, leading to complete battery exhaustion at 590 hours. In contrast, the Fixed strategy exhibits a more conservative energy profile, with a slower depletion rate of 200 mJ per hour. This more gradual energy use allows it to sustain operation for a longer period, lasting 675 hours. Therefore, this figure provides clear evidence that, from a purely longevity standpoint, the Fixed strategy extended the node's operational life by approximately 3.5 days (85 hours) over the Hyper-Adaptive strategy in this simulation.

4.2. Buffer Performance Evaluation

While battery life indicates system longevity, the data buffer length serves as a direct measure of a strategy's effectiveness in managing its primary task: data throughput. A growing buffer indicates that data is being generated faster than it can be transmitted, resulting in high latency and potential data loss.

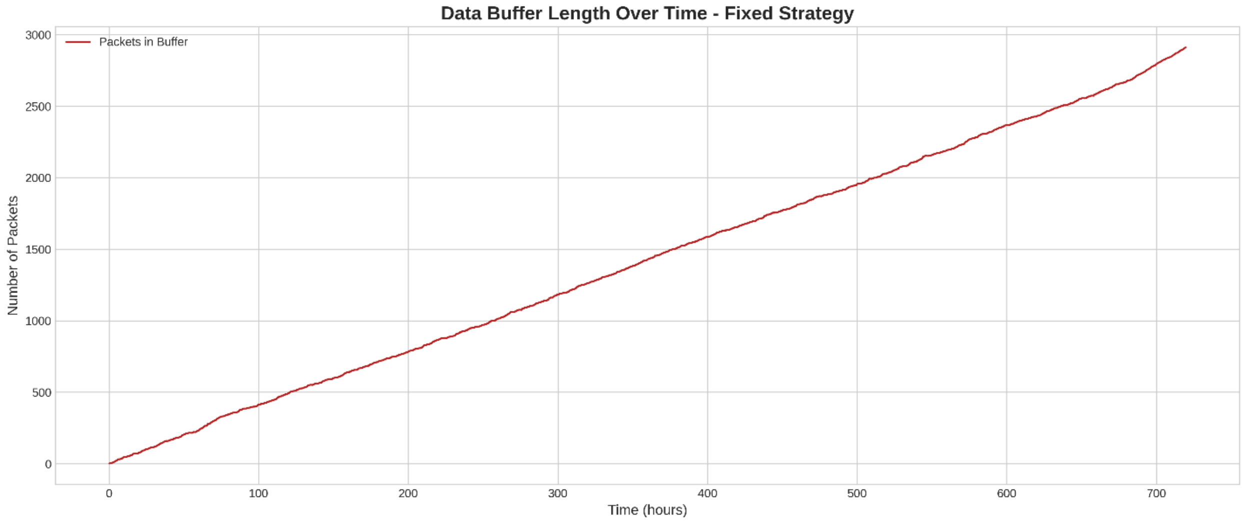

Fig. 11. Data buffer length over time: Fixed strategy.

Figure 11 illustrates the data buffer performance under the Fixed Strategy. The graph shows a persistent and nearly linear increase in the packet count throughout the entire 720-hour simulation. This upward trajectory distinctly indicates that the fixed transmission schedule is insufficient to service the incoming data workload. The buffer grows at a steady, unsustainable rate of approximately 4.1 packets per hour, culminating in a final backlog of nearly 2,900 packets. This phenomenon of continually expanding backlog confirms the inadequacy of the Fixed Strategy in managing data throughput. If this trend were to continue, a buffer overflow would be inevitable, resulting in the loss of critical data.

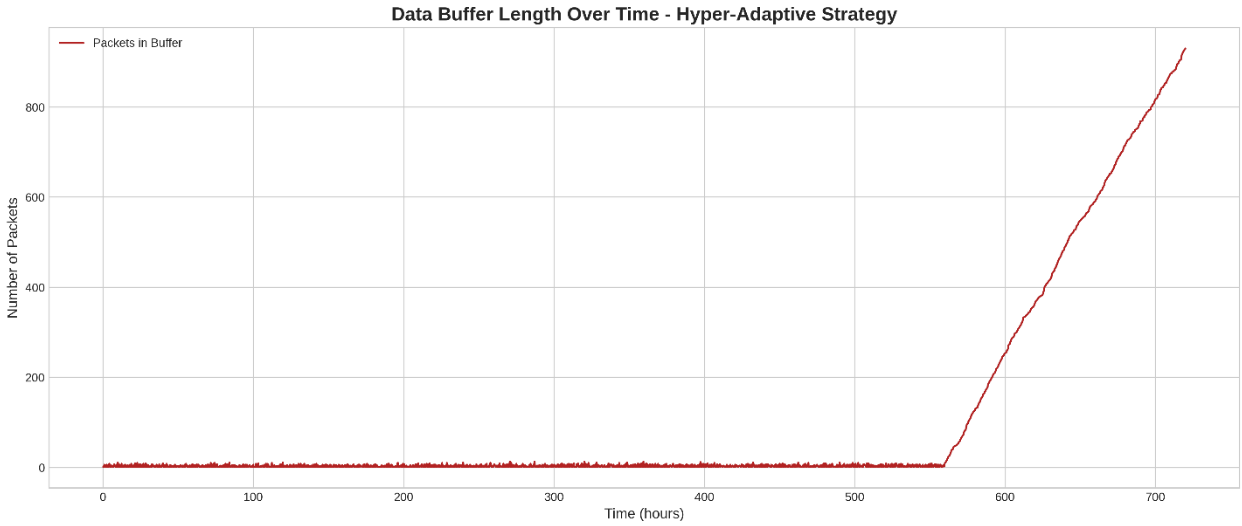

Fig. 12. Data buffer length over time: Hyper-adaptive strategy.

Figure 12 illustrates the data buffer performance under the Hyper-Adaptive Strategy. The behavior is split into two distinct phases. For the majority of its operational lifespan, from the start of the simulation until approximately hour 580, the strategy demonstrates highly effective throughput management. During these 24 days, the buffer length remains exceptionally low and stable, consistently holding fewer than 15 packets. This indicates that the adaptive transmission rate was sufficient to service the incoming data workload, preventing any significant backlog.

However, the graph exhibits a pronounced and abrupt transition around the 580-hour mark. At this point, coinciding directly with the battery depletion shown previously in Figure 8, the buffer length begins to increase at an exceedingly rapid, nearly vertical pace. This is not a gradual decline in performance but a catastrophic system failure; once the node exhausted its power, its ability to transmit data ceased entirely. The buffer subsequently accumulates all newly arriving packets, reaching a final size of over 900 packets by the end of the simulation.

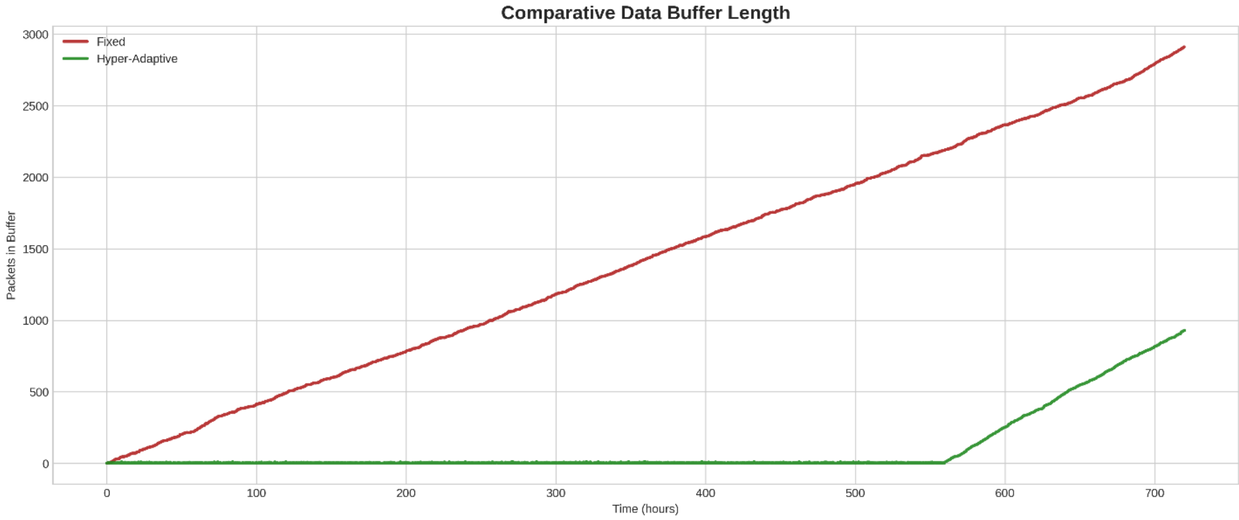

Fig. 13. Comparative data buffer length: Fixed vs. hyper-adaptive.

Figure 13 provides a compelling side-by-side comparison of data buffer performance, highlighting the profound impact of the chosen strategy on QoS. The disparity is both immediate and significant. The Fixed strategy exhibits a persistent failure from the outset, with its data buffer growing at a steady, linear rate of 4.1 packets per hour for the entire simulation, resulting in a final backlog of nearly 2,900 packets.

In stark contrast, the Hyper-Adaptive strategy demonstrates remarkable efficiency. For approximately 580 hours (24 days), it maintains a minimal backlog, with the buffer size staying consistently below 15 packets. The buffer only begins to accumulate after the node's power failure, and even then, its final backlog of just over 900 packets is less than a third of that accumulated by the Fixed strategy. This figure conclusively illustrates that, despite its shorter operational lifespan, the Hyper-Adaptive strategy's performance in its primary data handling task was vastly superior, successfully preventing data backlogs throughout its entire active operational period.

4.3. Energy Event Analysis

Analyzing the daily breakdown of energy consumption reveals the fundamental reasons behind each strategy's performance. By distinguishing between energy used for data transmission (consumed_tx) and energy consumed during the idle sleep state (consumed_sleep), the core operational logic and inherent trade-offs of each approach become clear.

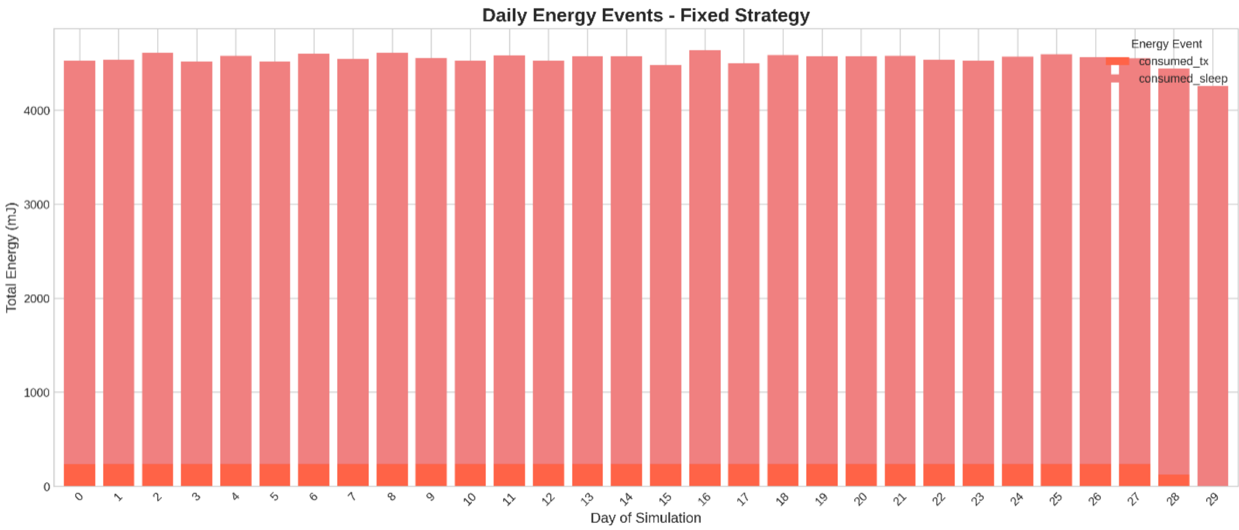

Fig. 14. Daily energy consumption breakdown: Fixed strategy.

Figure 14 provides a detailed breakdown of the daily energy consumption for the Fixed Strategy. The most prominent characteristic of this graph is the rigid consistency of its energy expenditure. Across the 30-day simulation, the total energy consumed each day remains nearly constant, hovering around 4,500 mJ. This static energy profile, maintained irrespective of solar availability or workload, is the core of the strategy's inefficiency.

A closer analysis of the stacked components reveals a critical insight into this inefficiency. The energy dedicated to the primary function of data transmission is minuscule, consistently accounting for less than 5% of the daily energy budget. The vast majority of energy is spent on the continuous power draw during sleep intervals. This disproportionate and inflexible energy allocation directly explains the unsustainable battery depletion observed previously; the system consistently expends a high, fixed amount of energy each day, with most of it contributing nothing to the primary mission of data delivery.

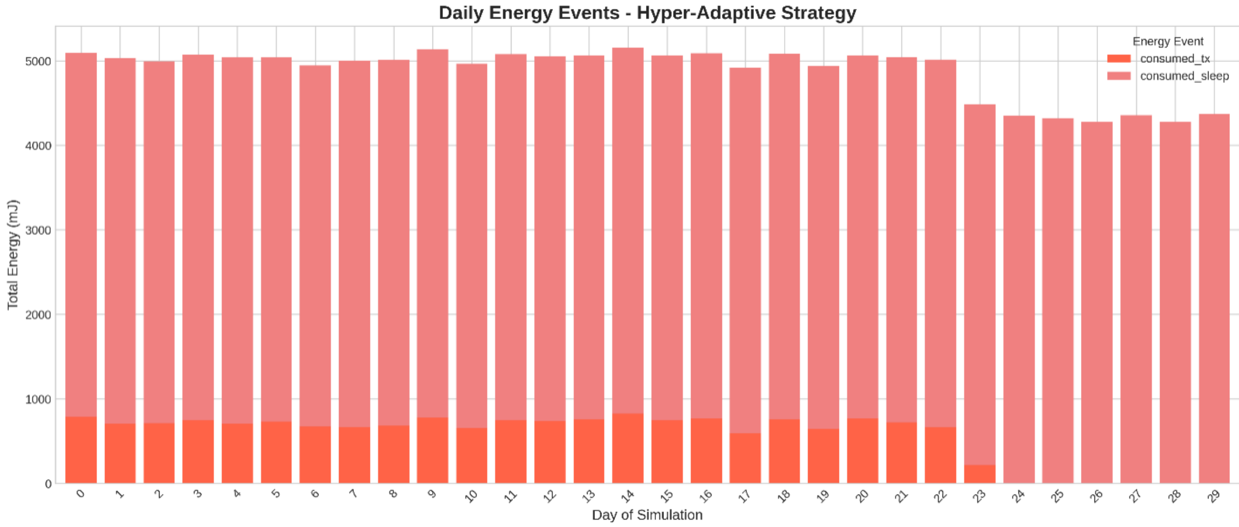

Fig. 15. Daily energy consumption breakdown: Hyper-adaptive strategy.

Figure 15 showcases the dynamic daily energy consumption profile of the Hyper-Adaptive Strategy. In stark contrast to the Fixed strategy's rigid profile, this graph reveals two distinct operational phases.

The first phase, lasting for the initial 22 days, is characterized by aggressive, high-throughput performance. During this period, the system operates at a high and relatively consistent energy level of approximately 5,100 mJ per day. Critically, the energy for data transmission constitutes a significant portion of this budget, averaging around 15-20% daily. This high expenditure on data transmission explains the strategy's superior performance in maintaining a clear data buffer.

The second phase begins with a critical shift around day 23, where the total daily energy consumption drops noticeably to under 4,500 mJ. This reduction is driven by a drastic cut in the energy allocated to transmission, which falls to less than 5% of the daily total. This change is a clear indication of the algorithm's adaptive survival mechanism activating in response to rapidly depleting battery reserves. Despite this late attempt to conserve energy, the high initial rate of consumption ultimately proves unsustainable, leading to the system's power failure.

4.4. Final System Performance

The ultimate evaluation of the two strategies hinges on their final performance outcomes, which strike a balance between system longevity and the core mission of data delivery. This section consolidates the key performance indicators to provide a definitive comparison.

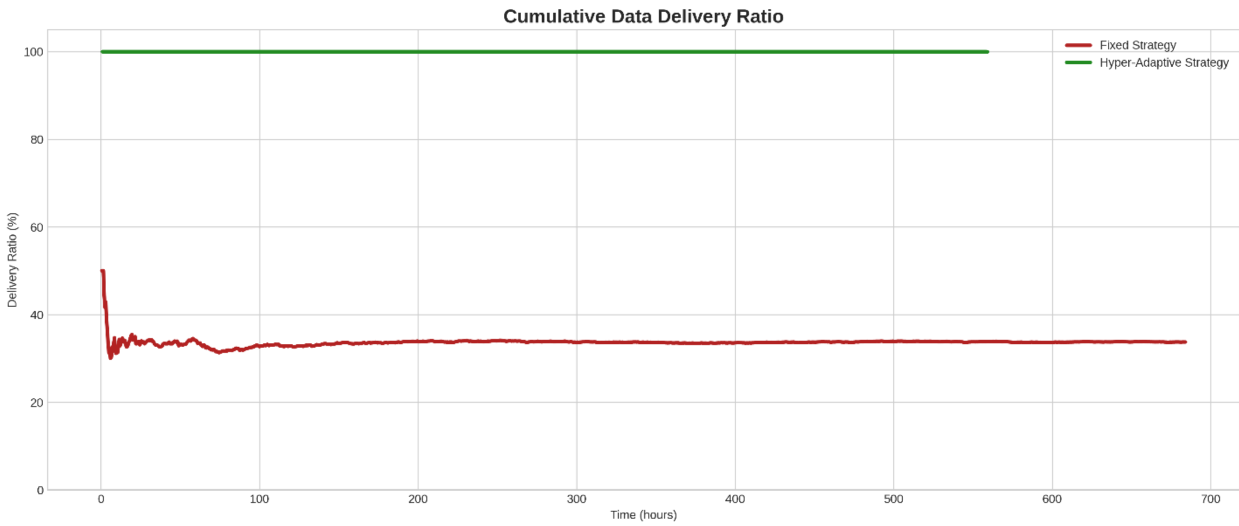

Fig. 16. Cumulative data delivery ratio: Fixed vs. hyper-adaptive.

Figure 16 presents the cumulative data delivery ratio (DDR), a key metric for evaluating the effectiveness of each strategy in its primary task. The comparison reveals a stark performance difference.

The Hyper-Adaptive strategy demonstrates perfect reliability during its operational window, maintaining a 100% DDR until its power failure around hour 590. This confirms its success in preventing data loss while active.

In stark contrast, the Fixed strategy's performance is abysmally poor. After some initial volatility, its DDR stabilizes at a consistently low level of approximately 32%. This means that despite its longer operational life, the Fixed strategy failed to transmit more than two-thirds of the data it generated, a direct consequence of its perpetually growing data buffer. This figure powerfully illustrates the core trade-off of the study: the Hyper-Adaptive strategy fulfilled its data delivery mission at the cost of sustainability, whereas the Fixed strategy sustained itself for longer but failed profoundly at its primary objective.

Table 3. A comparative summary of key performance indicators

Strategy	Data Delivery Ratio (%)	Final Battery Level (mJ)	Total Packets Generated	Total Packets Sent	Network Lifetime (Days)
Fixed	31.97%	2	4331	1385	28.1
Hyper-Adaptive	78.29%	1	4331	3391	24.6

Table 3 consolidates the final performance metrics for both strategies, providing a quantitative summary of the fundamental trade-off between QoS and network lifetime.

The Hyper-Adaptive strategy clearly prioritized QoS. It successfully transmitted 3,391 packets, achieving a final Data Delivery Ratio of 78.29%. However, this aggressive data transmission came at the cost of longevity, resulting in a shorter network lifetime of 24.6 days.

Conversely, the Fixed strategy prioritized longevity over QoS. Its less aggressive energy profile allowed it to operate for 28.1 days—3.5 days longer than its adaptive counterpart. This extended lifetime came at a severe cost to its primary function: it only managed to transmit 1,385 packets, resulting in a final DDR of just 31.97% and leaving nearly 70% of the generated data undelivered. The data unequivocally demonstrates the inherent compromise: Hyper-Adaptive optimizes for data throughput at the expense of operational duration, while Fixed extends its duration by profoundly failing at its core mission.

5. Discussion

The results of this study present a compelling narrative regarding the trade-offs inherent in designing energy-harvesting IoT systems. The markedly different outcomes of the Fixed and Hyper-Adaptive strategies provide a clear lesson on the necessity of intelligent, context-aware control.

The Fixed Duty Cycle strategy failed unequivocally, not due to a temporary lack of energy, but because of a fundamental design flaw. Its logic, predicated on an assumption of a constant and sufficient energy supply, is invalid in any solar-powered context. By activating at a fixed interval, it incurs a guaranteed net energy deficit over each 24-hour cycle, creating an irrecoverable downward trend in the battery level. Although it survived for 28.1 days, its exceedingly low Data Delivery Ratio of 31.97% indicates that it failed in its primary mission of delivering data, long before the battery was depleted.

In contrast, the Hyper-Adaptive strategy demonstrated superior performance in its primary objective of data delivery. Its success can be ascribed to several fundamental design principles. First, by making decisions based on its current battery level, it avoids the critical error of expending energy it does not possess. Second, the strategy mitigates wake-up costs by adopting a burst transmission approach. Instead of awakening for a single packet, it clears a substantial portion of its buffer in a single session, significantly enhancing energy efficiency per bit. This assertive throughput control resulted in near-zero data latency for over three weeks of the simulation.

Nevertheless, this comparison highlights the fundamental trade-off between QoS and network longevity. The assertiveness of the Hyper-Adaptive strategy is a double-edged sword; by pushing the limits of its energy resources to optimize throughput, it operates with a narrower safety margin. Consequently, its operational lifespan was 24.6 days, approximately four days shorter than the less efficient but more conservative Fixed strategy. The designation 'Hyper-Adaptive' accordingly reflects its design philosophy: to maximize performance, while accepting the inherent risk of a reduced lifespan. Its failure was not a defect in its logic but a direct and foreseeable consequence of its high-risk, high-reward conceptual framework.

This leads to an essential realization for system designers: the best algorithm is entirely dependent on the application context. For applications where near-real-time data is crucial and occasional downtime is acceptable, such as opportunistic environmental monitoring, the high-QoS Hyper-Adaptive strategy is more effective. Conversely, for applications emphasizing long-term reliability over immediate data delivery, such as critical infrastructure monitoring, a more conservative adaptive algorithm would be more suitable.

While this study provides a clear performance baseline, it is important to contextualize its contributions. The comparison was made against a non-adaptive baseline to establish the fundamental value of adaptive control. A direct simulation-based comparison with other adaptive heuristics, such as the AIMD-inspired FSM proposed by Giordano et al.

[19], would be a valuable next step. We hypothesize that while the FSM approach might offer greater stability by modulating the data generation rate, our burst-based method likely achieves higher energy efficiency per transmitted bit, a critical factor for data-intensive applications.

Furthermore, this study intentionally employed a single-node simulation to isolate and rigorously evaluate the core energy management logic, thereby eliminating confounding variables from network-level interactions. Extending this work to a multi-node network would introduce new challenges, such as Medium Access Control (MAC) layer contention and routing energy costs, which would undoubtedly impact the performance of both strategies.

6. Conclusion

This research evaluated adaptive duty cycling for EH-IoT sensor networks through simulation, comparing a Hyper-Adaptive algorithm with a Fixed-Cycle baseline. The findings demonstrate that intelligent control is essential for sensor nodes and underscore the critical trade-off between data throughput and sustainability. The Fixed Duty Cycle underperformed, achieving only 31.97% data throughput over 28.1 days. In contrast, the Hyper-Adaptive strategy attained a 78.29% Data Delivery Ratio by dynamically balancing energy consumption, amortizing wake-up costs, and clearing buffers, thereby ensuring near-optimal QoS for most of its operational life.

However, this high performance was achieved at the expense of longevity. The Hyper-Adaptive algorithm's aggressive, QoS-focused approach resulted in earlier battery depletion at 24.6 days. This highlights a fundamental conclusion: the most effective strategy is not solely the one that transmits the most data but rather the one that optimally balances an application's specific QoS requirements with the necessity for long-term network sustainability. This research effectively quantifies this trade-off, providing a valuable framework for designing and deploying robust and efficient energy-harvesting IoT networks.

While this study offers valuable insights, its limitations delineate several promising avenues for future research. The single-node simulation must be extended to a multi-node framework to examine how adaptive duty cycling influences network-level phenomena, such as data routing and MAC layer contention. Second, the reactive Hyper-Adaptive algorithm could be enhanced with predictive capabilities. Future work should investigate the integration of machine learning models, such as LSTM networks, to forecast solar energy availability from historical data or weather forecasts, enabling more proactive and conservative behavior during anticipated periods of low energy. Finally, validating these algorithms on physical hardware is an essential step toward real-world deployment, enabling an assessment of performance under the nuanced energy costs of actual electronic components.

References

[1] W. Chen, Q. Cao, B. Cao, and B. Jin, “An innovative coverage optimization method for smart information monitoring in agricultural IoT using the multi-strategy Pelican optimization algorithm”, Sci Rep, vol. 15, no. 1, p. 12634, 2025.

[2] M. R. Sarker, A. Riaz, M. S. H. Lipu, M. H. M. Saad, M. N. Ahmad, R. A. Kadir, and J. L. Olazagoitia, “Micro energy harvesting for IoT platform: Review analysis toward future research opportunities”, Heliyon, vol. 10, no. 6, 2024.

[3] M. U. Mushtaq, H. Venter, A. Singh, and M. Owais, “Advances in energy harvesting for Sustainable wireless sensor networks: challenges and opportunities”, Hardware, vol. 3, no. 1, p. 1, 2025.

[4] N. Munusamy, “Localization and Deployment Considerations into Quality of Service Optimization for Energy-Efficient Wireless Sensor Networks”.

[5] M. Imani, M. Ali, and H. R. Arabnia, “Power-saving asynchronous quorum-based protocols for maximal neighbour discovery”, arXiv preprint arXiv:2007.04003, 2020.

[6] V. Nkemeni, F. Mieyeville, G. S. Kuaban, P. Czekalski, K. Tokarz, W. B. Nsanyuy, E. M. Deussom Djomadji, M. L. Katche, P. Tsafack, and B. Zieliński, “Evaluation of green strategies for prolonging the lifespan of linear wireless sensor networks”, Sensors, vol. 24, no. 21, p. 7024, 2024.

[7] S. Hudda and K. Haribabu, “A review on WSN based resource constrained smart IoT systems”, Discover Internet of Things, vol. 5, no. 1, p. 56, 2025.

[8] M. Rajput and R. N. Yadav, “Machine and deep learning driven energy efficient clustering in IoT-WSNs: A Review”, IEEE Sens J, 2025.

[9] G. S. Kuaban, V. Nkemeni, and P. Czekalski, “An analytical framework for optimizing the renewable energy dimensioning of green IoT systems in pipeline monitoring”, Sensors, vol. 25, no. 10, p. 3137, 2025.

[10] F. Giuliano, A. Pagano, D. Croce, G. Vitale, and I. Tinnirello, “Dynamic transmission adaptation algorithms for battery-free LoRaWAN networks”, Internet of Things, p. 101706, 2025.

[11] M. Gerndt, M. Ispir, I. Nunez, and S. Benedict, “Energy-aware duty cycle management for solar-powered IoT devices”, Sensors, vol. 25, no. 14, p. 4500, 2025.

[12] Y. Azizi and A. Yazdizadeh, “Passivity-based adaptive control of a 2-DOF serial robot manipulator with temperature dependent joint frictions”, Int J Adapt Control Signal Process, vol. 33, no. 3, pp. 512–526, 2019.

[13] O. L. A. López, M. Ashraf, S. Nasser, G. M. de Jesus, R. K. Singh, M. C. Filippou, and J. Famaey, “Foundations for energy-aware zero-energy devices: from energy sensing to adaptive protocols”, arXiv preprint arXiv:2507.22740, 2025.

[14] K. Saurabh, M. M. Tripathi, and S. Mahapatra, “Efficient

utilization of energy in IoT devices using machine learning algorithms", Int. J. Exp. Res. Rev, vol. 47, pp. 133–145, 2025.

[15] R. Priyadarshi, R. R. Kumar, R. Ranjan, and P. V. Kumar, "AI-based routing algorithms improve energy efficiency, latency, and data reliability in wireless sensor networks", Sci Rep, vol. 15, no. 1, p. 22292, 2025.

[16] N. Charef, A. Ben Mnaouer, M. Aloqaily, O. Bouachir, and M. Guizani, "Artificial intelligence implication on energy sustainability in Internet of Things: A survey", Inf Process Manag, vol. 60, no. 2, p. 103212, 2023.

[17] O. Alamu, T. O. Olwal, and E. M. Migabo, "Machine learning applications in energy harvesting internet of things networks: A review", IEEE Access, 2025.

[18] S. Sarang, G. M. Stojanović, M. Drieberg, S. Stankovski, K. Bingi, and V. Jeoti, "Machine learning prediction based adaptive duty cycle MAC protocol for solar energy harvesting wireless sensor networks", IEEe Access, vol. 11, pp. 17536–17554, 2023.

[19] M. Giordano, S. Cortesi, P.-V. Mekikis, M. Crabolu, G. Bellusci, and M. Magno, "Energy-aware adaptive sampling for self-sustainability in resource-constrained IoT devices", in Proceedings of the 11th International Workshop on Energy Harvesting & Energy-Neutral Sensing Systems, 2023, pp. 65–71.

[20] R. Mohammadi and Z. Shirmohammadi, "DRDC: Deep reinforcement learning based duty cycle for energy harvesting body sensor node", Energy Reports, vol. 9, pp. 1707–1719, 2023.

[21] D. E. Ruíz-Guirola, O. L. A. López, S. Montejo-Sánchez, I. L. Mayorga, Z. Han, and P. Popovski, "Intelligent duty cycling management and wake-up for energy harvesting IoT networks with correlated activity", in 2024 58th Asilomar Conference on Signals, Systems, and Computers, IEEE, 2024, pp. 1812–1818.

[22] H. T. Tran, C. V Nguyen, H. T. T. Nguyen, and M. T. Nguyen, "Energy harvesting for devices in wireless sensor networks: A Review", EAI Endorsed Trans. Internet Things, vol. 9, no. 2, pp. 1–11, 2023.

[23] F. Shan, J. Luo, Q. Jin, L. Cao, W. Wu, Z. Ling, and F. Dong, "Optimal harvest-then-transmit scheduling for throughput maximization in time-varying RF powered systems", IEEE Journal on Selected Areas in Communications, 2024.