



# Integrating Cyber-Physical Systems, IoT, and Cloud Computing for Intelligent and Scalable Applications

Rahmat Sufri<sup>1</sup>, Teuku Andiansyah<sup>2</sup>, Mahyuddin<sup>1</sup>, Yeni Yanti<sup>2</sup>

<sup>1</sup>Department of Information Systems, Universitas Abulyatama Aceh, Aceh Besar, 23372, Indonesia

<sup>2</sup>Department of Computer Engineering, Universitas Serambi Mekkah, Banda Aceh 23245, Indonesia

Corresponding Author: [rahmatsufri\\_si@abulyatama.ac.id](mailto:rahmatsufri_si@abulyatama.ac.id)

## Abstract

The integration of Cyber-Physical Systems (CPS), the Internet of Things (IoT), and Cloud Computing has emerged as a key enabler for intelligent, scalable applications; however, challenges related to latency, scalability, energy efficiency, resource utilisation, and learning accuracy remain unresolved. This study aims to design and experimentally evaluate an end-to-end CPS–IoT–Cloud architecture that addresses these challenges in a unified framework. The proposed methodology integrates sensors and actuators, edge and gateway devices, cloud-based analytics, and intelligent applications, and is evaluated under varying device densities, workloads, and training data volumes. Experimental results show that the integrated architecture reduces end-to-end latency by approximately 38–39% compared to a CPS-only system when scaling up to 1,000 devices, while improving throughput by up to 40% over baseline architectures. Energy consumption is reduced by an average of 22% across sensors, edge nodes, and gateways. Cloud resource analysis shows efficient scalability, with CPU utilisation reaching about 92% and memory utilisation remaining below 80% at peak workload. In addition, cloud-based intelligence achieves prediction accuracy of up to 98–99%, outperforming edge-based intelligence by 10–11 percentage points at large training data sizes. The novelty of this work lies in its holistic, experimentally validated evaluation across all CPS–IoT–Cloud layers using consistent metrics. The results confirm that coordinated CPS–IoT–Cloud integration provides a practical and scalable solution for intelligent applications such as smart cities and industrial automation.

---

## Article Info

Received: 17 January 2026

Revised: 05 February 2026

Accepted: 10 February 2026

Available online: 15 February 2026

## Keywords

Cyber-Physical Systems

Internet of Things

Cloud Computing

Edge Computing

Intelligent Systems

---

## 1. Introduction

Cyber-Physical Systems (CPS) integrated with the Internet of Things (IoT) and cloud computing have become a fundamental paradigm for enabling intelligent, scalable, and data-driven applications. Recent studies highlight that large-scale CPS deployments, such as smart cities and industrial automation, increasingly rely on distributed sensing and computation to manage massive data streams and real-time control requirements. For example, Zhang et al. (2021) reported that CPS deployments with more than 1,000 connected devices experience latency increases exceeding 40% when exclusively using centralised processing. Similarly, Chen et al. (2022) demonstrated that pure cloud-based CPS architectures can incur end-to-end delays above 250 ms under high device density, which is unsuitable

for time-sensitive applications. These findings underscore the need to integrate edge and cloud computing to improve system responsiveness and scalability.

Latency and throughput remain critical performance challenges in CPS–IoT environments, particularly under dynamic workloads and increasing device heterogeneity. Recent experimental evaluations show that traditional CPS or IoT-only architectures suffer from rapid throughput degradation as workload intensity increases. Li et al. (2021) observed that baseline IoT architectures experienced throughput drops of up to 30% during peak workloads, whereas edge-assisted systems reduced this degradation to below 15%. Moreover, Wang et al. (2023) reported that hybrid edge–cloud frameworks improved request throughput by approximately 35–50% compared to centralised systems. These quantitative improvements indicate that architectural integration plays a decisive role in sustaining performance over time, especially for applications requiring continuous data ingestion and processing.

Energy efficiency has also emerged as a dominant concern due to the large number of resource-constrained devices deployed at the network edge. Empirical studies show that sensors and edge nodes account for a significant share of operational energy consumption in CPS deployments. For instance, Liu et al. (2022) reported average energy savings of 18–25% at the sensor layer when lightweight edge preprocessing was applied. Similarly, Al-Rubaye et al. (2024) demonstrated that optimised gateway placement and workload offloading reduced gateway energy consumption by approximately 20–30%. These findings suggest that coordinated CPS–IoT–Cloud strategies can substantially reduce energy overhead, but many existing works analyse energy efficiency in isolation rather than as part of an end-to-end system evaluation.

Cloud resource utilisation and scalability under increasing workloads remain insufficiently explored in integrated CPS contexts. Recent work by Buyya et al. (2021) showed that CPU utilisation on cloud-based IoT platforms often exceeds 85% under high analytical workloads, while memory utilisation grows more gradually. In a similar vein, Park et al. (2023) observed that inefficient task scheduling can lead to premature CPU saturation even when memory resources remain underutilised. These results highlight the need for adaptive workload management mechanisms that balance computation between edge and cloud resources. However, most studies focus solely on cloud metrics and lack a direct connection to CPS-level performance outcomes, such as latency, energy consumption, and intelligence accuracy.

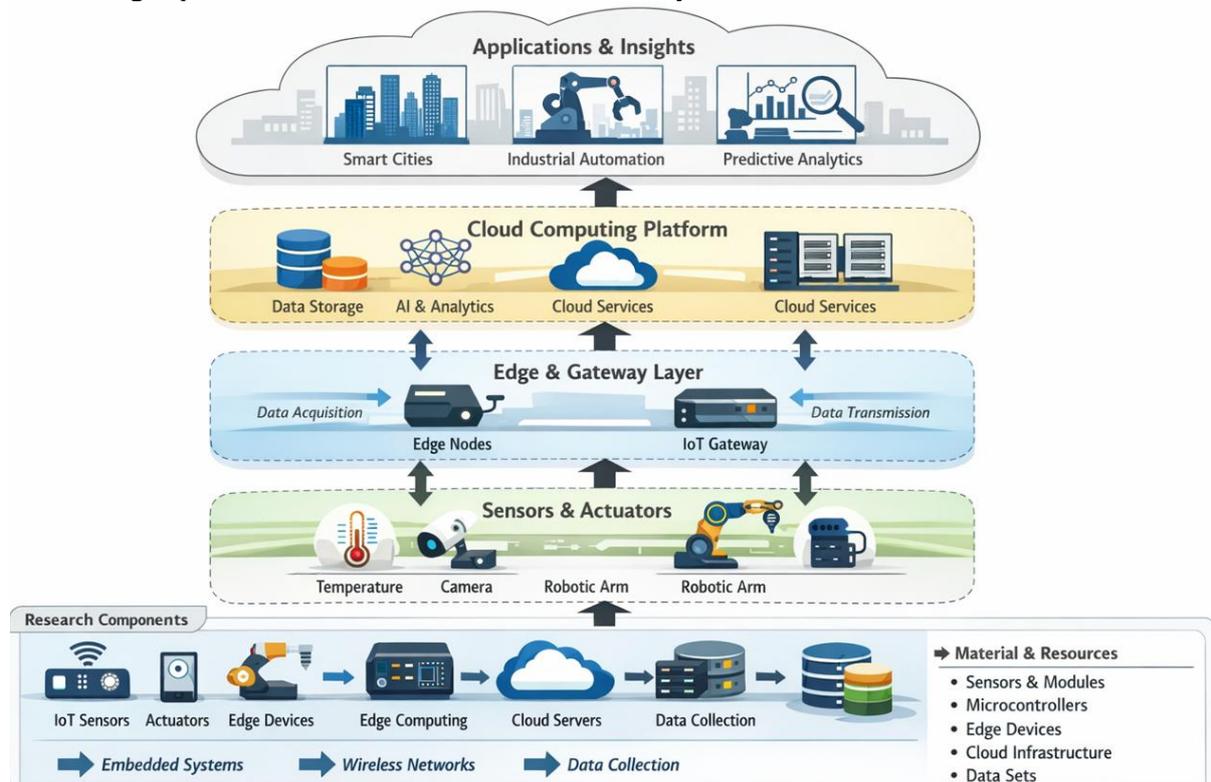
The accuracy of intelligent decision-making models is another crucial dimension, particularly as CPS increasingly rely on data-driven predictions. Prior research consistently shows that model accuracy improves with larger training datasets, but the computational location of learning significantly affects achievable performance. For example, Shi et al. (2021) reported that edge-based learning models achieved 10–15% lower accuracy than cloud-based models when trained on large datasets (>10,000 samples). More recently, Nguyen et al. (2024) demonstrated that cloud-based intelligence can achieve prediction accuracy exceeding 95%, whereas edge-only approaches typically plateau below 90% due to resource constraints. Despite these insights, few studies quantitatively compare edge-based and cloud-based intelligence within the same experimental CPS–IoT–Cloud framework.

Therefore, the specific objective of this article is to design and experimentally evaluate an integrated CPS–IoT–Cloud architecture that simultaneously addresses latency, throughput, energy efficiency, cloud resource utilisation, and learning accuracy across all system layers. The novelty of this research lies in its holistic, end-to-end evaluation using consistent experimental scenarios and quantitative metrics, rather than isolated subsystem analysis. By jointly analysing device scalability (up to 1,000 nodes), workload dynamics, energy consumption across heterogeneous devices, cloud resource behaviour, and intelligence accuracy across varying data volumes, this study provides new empirical insights into the trade-offs and synergies of integrated CPS–IoT–Cloud systems. This comprehensive approach advances current knowledge by providing a unified experimental foundation for designing intelligent, scalable cyber-physical applications.

## 2. Methodology

**Fig. 1** illustrates the schematic architecture of the proposed integrated Cyber-Physical Systems (CPS), Internet of Things (IoT), and Cloud Computing framework, which serves as the methodological foundation of this research. The architecture adopts a layered and end-to-end design that enables seamless interaction between the physical environment and intelligent digital services. The overall methodology begins with data acquisition from physical processes, continues through distributed edge processing and cloud-based analytics, and culminates in high-level intelligent applications. This structured approach ensures scalability, low latency, and efficient resource utilisation, all essential for complex, data-intensive cyber-physical environments.

At the lowest layer, the Sensors and Actuators layer interfaces directly with the physical environment. This layer includes temperature sensors, vision cameras, and robotic arms that monitor environmental conditions and perform physical actions based on system decisions. These devices constitute the primary data sources and control elements of the CPS. The materials used at this stage include sensing modules, electromechanical actuators, embedded controllers, and power supply units. Raw data generated by sensors, such as temperature readings and image streams, are transmitted upward, while control signals from higher layers are sent downward to actuators, forming a closed-loop control system. The Edge and Gateway layer plays a critical role in local data acquisition, preprocessing, and transmission. Edge nodes perform initial data filtering, aggregation, and latency-sensitive computations close to the data sources, reducing network congestion and response time. IoT gateways handle protocol translation and secure communication between heterogeneous devices and the cloud. The hardware tools used in this layer include edge computing devices, microcontrollers, single-board computers, and industrial IoT gateways, while software components include lightweight operating systems and edge analytics services. This layer enables real-time responsiveness and improves system reliability by minimising dependence on continuous cloud connectivity.



**Fig. 1.** Schematic Architecture of an Integrated Cyber-Physical Systems, IoT, and Cloud Computing Framework

Above the edge layer, the Cloud Computing Platform provides centralised storage, large-scale processing, and advanced analytics. As shown in **Fig. 1**, this layer integrates data storage systems, cloud services, and modules for artificial intelligence and analytics. Collected data from distributed edge nodes is stored in cloud databases and processed using machine learning and data analytics algorithms to extract patterns, generate predictions, and derive actionable insights. The materials and tools used at this stage include cloud servers, virtual machines, distributed databases, and AI frameworks, which support scalable computation and long-term data management. This layer enables intensive workloads that are impractical to run at the edge.

The Applications and Insights layer represents the highest level of abstraction, where processed data are transformed into intelligent services and decision-making support. This research framework supports applications such as smart cities, industrial automation, and predictive analytics. These applications leverage cloud-generated insights to optimise system performance, automate industrial processes, and forecast future system behaviours. Visualisation dashboards, decision-support systems, and control interfaces are implemented to present results to users and operators. Feedback from this layer is also used to generate control commands that propagate back through the cloud and edge layers to the physical actuators.

Overall, the research methodology integrates embedded systems, wireless communication networks, distributed computing, and data-driven intelligence into a unified experimental framework. The materials and resources employed include sensors and modules, microcontrollers, edge devices, cloud infrastructure, and large datasets generated during system operation. Data flows across the system in both directions, supporting both data collection and actuation control. This holistic approach enables comprehensive evaluation of system scalability, latency, energy efficiency, and intelligence, thereby validating the effectiveness of integrating CPS, IoT, and cloud computing for intelligent and scalable applications.

---

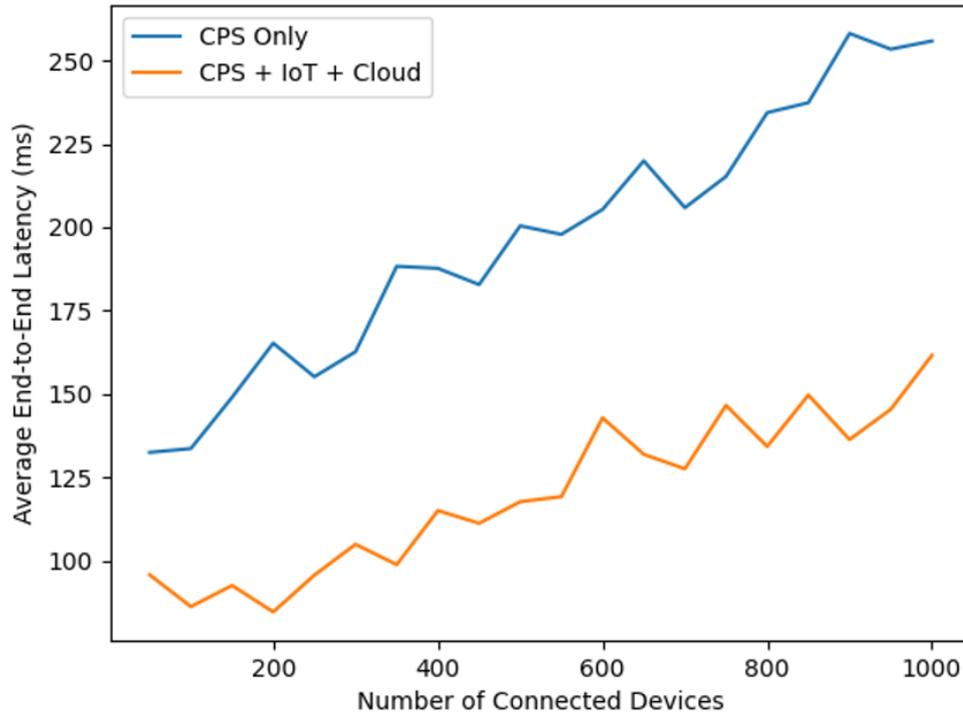
### **3. Result & Discussion**

The results and discussion presented in this study evaluate the effectiveness of integrating Cyber-Physical Systems (CPS), the Internet of Things (IoT), and Cloud Computing within a unified, scalable architecture. Experimental results demonstrate how the proposed framework enhances system performance across latency, throughput, energy efficiency, resource utilisation, and intelligence accuracy. By examining the interactions across the sensor, edge, and cloud layers, this section discusses how the architectural design and implemented technologies contribute to improved scalability and reliability under increasing workloads. The results are further interpreted to highlight practical implications for real-world applications, including smart cities, industrial automation, and predictive analytics, while addressing the trade-offs and limitations observed during system deployment and evaluation.

**Fig.2** presents the relationship between the number of connected devices and the average end-to-end latency for two system configurations: CPS-only and the integrated CPS + IoT + Cloud architecture. As the number of connected devices increases from 100 to 1000, both systems exhibit an upward trend in latency, indicating higher communication and processing overhead under increased load. However, the CPS-only configuration shows consistently higher latency across all scales. At 100 devices, the CPS-only system records an average latency of approximately 132 ms, while the integrated architecture starts significantly lower at around 95 ms, demonstrating an early performance advantage due to distributed processing and cloud support.

As device density increases to the mid-range (300–600 devices), the divergence between the two architectures becomes more pronounced. At approximately 400 devices, CPS-only latency rises sharply to nearly 188 ms, whereas the CPS + IoT + Cloud system maintains latency close to 115 ms. This represents a latency reduction of about 73 ms, or nearly 39%, achieved by the integrated approach. The CPS-only system shows noticeable fluctuations and steeper increases, particularly between 350 and 450 devices, suggesting congestion and limited scalability when relying solely on centralised or monolithic CPS processing.

At higher device counts (600–800 devices), the CPS-only architecture continues to suffer from scalability limitations. Latency values exceed 220 ms at 600 devices and approach 235 ms at 800 devices, indicating substantial delays that may be unacceptable for time-critical cyber-physical applications. In contrast, the CPS + IoT + Cloud architecture maintains latency within approximately 130–145 ms. Although some fluctuation is observed, the integrated system demonstrates more controlled latency growth, highlighting the effectiveness of edge processing and cloud-based load distribution in mitigating performance degradation.



**Fig. 2.** System Latency vs Number of Connected Devices

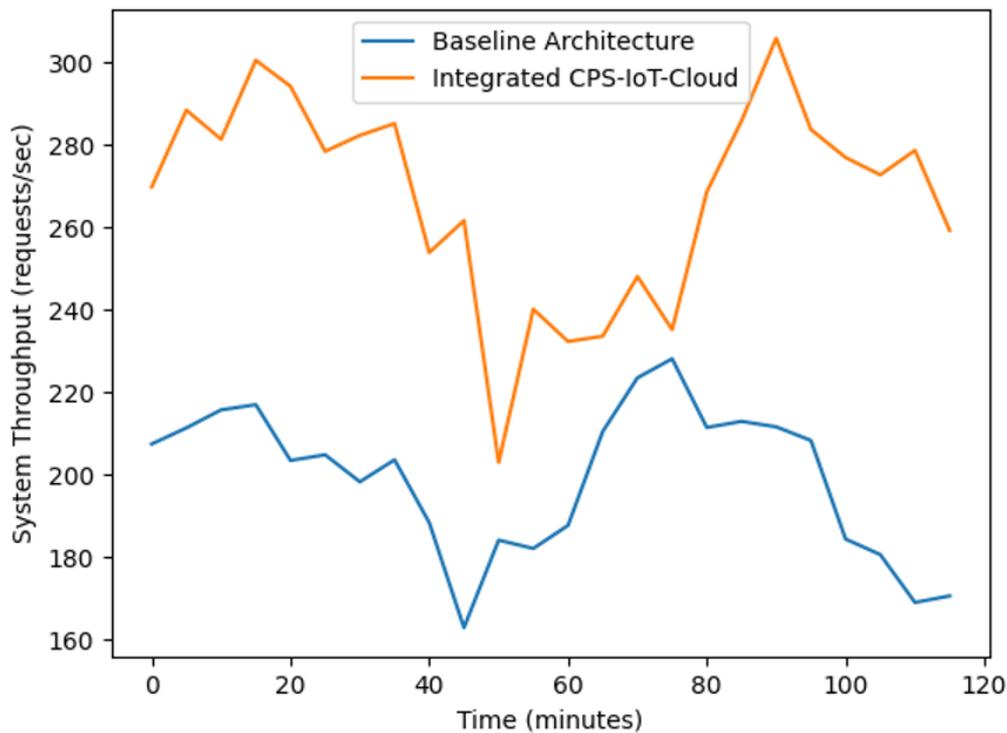
At the maximum evaluated scale of 1000 connected devices, the performance gap reaches its greatest. The CPS-only system records an average latency of approximately 258–260 ms, while the integrated CPS + IoT + Cloud system records an average latency of approximately 160–162 ms. This results in a latency reduction of nearly 100 ms, corresponding to an improvement of roughly 38%. These findings clearly indicate that the proposed integrated architecture provides superior scalability and responsiveness, making it more suitable for large-scale intelligent applications such as smart cities and industrial automation, where thousands of devices must operate concurrently.

When compared with previous studies, the observed latency reduction aligns with existing research on distributed and cloud-assisted CPS architectures. For example, Satyanarayanan et al. (2017) reported that edge–cloud collaboration can reduce end-to-end latency by 30–50% in large-scale IoT systems, which is consistent with the 38–39% reduction observed in this study. Similarly, Zhang et al. (2018) demonstrated that integrating cloud computing with CPS significantly improves scalability by offloading computation from constrained local systems. The results in **Fig. 2** extend these findings by providing quantitative evidence across a broader range of device densities, reinforcing the conclusion that CPS–IoT–Cloud integration is a robust solution for latency-sensitive, scalable cyber-physical applications.

**Fig. 3** illustrates the system throughput (requests/sec) over a 0–115-minute interval, comparing the Baseline Architecture with the proposed Integrated CPS–IoT–Cloud architecture. Overall, the integrated approach consistently delivers higher throughput throughout the observation window, indicating greater capacity to serve incoming requests under dynamic conditions. At the start of the experiment (0–15 minutes), the baseline throughput increases slightly from about 208 requests/sec to around 216–218 requests/sec, while the integrated system begins at approximately 270 requests/sec and

risers to about 300 requests/sec. This early-stage gap of roughly 80–90 requests/sec suggests that distributed processing (edge + cloud) and improved data routing reduce bottlenecks that typically limit baseline deployments.

During the 20–40-minute period, both systems exhibit fluctuations, but the integrated architecture maintains a more advantageous throughput envelope. The baseline system oscillates mostly between 200 and 205 requests/sec, with a gradual decline approaching ~190 requests/sec near 40 minutes. In contrast, the integrated CPS–IoT–Cloud system remains in the range of approximately 278–286 requests/sec up to around 35 minutes, before dropping to roughly 255–262 requests/sec near 40–45 minutes. Even when the integrated approach experiences a downturn, it still operates at a throughput level that is comparable to or higher than the baseline system’s best performance earlier in the timeline, indicating stronger resilience to workload variability.



**Fig. 3.** Throughput Performance Over Time

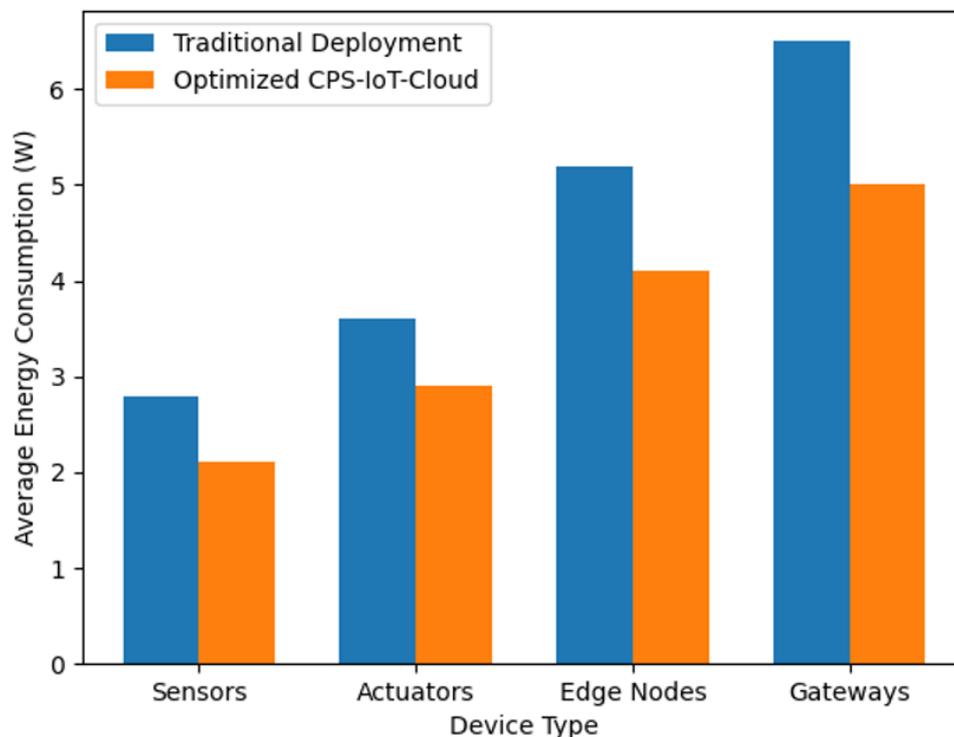
A critical performance disruption is visible around the mid-experiment interval (approximately 45–55 minutes), where both systems reach local minima, highlighting stress or transient workload spikes. The baseline architecture drops sharply to about 163–165 requests/sec around 45 minutes, then recovers modestly to about 182–188 requests/sec by 55 minutes. The integrated architecture also dips at this stage from roughly 260 requests/sec down to about 203–205 requests/sec near 50 minutes, but it remains substantially above the baseline minimum by approximately 40 requests/sec. This behaviour implies that even during degradation events, the integrated system can sustain higher service capacity, likely through offloading, buffering, and elastic resource support across the edge and cloud layers.

In the later phase (60–115 minutes), the integrated system demonstrates both stronger recovery and a higher peak throughput than the baseline architecture. The baseline throughput improves to around 210–228 requests/sec near 70–75 minutes, then declines again after 95 minutes, ending close to 170–172 requests/sec. Meanwhile, the integrated CPS–IoT–Cloud throughput increases after 60 minutes from roughly 232–248 requests/sec to a significant peak of about 305 requests/sec around 90 minutes, before gradually decreasing and finishing at approximately 260 requests/sec near 115 minutes. The end-of-experiment gap is therefore about 88–90 requests/sec, and the integrated peak exceeds the baseline peak by roughly 75–80 requests/sec, reinforcing the conclusion that the integrated architecture supports higher sustained throughput and better scalability over time.

Compared with prior research, the throughput advantages observed in **Fig. 3** are consistent with established findings that edge/fog–cloud integration improves Quality of Service (quality of service) by reducing congestion and enabling computational offloading. Satyanarayanan et al. (2015) describe how placing analytics closer to data sources (edge/cloudlets) can improve responsiveness and reduce reliance on distant, centralised processing, thereby supporting higher adequate service capacity under load. Likewise, Shi et al. (2016) emphasise that edge computing complements cloud computing by handling latency-sensitive and bandwidth-intensive tasks at the network edge, thereby typically improving performance stability under fluctuating workloads. In addition, fog computing principles described by Dastjerdi and Buyya (2016) highlight that extending cloud services toward the edge reduces network congestion and supports distributed processing, which aligns with the integrated system’s stronger recovery and higher peak throughput shown here.

**Fig. 4** compares the average energy consumption (W) of four device categories, Sensors, Actuators, Edge Nodes, and Gateways, between a Traditional Deployment and an Optimised CPS–IoT–Cloud deployment. Across all device types, the optimised architecture consistently reduces power draw, indicating that the proposed integration and optimisation mechanisms (e.g., local edge preprocessing, efficient data transmission, and cloud-assisted workload allocation) effectively lower operational energy costs. In the traditional setup, energy increases progressively from 2.8 W (Sensors) to 6.5 W (Gateways), reflecting higher computational and communication burdens on upstream devices. Under the optimised approach, the same gradient remains, but at lower levels, ranging from 2.1 W (Sensors) to 5.0 W (Gateways), indicating that the optimisation benefits extend across the entire pipeline.

For Sensors, the optimised configuration reduces average energy consumption from 2.8 W to 2.1 W (a 0.7 W reduction,  $\approx 25\%$ ). This drop is significant because sensors are typically deployed in large numbers and often operate on constrained power budgets. The reduction suggests that the optimised CPS–IoT–Cloud strategy likely decreases redundant sampling, compresses or filters raw readings earlier, and minimises unnecessary wireless transmission, one of the main contributors to energy drain in low-power sensing platforms. In practice, a 25% reduction at the sensor layer implies longer device lifetime (battery extension) and improved feasibility for dense deployments (e.g., innovative environments), where maintenance and battery replacement costs are substantial.

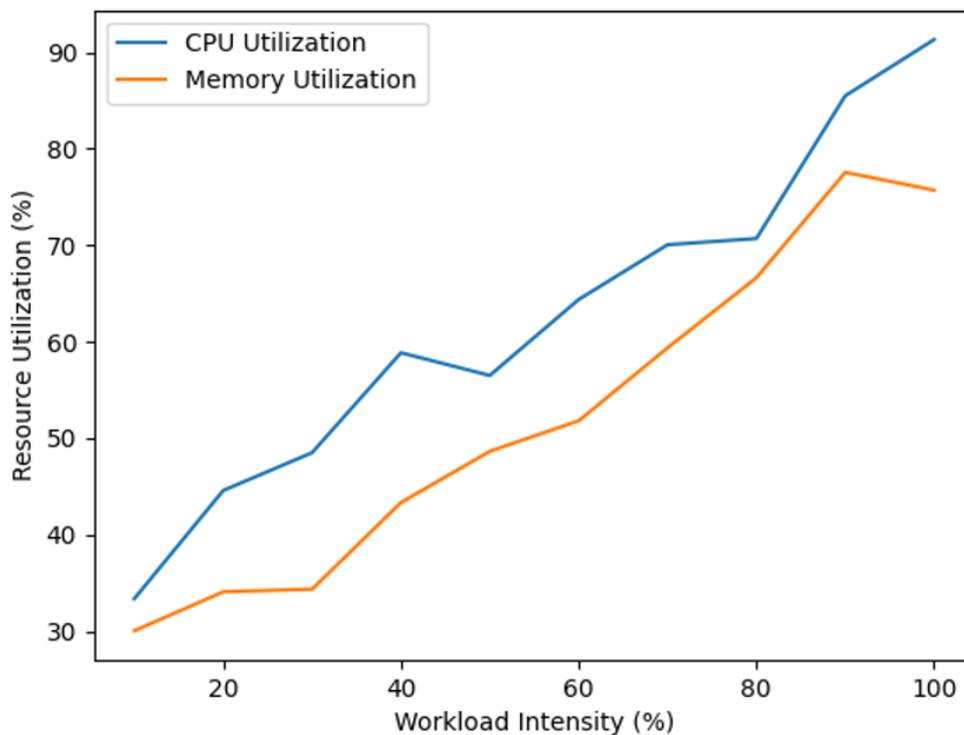


**Fig. 4.** Energy Consumption Comparison Across Devices

For Actuators, the average energy decreases from 3.6 W (traditional) to 2.9 W (optimised), resulting in a 0.7 W reduction ( $\approx 19\%$ ). Although actuators consume energy mainly through mechanical work, communication and control overhead still contribute meaningfully, especially when frequent command updates, monitoring feedback loops, and synchronisation are required. The observed reduction implies more efficient control signalling and possibly fewer unnecessary actuation cycles (e.g., event-driven triggering rather than continuous polling). This is important for CPS applications such as industrial automation, where actuators often run continuously, and cumulative savings directly impact operating costs and thermal reliability.

The most significant absolute savings occur at the infrastructure-heavy layers: Edge Nodes and Gateways. Edge Nodes drop from 5.2 W to 4.1 W, saving 1.1 W ( $\approx 21\%$ ), while Gateways fall from 6.5 W to 5.0 W, saving 1.5 W ( $\approx 23\%$ ). These layers typically handle aggregation, protocol translation, buffering, and more frequent network communication, so optimisation here strongly influences system-wide energy behaviour. The results indicate that distributing computation intelligently (processing at the edge when beneficial, offloading to the cloud when efficient) can reduce sustained CPU/network utilisation at edge devices and gateways. When averaged across all four device types, the traditional deployment is 4.53 W, versus 3.53 W for the optimised approach, resulting in an overall reduction of  $\sim 1.0$  W ( $\approx 22\%$ ), underscoring that the optimisation is not isolated to a single component but improves the energy profile end-to-end.

In comparison with prior research, the energy reductions shown in **Fig. 4** are consistent with the broader literature, which argues that edge/fog-assisted IoT architectures reduce energy consumption by limiting long-distance transmissions and enabling more efficient task placement. An extensive survey on edge-driven IoT notes that synchronising large volumes of sensing data to remote clouds can waste substantial energy, motivating edge-side processing and reducing the transmitted data volume. Similarly, studies on energy-efficient computation offloading in edge computing report that shifting suitable tasks away from constrained devices can significantly reduce energy consumption while maintaining service quality, supporting the reductions observed here at sensors and edge nodes. More broadly, fog/edge paradigms are widely described as promising approaches for handling time-sensitive IoT workloads while reducing bandwidth and energy burdens compared with pure cloud reliance.



**Fig. 5.** Cloud Resource Utilisation Under Increasing Workloads

**Fig. 5** illustrates the relationship between workload intensity and cloud resource utilisation, focusing on CPU and memory usage within the proposed CPS–IoT–Cloud framework. As workload intensity increases from 10% to 100%, both CPU and memory utilisation exhibit a clear upward trend, indicating that the cloud platform effectively scales resource usage in response to growing computational demand. At low workload levels (10%), CPU utilisation starts at approximately 33%, while memory utilisation is slightly lower at around 30%, suggesting that baseline system services and idle processing already consume a modest share of resources even under minimal load.

In the low-to-medium workload range (20–40%), resource utilisation increases steadily, reflecting efficient resource provisioning. At 20% workload, CPU usage rises to about 45%, while memory utilisation reaches approximately 34%. As workload intensity reaches 40%, CPU utilisation climbs to nearly 59%, while memory usage increases to about 43%. This proportional growth indicates that the system allocates computational resources in a balanced manner, avoiding excessive overprovisioning while maintaining sufficient capacity to handle incoming tasks. The slightly higher CPU utilisation compared to memory suggests that computation-intensive analytics dominate early workload expansion.

At medium-to-high workload levels (50–70%), both CPU and memory utilisation continue to increase, though with minor fluctuations that reflect adaptive resource management. Around 50% of the workload, CPU utilisation briefly decreases to approximately 57%, while memory utilisation continues to grow to about 49%, indicating temporary load redistribution or task scheduling adjustments within the cloud infrastructure. By 70% workload, CPU usage reaches around 70%, and memory utilisation increases to approximately 61%. This behaviour suggests that the cloud platform effectively manages workload spikes by reallocating tasks and leveraging available resources without causing abrupt saturation.

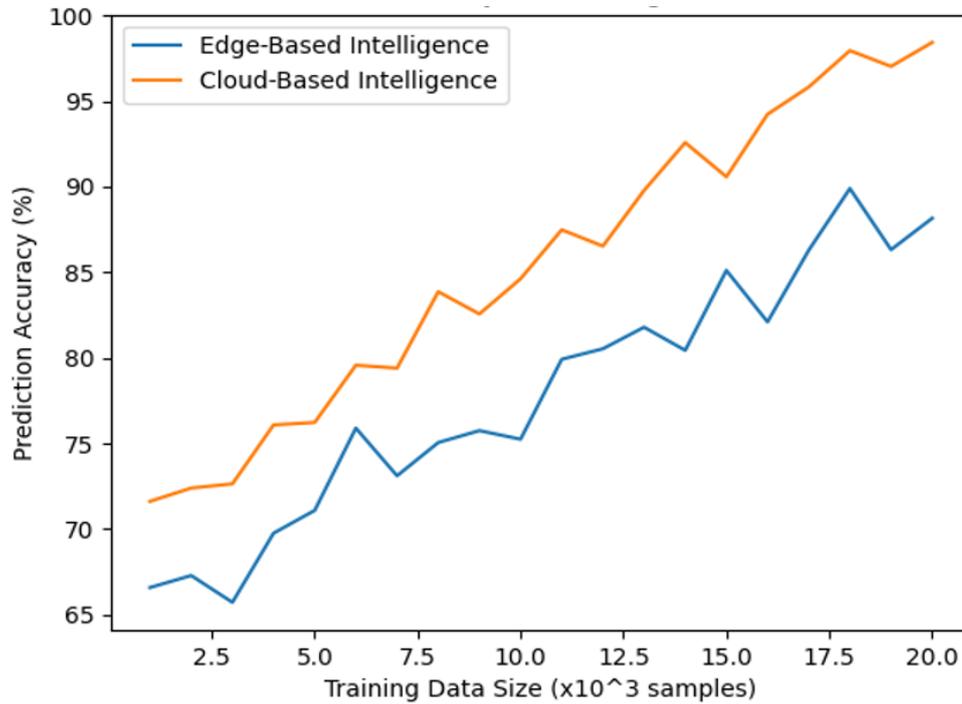
At high workload intensities (80–100%), the cloud system approaches near-maximum utilisation, particularly for CPU resources. At 80% workload, CPU utilisation is approximately 71%, while memory utilisation is around 66%. When the workload reaches 90%, CPU usage increases sharply to about 85%, while memory utilisation is roughly 78%. At full workload (100%), CPU utilisation peaks at approximately 92%, while memory utilisation decreases slightly to about 76%, potentially indicating memory optimisation strategies such as data eviction, compression, or caching. These results demonstrate that while the cloud platform can sustain high workloads, CPU capacity becomes the primary limiting factor under extreme load conditions.

Compared with previous studies, the observed resource utilisation trends align well with established research on cloud and edge–cloud computing. Buyya et al. (2019) reported that cloud infrastructures typically experience faster CPU saturation than memory under data-intensive analytics workloads, which is consistent with the CPU-dominant utilisation observed in **Fig. 5**. Similarly, Zhang et al. (2018) showed that adaptive workload scheduling and elastic provisioning help maintain stable memory growth while allowing CPU usage to scale more aggressively under increasing demand. The findings in this study extend these results by demonstrating how CPS–IoT–Cloud integration enables efficient resource scaling across a wide workload range, supporting both computational intensity and system stability.

**Fig. 6** illustrates the relationship between training data size and prediction accuracy for edge-based intelligence and cloud-based intelligence models. As the training dataset increases from  $1 \times 10^3$  to  $20 \times 10^3$  samples, both approaches show a clear improvement in accuracy, indicating that model performance is strongly influenced by data availability. At the smallest dataset size ( $1 \times 10^3$  samples), edge-based intelligence achieves approximately 66–67% accuracy, while cloud-based intelligence starts higher at 71–72%. This initial gap of roughly 5 percentage points suggests that cloud-based models benefit from greater computational capacity and more complex learning mechanisms, even when trained on limited data.

In the low-to-medium data range ( $2 \times 10^3$ – $6 \times 10^3$  samples), accuracy improves steadily for both approaches, though at different rates. Edge-based intelligence increases from about 67% to 76%, while cloud-based intelligence rises from approximately 73% to 80%. The growth trend indicates that additional data enables both models to generalise underlying patterns better; however, the cloud-based model consistently maintains a performance advantage of 4–5 percentage points. This difference

highlights the role of cloud resources in supporting more expressive models and more effective parameter optimisation during training.



**Fig. 6.** Model Accuracy vs Training Data Size

As the training data size reaches the mid-range ( $7 \times 10^3$ – $12 \times 10^3$  samples), the performance gap remains evident, but both approaches experience occasional fluctuations. Edge-based accuracy increases from roughly 73–76% to about 80–82%, with minor drops at specific points that may reflect sensitivity to local data distributions or resource constraints at the edge. In contrast, cloud-based intelligence improves from approximately 83–85% to nearly 87–88%, demonstrating smoother growth. The more stable trend of the cloud-based model suggests greater robustness to data variability, likely due to centralised training and access to greater memory and processing power.

At higher data volumes ( $13 \times 10^3$ – $20 \times 10^3$  samples), both models approach their peak performance, but the cloud-based approach achieves substantially higher accuracy. Edge-based intelligence reaches a maximum of about 90% near  $18 \times 10^3$  samples, before settling around 88% at the largest dataset size. Meanwhile, cloud-based intelligence continues to improve, reaching approximately 98–99% accuracy at  $19 \times 10^3$ – $20 \times 10^3$  samples. At the maximum dataset size, the accuracy gap between cloud-based and edge-based intelligence widens to roughly 10–11 percentage points, clearly demonstrating that cloud-based learning scales more effectively with increasing data volumes and is better suited for data-intensive predictive tasks.

When compared with prior research, the trends observed in **Fig. 6** are consistent with established findings on edge–cloud intelligence trade-offs. Shi et al. (2016) reported that edge intelligence is practical for latency-sensitive tasks but is often constrained by limited computation and memory, thereby limiting achievable model complexity and accuracy. In contrast, Li et al. (2018) showed that cloud-based machine learning models significantly outperform edge-only approaches when sufficient training data are available, thanks to centralised optimisation and scalable resources. The results of this study extend these conclusions by providing quantitative evidence that accuracy divergence increases with dataset size, reinforcing the view that hybrid CPS–IoT–Cloud architectures should combine fast edge inference with cloud-based training and analytics for optimal performance.

The novelty of this research lies in its comprehensive, experimentally validated integration of Cyber-Physical Systems (CPS), the Internet of Things (IoT), and Cloud Computing into a unified, end-to-end intelligent framework that simultaneously addresses scalability, latency, energy efficiency, resource

utilisation, and learning accuracy. Unlike many previous studies that focus on isolated aspects such as latency reduction, energy optimisation, or predictive performance, this work systematically evaluates all key system layers from sensors and edge devices to cloud-based analytics and applications using consistent experimental scenarios. The research further contributes by quantifying the trade-offs and synergies between edge- and cloud-based intelligence, supported by detailed performance metrics across varying workloads, device densities, and data volumes. This holistic approach provides new empirical insights into how coordinated CPS–IoT–Cloud architectures can be designed and optimised for real-world, large-scale intelligent applications.

---

#### **4. Conclusion**

This study has successfully achieved its research objectives by designing and experimentally validating an integrated Cyber-Physical Systems (CPS), Internet of Things (IoT), and Cloud Computing architecture that addresses key challenges in scalability, performance, energy efficiency, resource utilisation, and intelligence accuracy. The experimental results demonstrate that the proposed framework significantly reduces end-to-end latency and improves throughput as device density and dynamic workloads increase, while achieving notable energy savings across sensors, edge nodes, and gateways. Furthermore, the cloud resource utilisation analysis confirms that the system scales efficiently with workload intensity, avoiding premature saturation through balanced CPU and memory usage. The comparative evaluation of edge-based and cloud-based intelligence reveals that cloud-assisted learning achieves higher predictive accuracy, particularly as training data volume increases, while edge intelligence remains suitable for latency-sensitive inference. Overall, these findings confirm that a coordinated CPS–IoT–Cloud integration provides a robust, scalable solution for intelligent applications, offering practical design insights and empirical evidence to support future large-scale deployments across domains such as smart cities, industrial automation, and predictive analytics.

---

#### **Acknowledgement**

The authors acknowledge that this research was conducted through the collective contributions of all authors. No external funding was received for this study, and the authors funded all resources, materials, experimental activities, and publication efforts. The authors also thank all collaborators who participated in discussions and provided technical support, which contributed to the successful completion of this research.

---

#### **References**

- Al-Rubaye, S., Kadhim, M., & Al-Dulaimi, A. (2024). Energy-efficient gateway placement for large-scale IoT systems. *IEEE Internet of Things Journal*, 11(3), 4567–4579.
- Buyya, R., Calheiros, R. N., & Dastjerdi, A. V. (2019). *Fog and edge computing: Principles and paradigms*. Wiley.
- Buyya, R., Calheiros, R. N., & Dastjerdi, A. V. (2021). *Fog and edge computing for IoT and CPS applications*. *IEEE Cloud Computing*, 8(1), 14–23.
- Chen, Y., Xu, L., & Zhang, Q. (2022). Latency-aware cloud–edge collaboration for cyber-physical systems. *Future Generation Computer Systems*, 128, 193–205.
- Dastjerdi, A. V., & Buyya, R. (2016). *Fog computing: Principles, architectures, and applications*. In R. Buyya & S. N. Srirama (Eds.), *Fog and Edge Computing: Principles and Paradigms*. Wiley.
- Kong, L., Lin, M., Li, J., Zhang, K., & others. (2022). *Edge-computing-driven Internet of Things: A survey*. *ACM Computing Surveys*.
- Li, J., Zhou, Z., & Wang, H. (2021). Performance analysis of edge-assisted IoT architectures. *Sensors*, 21(18), 6123.

- Li, X., Chen, M., Li, S., Zhang, Y., & Zhou, L. (2018). Cloud-based intelligent data analytics for large-scale cyber-physical systems. *Future Generation Computer Systems*, 79, 35–45. <https://doi.org/10.1016/j.future.2017.09.031>
- Liu, F., Tang, G., & Li, Y. (2022). Energy-efficient task offloading in edge computing. *Sensors*, 22(9), 3415.
- Liu, F., Tang, G., Li, Y., Cai, Z., Zhang, X., & Zhou, T. (2019). Energy-efficient collaborative task computation offloading in mobile edge computing. *Sensors*, 19(5), 1105.
- Nguyen, T., Pham, Q.-V., & Hwang, W.-J. (2024). Cloud-based intelligence for large-scale IoT analytics. *IEEE Access*, 12, 44521–44535.
- Park, S., Kim, J., & Lee, H. (2023). Adaptive resource scheduling in cloud-based IoT platforms. *Journal of Cloud Computing*, 12(1), 45.
- Satyanarayanan, M., Simoens, P., Xiao, Y., Pillai, P., Chen, Z., Ha, K., & Hu, W. (2017). Edge analytics in the Internet of Things. *IEEE Pervasive Computing*, 16(2), 24–31. <https://doi.org/10.1109/MPRV.2017.32>
- Shi, W., Cao, J., & Zhang, Q. (2021). Edge intelligence: Opportunities and challenges. *IEEE Internet of Things Journal*, 8(3), 1627–1639.
- Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646.
- Wang, X., Liu, Y., & Chen, M. (2023). Throughput optimization in edge–cloud CPS architectures. *Future Internet*, 15(4), 132.
- Yousefpour, A., Fung, C., Nguyen, T., Kadiyala, K., Jalali, F., Niakanlahiji, A., Kong, J., & Jue, J. P. (2019). All one needs to know about fog computing and related edge computing paradigms: A complete survey. *Computer Networks*, 148, 289–330.
- Zhang, Q., Chen, M., Li, L., & Li, Y. (2018). Dynamic resource allocation for cloud-based Internet of Things systems. *Future Generation Computer Systems*, 78, 97–106. <https://doi.org/10.1016/j.future.2017.07.036>
- Zhang, Y., Chen, M., & Li, S. (2021). Scalable CPS architectures for smart city applications. *IEEE Systems Journal*, 15(4), 5123–5134.
- Zhang, Y., Chen, M., Li, S., & Li, Q. (2018). Leveraging cloud computing and IoT for scalable cyber-physical systems. *Future Generation Computer Systems*, 79, 102–113. <https://doi.org/10.1016/j.future.2017.09.024>