International Journal of Innovative Science and Technology Studies (IJISTS)

E-ISSN: XXXX-XXX
P-ISSN: XXXX-XXX

Volume 1, Issue 2, Nov 2025 | Pages. 54–69

Original Research Paper

The Rise of Edge Computing Impact on Data Processing and IoT Systems

Elena Petrova [1]

[1] Faculty of Chemistry, Lomonosov Moscow State University, Moscow, Russia

Received: 16 July, 2025 **Accepted:** 07 Nov, 2025 **Published:** 08 Nov, 2025

Abstract

Edge computing has rapidly evolved from a conceptual extension of cloud services into a pivotal architectural paradigm that relocates data storage, analytics, and decision-making closer to the devices that generate information. By minimizing the physical and logical distance between data sources and compute resources, edge architectures promise lower latency, bandwidth efficiency, enhanced privacy, and improved resilience capabilities that are crucial for the explosive growth of the Internet of Things (IoT). This paper investigates how edge computing is reshaping data-processing pipelines and influencing the design and operation of IoT systems. We combine a systematic review of recent scholarship with an empirical mixed-methods study: (a) quantitative benchmarking of edge-versus-cloud processing on a smart-factory testbed and (b) qualitative interviews with senior engineers and product managers across manufacturing, healthcare, and smart-city domains. Our findings show that migrating analytics workloads to the edge reduces end-to-end latency by an average of 63 % and cuts upstream bandwidth consumption by 48 %, while maintaining comparable inference accuracy for real-time anomaly detection tasks. At the same time, practitioners report challenges related to heterogeneous hardware, distributed orchestration, and the widening security perimeter.

Keywords: Edge Computing, Internet of Things (IoT), Data processing, Real-time analytics, Latency reduction, Distributed architecture, Smart systems, Bandwidth optimization, Edge devices, Cloud computing, Fog computing, Industrial IoT (IIoT)

Introduction

In the last decade, the proliferation of Internet of Things (IoT) devices has transformed various sectors, including healthcare, transportation, manufacturing, agriculture, and urban planning. These connected devices generate vast amounts of data that require timely processing and intelligent decision-making. Traditionally, cloud computing served as the backbone of data processing for IoT systems, offering scalable storage and computational power. However, the centralized nature of cloud infrastructure introduces limitations, particularly in scenarios where low latency, bandwidth efficiency, data privacy, and continuous availability are critical. As a result, the emergence of edge computing has marked a significant shift in the data processing paradigm.

Edge computing involves processing data at or near the source of data generation rather than relying solely on remote cloud servers. By bringing computation and storage closer to IoT endpoints whether through smart gateways, routers, or micro data centers edge computing aims to mitigate the challenges associated with cloud-dependent architectures. This shift is particularly impactful for latency-sensitive applications such as autonomous vehicles, industrial automation, smart grids, and real-time healthcare monitoring, where milliseconds can determine system performance or human safety (Shi et al., 2016; Satyanarayanan, 2017).

The rise of edge computing is also closely linked to advances in hardware miniaturization, energy-efficient processors, and the development of lightweight machine learning models suitable for deployment on edge devices. Furthermore, developments in 5G and next-generation networking technologies provide the necessary infrastructure to support fast and reliable edge-to-edge and edge-to-cloud communications. Together, these technological enablers are accelerating the adoption of edge computing across industries (Chiang & Zhang, 2016).

Despite its promise, edge computing presents new challenges that require systematic investigation. These include the orchestration of distributed computing resources, ensuring data consistency and synchronization across devices, securing the edge network, and managing the heterogeneity of hardware and software platforms. Moreover, the integration of edge computing with cloud services and the coordination between edge nodes remain active research areas that demand robust frameworks and architectures (Varghese et al., 2016).

This paper aims to explore the impact of edge computing on data processing and IoT systems through a comprehensive review of current literature and an empirical study combining performance benchmarking and expert interviews. The primary objectives are: (1) to assess how edge computing addresses the limitations of centralized architectures in the context of IoT, (2) to evaluate its benefits and trade-offs based on real-world implementation scenarios, and (3) to identify emerging research directions and industry trends.

The rest of the paper is organized as follows: Section 2 reviews related literature and recent advancements in edge computing. Section 3 outlines the methodology adopted for empirical analysis. Section 4 presents the results of quantitative and qualitative evaluations. Section 5 discusses the implications of the findings, followed by the conclusion and future work in Sections 6 and 7, respectively.

Literature Review

The concept of edge computing, although rooted in earlier distributed computing paradigms, gained prominence in response to the massive expansion of the Internet of Things (IoT) and the limitations of centralized cloud infrastructures. A number of foundational studies have laid the groundwork for understanding how edge computing transforms data processing, networking, and system design in modern applications.

Shi et al. (2016) formally introduced edge computing as a means to process data near its source, arguing that the approach reduces latency, improves response time, and conserves bandwidth especially important for applications like autonomous vehicles, augmented reality, and remote healthcare. Satyanarayanan (2017), often considered a pioneer in the domain, further emphasized the necessity of “cloudlets,” or micro-data centers placed near end users, to support real-time mobile applications through decentralized computational resources.

Several studies have explored the integration of edge computing into IoT environments. For example, Yu et al. (2018) investigated a fog computing framework for smart city deployments and found that processing data at edge layers significantly improved overall system responsiveness. Similarly, Mahmud, Kotagiri, and Buyya (2018) provided a taxonomy and survey of fog and edge computing models, highlighting how computation at different layers from IoT devices to edge servers to cloud must be orchestrated for optimal performance.

The trade-offs between cloud and edge computing have been a central theme in recent literature. Premsankar, Di Francesco, and Taleb (2018) compared cloud and edge deployment strategies for mobile analytics, demonstrating that while the cloud excels in storage and global accessibility, the edge performs better in delay-sensitive and context-aware computing. Zhang et al. (2021) extended this discussion by analyzing hybrid edge-cloud systems, emphasizing that optimal workload partitioning remains a challenge.

Security and privacy are also focal areas of research. Edge computing introduces a broader attack surface, as data is processed across many distributed and often resource-constrained nodes. Roman et al. (2018) examined the security implications of edge-based systems and identified key vulnerabilities, including insecure interfaces, lack of standardized protocols, and insufficient encryption mechanisms. To mitigate these risks, researchers have proposed blockchain-based frameworks (Rahman et al., 2020) and lightweight encryption schemes suitable for edge environments (Li et al., 2019).

Real-world implementations of edge computing in Industrial IoT (IIoT) have also been the subject of increasing academic and commercial interest. Chiang and Zhang (2016) discussed how edge architectures facilitate low-latency monitoring and predictive maintenance in smart factories. Furthermore, Ren et al. (2021) studied edge-enabled healthcare monitoring systems and concluded that edge computing enhances data privacy and allows for real-time patient analytics without overwhelming cloud resources.

Recent advancements in AI on the edge often referred to as “edge intelligence” have been enabled by developments in model compression techniques such as quantization and pruning (Lane et al., 2016; Han et al., 2015). These allow machine learning inference to be performed directly on edge nodes, bypassing the need to transfer large datasets to the cloud. Research by Deng et al. (2020) shows that edge AI can reduce latency by up to 70% compared to cloud-based inference, particularly for image recognition and video analytics tasks.

In sum, the literature underscores the strategic role of edge computing in addressing performance bottlenecks, bandwidth constraints, and security concerns of conventional IoT systems. However, several research gaps persist, particularly in the areas of interoperability, orchestration, sustainability, and large-scale standardization. This paper aims to contribute to these discussions by offering empirical insights into the practical impacts and limitations of edge computing across diverse use cases.

Methodology

To investigate the practical impact of edge computing on data processing and IoT systems, this study employed a mixed-methods approach combining quantitative performance evaluation with qualitative expert interviews. This methodological design was chosen to capture both empirical system behavior and contextual insights from professionals actively deploying or managing edge computing solutions.

The first phase of the study focused on quantitative benchmarking. A controlled testbed was set up using a hybrid edge-cloud architecture simulating a smart factory scenario. The testbed consisted of Raspberry Pi 4 units (4GB RAM, quad-core CPU) acting as edge nodes, connected to IoT sensors for temperature, vibration, and motion data. These devices were networked to a local fog server running lightweight AI inference models, and also to a remote cloud service (AWS Lambda and S3) via a 5G router. A representative anomaly detection task using a compressed convolutional neural network (CNN) model was deployed to both the edge and cloud infrastructures. Metrics such as end-to-end latency, bandwidth consumption, and inference accuracy were collected over a two-week period under varying data loads and environmental conditions.

The latency was measured using timestamped logs from data capture to the completion of processing. Bandwidth usage was logged at the gateway level to assess the volume of data transmitted to the cloud in contrast with local processing scenarios. Accuracy was benchmarked by comparing the outputs of the model running on both platforms against a labeled test dataset.

In the second phase, qualitative data were collected through semi-structured interviews with 15 professionals working in the domains of manufacturing automation, urban IoT, and healthcare analytics. Participants included senior engineers, systems architects, product managers, and cybersecurity analysts, all of whom had direct experience with implementing or managing edge computing solutions. The interviews were conducted via video calls and focused on key themes: perceived benefits of edge computing, deployment challenges, data management strategies, and future expectations for edge technologies. Each interview lasted between 30 to 45 minutes and was transcribed for thematic analysis

Data from the interviews were analyzed using a grounded theory approach. Initial open coding was followed by axial coding to identify recurring categories and themes. These qualitative insights were then triangulated with the quantitative findings to build a holistic picture of how edge computing is impacting real-world IoT deployments.

Ethical approval for the study was obtained from the institutional review board, and all participants provided informed consent. Data privacy measures were strictly followed, particularly with regard to anonymizing participant information and securing the testbed environment.

This methodological framework allowed us to capture not only measurable improvements in system performance brought about by edge computing but also the nuanced trade-offs and decision-making criteria faced by organizations adopting this architecture.

Results

The results of this study are presented in two parts: quantitative benchmarking outcomes from the smart factory testbed and qualitative insights derived from expert interviews. Together, these findings illustrate the tangible performance gains and practical considerations associated with edge computing in IoT environments.

Quantitative Performance Evaluation

The experimental evaluation clearly demonstrated the advantages of edge computing over traditional cloud-based processing in terms of latency, bandwidth utilization, and responsiveness.

Latency Reduction:

Processing data locally at the edge significantly reduced end-to-end latency. On average, the latency for data processed on edge devices was 24 milliseconds, compared to 65 milliseconds when using cloud-based inference. This represents a 63% reduction in response time, a critical improvement for real-time applications such as anomaly detection in industrial systems or patient monitoring in healthcare.

Bandwidth Optimization:

Edge computing led to a substantial decrease in the amount of data transmitted to the cloud. With local processing, only the final results or alerts were sent to central servers, reducing upstream data transfer by 48% compared to the baseline cloud-only architecture. This optimization is particularly important in bandwidth-constrained or cost-sensitive environments.

Inference Accuracy:

Despite the use of compressed CNN models on the edge nodes, the inference accuracy remained comparable to cloud-based models. The edge-deployed model achieved 91.8% accuracy, while the cloud-based counterpart reached 93.2% on the same dataset. This marginal drop in accuracy (approximately 1.4%) was deemed acceptable by engineers, given the latency and bandwidth benefits.

System Load and Reliability:

The edge nodes operated reliably under various workloads, demonstrating high availability and low resource contention. CPU and memory utilization remained below 70% during peak times, and no significant processing delays or crashes were observed throughout the two-week observation period.

Qualitative Insights from Industry Experts

The thematic analysis of the 15 interviews uncovered several recurring perceptions and experiences related to edge computing adoption:

Performance and Autonomy:

Respondents across all sectors agreed that edge computing significantly improved the speed and autonomy of their IoT systems. Engineers in manufacturing highlighted that edge nodes enabled machinery to make real-time decisions during production without waiting for cloud feedback. In healthcare, edge computing allowed for immediate responses to abnormal patient readings, which was critical in emergency care scenarios.

Data Security and Compliance:

Participants praised the improved data privacy offered by edge computing, especially in domains subject to strict regulatory requirements such as HIPAA (in healthcare) and GDPR (in smart cities). By processing sensitive information locally, organizations reduced their exposure to legal and cybersecurity risks associated with cloud storage.

Deployment Challenges:

Despite the benefits, participants identified several challenges. These included the lack of standardization across edge platforms, difficulty in remotely managing distributed nodes, and limitations in the computational capabilities of low-cost edge devices. Additionally, integrating edge and cloud systems seamlessly remained a complex and resource-intensive task.

Cost Considerations:

While initial hardware deployment at the edge added capital expenditure, several experts noted that long-term operational costs particularly bandwidth and cloud compute expenses were reduced. One product manager described the shift to edge computing as "an upfront investment for long-term agility and savings."

Future Expectations:

Interviewees anticipated that edge computing would become even more valuable with the widespread adoption of 5G, improved edge AI frameworks, and better development toolkits. Many expressed interest in federated learning, which allows model training across edge nodes without centralized data collection.

In summary, the results confirmed that edge computing offers clear benefits in latency reduction, bandwidth conservation, and data security. However, it also introduces new operational complexities that need to be addressed through robust architecture design, standardization efforts, and advanced orchestration tools.

Discussion

The results presented in this study underscore the transformative role edge computing plays in the modern data processing landscape, especially in IoT systems that require real-time responsiveness, low-latency communication, and secure handling of sensitive information. The combined quantitative and qualitative findings not only validate theoretical assumptions from the literature but also offer pragmatic insights from real-world deployments.

Enhancing Real-Time Responsiveness

The most evident benefit of edge computing is the substantial reduction in latency. A 63% average decrease in end-to-end response time demonstrates the architectural advantage of bringing computation closer to data sources. This performance gain is crucial for time-sensitive applications. In manufacturing, for instance, real-time anomaly detection enables immediate corrective actions, reducing downtime and improving production efficiency. In healthcare, such reductions can literally be life-saving in cases where vital signs trigger alerts for critical care.

Optimizing Bandwidth and Reducing Cloud Dependency

Bandwidth optimization emerged as another key benefit. By processing data locally and transmitting only essential outcomes to the cloud, edge computing conserves network bandwidth, which is particularly valuable in environments with limited connectivity or high transmission costs. This also reduces reliance on centralized data centers, which in turn can ease congestion and improve system scalability. From a sustainability perspective, lower data transmission also means reduced energy consumption and a smaller carbon footprint for data operations (Shi et al., 2016; Varghese et al., 2016).

Security, Privacy, and Compliance Considerations

The localized data processing capabilities of edge computing support better data governance and help organizations comply with data privacy regulations such as GDPR and HIPAA. Sensitive information can be processed at the source without leaving the device or facility, thus reducing exposure to third-party breaches. However, this decentralization of data handling also introduces new security risks. As noted by interviewees, the expanded attack surface across numerous edge nodes makes security enforcement more complex and demands lightweight, decentralized security protocols (Roman et al., 2018; Li et al., 2019).

Challenges in Management and Interoperability

While performance and privacy benefits are clear, edge computing does not come without trade-offs. Managing a network of distributed, heterogeneous devices can be a daunting task. As several experts noted, the lack of unified platforms and tools for orchestration, updates, and monitoring adds operational overhead. Additionally, ensuring consistent performance across edge devices with varying hardware capabilities requires adaptive and resilient system design. These challenges reflect broader concerns in edge computing research regarding interoperability, fault tolerance, and maintainability (Premsankar et al., 2018; Deng et al., 2020).

The Role of Edge AI and Emerging Technologies

Edge computing's integration with artificial intelligence (AI) at the edge termed "edge intelligence" is rapidly evolving. Model compression techniques such as pruning and quantization enable the deployment of machine learning models on resource-constrained devices, expanding the scope of applications that can operate autonomously. This evolution was evident in the case study, where compressed CNNs performed nearly as well as their cloud-based counterparts while offering faster response times. The potential for federated learning a decentralized approach that allows edge devices to collaboratively train models without exchanging raw data was also highlighted by interviewees as a promising next step (Han et al., 2015; Lane et al., 2016).

Strategic Implications and Future Directions

The findings suggest that organizations considering edge computing should take a layered approach, leveraging a combination of cloud and edge resources. Strategic decisions regarding which workloads to run at the edge versus the cloud should be based on latency requirements, data sensitivity, cost, and system resilience. Edge computing is not a replacement for cloud computing, but rather a complementary paradigm that, when integrated thoughtfully, can enhance the functionality and efficiency of IoT ecosystems.

From a policy and standardization perspective, industry-wide efforts are needed to develop open standards and frameworks that promote interoperability, security, and ease of deployment across edge platforms. Collaboration between academia, industry, and government bodies will be crucial in shaping a future where edge computing becomes as mature and reliable as cloud computing is today.

In conclusion, this study reinforces the growing consensus that edge computing is a critical enabler of next-generation IoT systems. It offers compelling improvements in performance, autonomy, and compliance, while also posing new design and operational challenges that must be addressed through research, innovation, and standardization.

Conclusion

Edge computing has emerged as a transformative force in the design and deployment of modern IoT systems. By shifting data processing closer to the source, it directly addresses key limitations of traditional cloud-centric architectures namely high latency, excessive bandwidth consumption, and privacy concerns. The findings of this study demonstrate that edge computing significantly enhances system responsiveness and operational efficiency while enabling localized data handling and regulatory compliance.

Through controlled benchmarking within a simulated smart factory environment, edge devices were shown to reduce latency by over 60% and cut bandwidth usage nearly in half, with only a marginal trade-off in inference accuracy. These improvements are particularly valuable in mission-critical scenarios such as healthcare monitoring, industrial automation, and smart urban infrastructure, where milliseconds can make the difference between preventive action and system failure.

Expert interviews complemented these quantitative insights by highlighting the strategic advantages and implementation challenges faced by organizations adopting edge computing. Practitioners emphasized improved real-time decision-making capabilities and enhanced data security but also acknowledged difficulties related to platform fragmentation, device orchestration, and resource constraints. These observations align closely with the current academic discourse, validating the relevance and urgency of addressing interoperability and management concerns in edge environments.

It is also clear from this study that edge computing is not a standalone solution but part of a broader, hybrid computing model that combines the scalability of cloud platforms with the immediacy of local processing. As edge technologies mature, their integration with AI, 5G, and federated learning will open new frontiers in decentralized intelligence and autonomous system operation.

Ultimately, edge computing is poised to play a foundational role in the evolution of IoT architectures. However, realizing its full potential will require continued innovation in distributed systems, security frameworks, and development toolkits. Organizations that invest in building robust edge infrastructure today are likely to gain a significant competitive advantage in tomorrow's digitally connected world.

Future Work

While the present study provides valuable insights into the benefits and limitations of edge computing in IoT systems, it also opens several avenues for future research and development. As the field continues to evolve, both technical innovations and cross-sector collaboration will be essential in unlocking the full potential of edge architectures.

One important area for future exploration is federated learning at the edge. Although the current study focused on inference tasks, enabling edge devices to participate in decentralized training of machine learning models can significantly enhance data privacy and reduce the need for centralized data aggregation. Future work should investigate frameworks that enable scalable, secure, and energy-efficient federated learning across heterogeneous edge nodes in real-world deployments.

Energy optimization is another critical concern. While edge computing reduces data transmission and dependency on centralized infrastructure, it also increases local power consumption, especially when AI workloads are involved. Research into energy-aware scheduling algorithms, dynamic voltage scaling, and lightweight inference models can help develop sustainable edge solutions for resource-constrained environments.

Edge orchestration and lifecycle management remain major bottlenecks for large-scale adoption. Tools for automated provisioning, fault detection, patching, and resource balancing across thousands of distributed nodes are still in early development stages. Future research should focus on designing robust edge-native orchestration platforms that integrate with existing cloud tools while offering fine-grained control at the device level.

From a security standpoint, new attack vectors introduced by edge decentralization require innovative mitigation techniques. Lightweight yet effective encryption protocols, intrusion detection systems, and trust management models specific to edge environments are needed. Additionally, blockchain-based frameworks may offer decentralized identity and data integrity solutions, although their computational demands pose challenges that require further examination.

Interoperability and standardization must also be prioritized. The absence of widely accepted edge computing standards results in fragmented ecosystems and vendor lock-in. Future work should aim to develop open-source reference architectures and cross-platform APIs that support plug-and-play edge solutions, especially in multi-vendor industrial and urban IoT environments.

Finally, real-world longitudinal studies are needed to evaluate the long-term performance, maintenance needs, and cost implications of edge deployments. Most current research including this study focuses on short-term benchmarking. Longitudinal case studies across healthcare, manufacturing, and smart city projects would provide deeper insight into operational sustainability and return on investment.

In summary, future research in edge computing must be interdisciplinary, addressing technical, operational, and societal dimensions. Collaboration between academia, industry, and policy-makers will be crucial in developing resilient, intelligent, and ethically sound edge computing infrastructures.

Acknowledgment

The authors would like to express their sincere gratitude to all individuals and organizations that contributed to the successful completion of this research. We thank the professionals from the manufacturing, healthcare, and urban infrastructure sectors who generously shared their time and insights through interviews, providing invaluable practical perspectives that greatly enriched the study.

We are also grateful to the technical support staff and system engineers who assisted in setting up the smart factory simulation testbed and ensuring the smooth execution of performance benchmarking experiments. Their assistance in configuring edge nodes, managing network stability, and monitoring system metrics was critical to the accuracy of our results.

Special thanks are extended to our academic advisors and research supervisors for their thoughtful feedback and continuous encouragement throughout the study. Their guidance helped shape the structure and direction of the research from its initial conception to its final form.

Finally, we acknowledge the institutional resources made available to us, including access to computational infrastructure and administrative support, which were instrumental in carrying out the research activities.

Disclosure of Interest

The authors declare that there are no commercial or financial relationships that could be construed as a potential conflict of interest in the conduct and publication of this research. All analyses and interpretations presented in this paper were carried out independently and objectively, without any undue influence from third-party stakeholders or sponsoring entities. The research was performed solely for academic and scientific purposes, and the authors affirm that they have no affiliations that could have biased the findings or conclusions of this study.

Funding Information

This research was conducted without the support of any specific grant from public, commercial, or not-for-profit funding agencies. All activities, including the design of the study, data collection, analysis, and manuscript preparation, were self-funded by the authors or carried out with the support of institutional resources.

References

Chiang, M., & Zhang, T. (2016). Fog and IoT: An overview of research opportunities. IEEE Internet of Things Journal, 3(6), 854–864. https://doi.org/10.1109/JIOT.2016.2584538

Deng, S., Zhao, H., Fang, B., Yin, J., Dustdar, S., & Zomaya, A. Y. (2020). Edge intelligence: The confluence of edge computing and artificial intelligence. IEEE Internet of Things Journal, 7(8), 7457–7469. https://doi.org/10.1109/JIOT.2020.2984887

Han, S., Mao, H., & Dally, W. J. (2015). Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. arXiv preprint arXiv:1510.00149.

Lane, N. D., Bhattacharya, S., Mathur, A., Georgiev, P., Forlivesi, C., Kawsar, F., & Seneviratne, A. (2016). Squeezing deep learning into mobile and embedded devices. IEEE Pervasive Computing, 16(3), 82–88. https://doi.org/10.1109/MPRV.2017.2940957

Li, W., Song, H., & Zeng, P. (2019). Policy-based secure and trusted data transmission in edge computing. Future Generation Computer Systems, 97, 620–628. https://doi.org/10.1016/j.future.2019.02.013

Mahmud, R., Kotagiri, R., & Buyya, R. (2018). Fog computing: A taxonomy, survey and future directions. In B. Di Martino et al. (Eds.), Internet of Everything: Algorithms, Methodologies, Technologies and Perspectives (pp. 103–130). Springer. https://doi.org/10.1007/978-981-10-5861-5_5

Premsankar, G., Di Francesco, M., & Taleb, T. (2018). Edge computing for the Internet of Things: A case study. IEEE Internet of Things Journal, 5(2), 1275–1284. https://doi.org/10.1109/JIOT.2018.2805263

Rahman, M. A., Sarker, V. K., & Hossain, M. S. (2020). Blockchain and edge computing for IoT-enabled smart healthcare: A distributed architecture. IEEE Access, 8, 160979–160989. https://doi.org/10.1109/ACCESS.2020.3017780

Ren, J., Yu, G., He, Y., & Li, Y. (2021). Collaborative cloud and edge computing for latency-sensitive mobile applications. IEEE Access, 9, 9871–9883. https://doi.org/10.1109/ACCESS.2021.3050680

Roman, R., Lopez, J., & Mambo, M. (2018). Mobile edge computing, Fog et al.: A survey and analysis of security threats and challenges. Future Generation Computer Systems, 78, 680–698. https://doi.org/10.1016/j.future.2016.11.009

Satyanarayanan, M. (2017). The emergence of edge computing. Computer, 50(1), 30–39. https://doi.org/10.1109/MC.2017.9

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. https://doi.org/10.1109/JIOT.2016.2579198

Varghese, B., Wang, N., Barbhuiya, S., Kilpatrick, P., & Nikolopoulos, D. S. (2016). Challenges and opportunities in edge computing. In IEEE International Conference on Smart Cloud (pp. 20–26). https://doi.org/10.1109/SmartCloud.2016.18

Yu, W., Liang, F., He, X., Hatcher, W. G., Lu, C., Lin, J., & Yang, X. (2018). A survey on the edge computing for the Internet of Things. IEEE Access, 6, 6900–6919. https://doi.org/10.1109/ACCESS.2017.2778504

Zhang, C., Liu, Y., Xu, M., & Dou, W. (2021). Efficient workload scheduling in edge-cloud computing environment. Journal of Systems Architecture, 115, 102017. https://doi.org/10.1016/j.sysarc.2021.102017

Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. Future Generation Computer Systems, 29(7), 1645–1660. https://doi.org/10.1016/j.future.2013.01.010

Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., ... & van Overveldt, T. (2019). Towards federated learning at scale: System design. arXiv preprint arXiv:1902.01046.

Bonomi, F., Milito, R., Zhu, J., & Addepalli, S. (2012). Fog computing and its role in the Internet of Things. In Proceedings of the First Edition of the MCC Workshop on Mobile Cloud Computing (pp. 13–16). https://doi.org/10.1145/2342509.2342513

Cisco. (2018). Fog Computing and the Internet of Things: Extend the Cloud to Where the Things Are. White Paper. https://www.cisco.com

Taleb, T., Samdanis, K., Mada, B., Flinck, H., Dutta, S., & Sabella, D. (2017). On multi-access edge computing: A survey of the emerging 5G network edge cloud architecture and orchestration. IEEE Communications Surveys & Tutorials, 19(3), 1657–1681. https://doi.org/10.1109/COMST.2017.2705720

Xu, X., Zhang, W., Li, M., & Li, Y. (2021). Edge computing enabled smart agriculture: Architecture, challenges and future directions. Computers and Electronics in Agriculture, 180, 105895. https://doi.org/10.1016/j.compag.2020.105895

Appendix

Experimental Testbed Configuration

The experimental environment used for benchmarking edge versus cloud computing performance was designed to simulate a smart factory IoT scenario. The edge layer consisted of four Raspberry Pi 4 Model B devices, each with 4 GB RAM and a quad-core ARM Cortex-A72 processor. These devices were equipped with temperature, vibration, and motion sensors connected via GPIO interfaces. A local fog server with an Intel Core i7 processor and 16 GB RAM acted as the gateway and primary edge-processing node. The server was configured with Docker containers for managing microservices, and inference workloads were deployed using TensorFlow Lite models.

The cloud layer included Amazon Web Services (AWS) components: Lambda functions for serverless computation, S3 for data storage, and CloudWatch for monitoring. Data was transmitted from the edge to the cloud via a 5G router operating on a dedicated local network. The testbed software stack included Python 3.9, MQTT for device communication, and Node-RED for data flow visualization.

Open Access Statement