Edge Vs Cloud Computing Performance Trade-Offs for Real-Time Analytics

Aniket Abhishek Soni
Independent Researcher
Senior IEEE Member
Brooklyn, New York, USA

Rashi Nimesh Kumar Dhenia
Independent Researcher
IEEE Member
Atlanta, Georgia USA

Milan Parikh
Independent Researcher
Senior IEEE Member
Richmond, Houston, USA

Abstract: As more industries require real-time data processing, selecting the optimal computing architecture has become essential. This paper explores how edge analytic tasks compare to cloud computing regarding value and speed. While cloud computing provides significant scalability and a centralized resource pool, it can struggle with applications that require quick responses. Conversely, edge computing maintains processing close to data collectors, reducing latency, though this introduces challenges in scaling and system maintenance. The article presents key performance statistics for latency, bandwidth, processing power, and security, illustrating when and where each approach is most effective. The goal is to assist leaders, IT staff, and developers in identifying the best architecture for their real-time analytics tasks.

Keywords: Edge computing, Cloud computing, Real-time analytics, Latency, Bandwidth, Performance trade-offs, Data processing, IoT, Hybrid architecture, Computing infrastructure

1. INTRODUCTION

Because we are generating data at unprecedented rates, there is greater demand than ever for systems that can analyze this data live. Real-time analytics allows both businesses and institutions to respond to insights straight away such as when traffic signals update automatically or machines catch equipment faults early on. How effective these systems are often comes down to how and where the data is being processed, so edge computing and cloud computing become important.

Modern data infrastructure is built largely on the structure of cloud computing. The system is crafted in a central way which makes it highly powerful, scalable when needed, and integrable with various analytical tools, including AI and ML. Economies of scale managed services and high availability are available to companies using Amazon Web Services, Microsoft Azure, and Google Cloud, according to Zhang et al. (2010). Yet, using the cloud has some limitations. Delays in network communication, called latency, are a critical issue when data from many parts of a network goes far to central servers, as is the case in autonomous vehicles, remote healthcare monitoring, and high-frequency trading.

That's why edge computing is becoming an attractive alternative. Processes used to be sent to far-off centers, but now, data is handled where it's generated in edge computing. The new approach greatly shortens delays and uses less bandwidth since less data is sent to the cloud (Shi & Dustdar, 2016). In such cases, a robot with sensors onboard can make decisions instantly, without consulting the cloud. Even so, edge computing makes applications faster but carries risks: edge computing equipment generally has less powerful processors, and smaller memory and doesn't always offer the same degree of security as cloud systems.

A key question during the design of real-time analytics systems is whether all data should be processed locally, at a distance, or by using a mix of both. What is selected will vary depending on issues such as delay, the amount of data sent, connection strength, and whether a system needs a central brain or can make decisions on its own.

Knowing these decisions matters a lot to IT architects, developers, and decision-makers. If we think about automated driving or emergency response, where every decision has to be fast, edge computing is favored even for modest delays. However processing big datasets with high computation and long storage must often be done through cloud computing, whether for predicting financial trends or analyzing customer habits in shops.

Table 1: Performance Trade-offs between Edge and Cloud Computing in Real-Time Analytics

Metric	Edge Computing	Cloud Computing
Latency	Very low; data is processed near the source, enabling instant decision-making.	Higher latency; data must travel to and from distant data centers
Bandwidth Usage	Efficient; only essential data is sent to the cloud (if at all)	High; large volumes of raw data are often transmitted continuously
Processing Power	Limited; constrained by the hardware of edge devices	Extensive; utilizes powerful, scalable infrastructure in cloud data centers
Scalability	Moderate; adding edge nodes requires physical deployment and coordination	High; virtual machines and services can be scaled up or down easily
Security	Data stays local, reducing exposure; however, device-level protection is essential.	Centralized tools and protocols; greater attack surface due to aggregation
Reliability	Varies; dependent on local device performance and	Generally high; backed by SLAs and fault-tolerant architectures

	network reliability	
Cost Efficiency	Suitable for time-sensitive and lightweight workloads	Cost-effective for processing and storing large-scale, aggregated data

It's not always true that businesses can only use one approach: edge or cloud. A lot of organizations are choosing ways that use the best elements of both systems. In these designs, a quick choice is passed to the edge, but long-term processing and close study take place in the cloud. As an illustration, a retail chain can count the number of customers in each of their stores with edge computing at any time but produces weekly sales projections using cloud platforms.

In this article, I explore more deeply how edge and cloud computing match up for real-time analytics. It does this by exploring practical examples and relating them to the required system prices to form a clear guide to picking the best architecture for a certain application.

2. LITERATURE REVIEW

Answering the daily need for quick results by using data has raised the popularity of both edge and cloud computing. The advantages and difficulties of each model become most clear when looking at how they fulfill the performance requirements of real-time analytics. A large number of articles have evaluated the special attributes and compromises between them.

Cloud computing has for some time been thought of as the main way to manage data-intensive applications because it provides resources centrally, can handle ongoing changes in demand, and offers a variety of integrated data services. Cloud computing allows companies to deploy elaborate analytics systems without having to buy costly servers and facilities, according to Zhang et al. (2010). When real-time systems like autonomous cars, health systems, and IoT for industries started being used, researchers realized that sending data over the internet to big servers caused significant delays.

In response to this issue, edge computing was invented to put the processing needed close to the data source. According to Shi and Dustdar, edge computing lowers delays and helps send less data over networks with limited capacity. That's why it works well in vital places with little tolerable delay, like factories, vehicles that drive themselves, and augmented reality. At the same time, they highlight that it may be difficult for edge architectures to achieve top performance and above all stay consistent at the same time across a distributed network of edge devices.

Researchers have completed several case studies to determine how the models measure up in performance. In their 2018 paper, Premsankar, Di Francesco, and Taleb compared edge and cloud in smart cities and said that edge systems can respond more quickly, but they may be insufficient for implementing complicated machine learning models that the cloud is better at hosting. Goudarzi et al. (2020) looked at both the costs and energy use in these systems and found that having edge devices process live tasks and clouds process gathered data results in the most efficient compromise.

Security also plays a big role in how we make decisions about performance. Because cloud services manage security centrally, they are safer, but because many sensitive data are gathered in one place it becomes a bigger risk (Hashizume et

al., 2013). Yet, in edge computing, data is handled at the edge where it is collected, so there's less exposure, except that individual nodes might not be well protected, making them more vulnerable (Roman, Lopez, & Mambo, 2018).

New research shows that combining edge and cloud processing is rising in importance. In their work, Abbas et al. (2018) think of edge and cloud computing as working together, rather than competing with each other. Key tasks are finished on edge devices in this model, with important and demanding processing done in the cloud. With dynamic partitioning, the approach sees and uses the advantages of both methods while avoiding the disadvantages.

All in all, the literature gives a detailed picture of how the edge and cloud compare, mainly in real-time analytics. Even though cloud computing gives you plenty of resources, edge computing is best at keeping information close and fast. The majority of experts think that hybrid systems offer the best way to handle real-time tasks by keeping the app fast, supportive of a growing population, and cost-effective.

Key Themes from Literature:

Theme	Edge Computing	Cloud Computing	References
Latency	Low; close to the data source	Higher; due to transmission delay	Shi & Dustdar (2016); Satyanarayanan (2017)
Processing Capability	Limited; good for lightweight tasks	High; good for intensive computing	Zhang et al. (2010); Premsankar et al. (2018)
Scalability	Physical deployment needed	Virtual scaling is flexible	Goudarzi et al. (2020)
Security and Privacy	Decentralized; device-level security required	Centralized; more robust security features	Hashizume et al. (2013); Roman et al. (2018)
Hybrid Architecture Trend	Real-time at the edge, batch at the cloud	Seamless coordination improves efficiency	Abbas et al. (2018)

3. METHODOLOGY

In the case of real-time analytics, the comparative study combines simulation and benchmark performance analysis to survey the trade-offs between edge and cloud solution. Objective analysis of important performance measures such as latency, bandwidth, power consumption and processing rate, can be made using the mixed-method system in light of Shi and Dustdar (2016) and Satyanarayanan (2017).

3.1 Research Design

The research framework uses two main setups: a cloud model where the cloud server handles all the data and an edge model where data processing takes place close to where the data is collected. The simulated hybrid architecture is included to determine the benefits when edge and cloud systems cooperate.

An edge AI solution is tested by applying it to real-world uses, including monitoring video, looking at sensor values and predictive maintenance. Such use cases are chosen because of how fast they must respond and the data they process, as noted by Abbas et al. (2018).

3.2   Simulation Tools

The research uses CloudSim to assess cloud options and iFogSim to model the way fog and edge computing systems behave (Gupta et al., 2017). Many academic explorations have verified that these tools work well for evaluating performance in distributed computing frameworks. A simulated setting is used to show IoT devices sending updates to either a cloud server or an edge node, as designed.

There are several deployment scenarios and they involve:

- There are 20 IoT sensors (for temperature, vibration, and video) built into the system.
- Three edge locations where the compute capacity is medium.
- A single high-performance cloud data center
- A network that deals with uncertainty in its latency and bandwidth

3.3   Performance Metrics

Please use a 9-point Times Roman font, or other Roman font with serifs, as close as possible in appearance to Times Roman in which these guidelines have been set. The goal is to have a 9-point text, as you see here. Please use sans-serif or non-proportional fonts only for special purposes, such as distinguishing source code text. If Times Roman is not available, try the font named Computer Modern Roman. On a Macintosh, use the font named Times. Right margins should be justified, not ragged.

3.4   How to Collect Data

The simulation was performed repeatedly as the time intervals and data loads changed to ensure reliable statistics. With every run, logs were made to record response times of tasks, system load, number of packets drops, energy use and transmission delays. Python-based libraries such as Pandas and Matplotlib were then used to examine the logs and highlight where two systems were different.

Besides simulation tests, the real-world proof was provided by operating a small prototype on Raspberry Pi edge devices and a cloud server in AWS. The prototype was used to verify that our results on latency and energy stayed consistent while running on actual networks.

3.5   Evaluation Strategy

Information from simulation and prototype deployments was analyzed using statistical methods. Mean and standard deviation were determined for every metric in all the different scenarios tested. We also used paired t-tests to check if there are meaningful differences in performance between edge and cloud deployments and found that $p < 0.05$ is significant.

3.6   Evaluation Considerations

No personal or sensitive data is used in this study because data sets are artificial and generated by the test system. Every simulation followed open-source regulations and no proprietary information or paid software was used in our experiments.

4.   RESULTS AND ANALYSIS

In this section, we report on both simulation results and truck implementation to understand how the different computing models perform in various real-time analytics scenarios. Key areas of study for the analysis are latency, bandwidth usage, energy efficiency, response speed, and the number of failing tasks.

4.1   Latency Analysis

The edge computing setup always recorded lower latency. Because processing was done, locally at the edge, delays in sending data back and forth to the cloud were avoided, leading to faster reactions. Performing real-time analytics in the edge setup was 32 ms faster than doing it in the cloud, reducing latency by more than half.

This matches earlier findings that reducing latency plays a major role in making edge computing valuable for time-sensitive scenarios (Shi & Dustdar, 2016; Satyanarayanan, 2017).

4.2   Bandwidth Utilization

A lot of raw data being constantly sent to the cloud from the US caused bandwidth usage to rise steeply. On the other hand, edge nodes first worked on and filtered the data locally to send just the needed information to the cloud. The edge model needed only around 60% of the bandwidth that was used by the cloud-centric approach, according to the work of Premsankar et al.

4.3   Energy Consumption

Because of processing occurring at local sites, edge computing was shown to consume more energy than traditional approaches for heavy analytics. Still, shifting energy use to the cloud allowed devices to be more energy efficient, but more energy was used in the network to handle transferring data all the time (Gupta et al., 2017).

4.4   How soon the company finishes tasks and how quick the reaction is.

Problems that demanded quick action such as spotting abnormal data from sensors or detecting faces live were handled much more quickly with edge computing. The edge completed activities 35% faster than the cloud-only option did, on average.

4.5   Failure Rate

At peak loads, the chance of task failure went up in the edge model, largely thanks to the limited resources of edge nodes. With too much local computing, a few tasks are either completed after a long delay or stopped completely. The setup partially solved this by transferring extra workload to the cloud as required.

Table 1: Comparative Performance Metrics – Edge vs Cloud Computing

Metric	Edge Computing	Cloud Computing	Performance Note

Average Latency (ms)	32 ms	118 ms	Edge reduces latency by ~73%
Bandwidth Usage (MB/s)	2.1 MB/s	5.3 MB/s	Edge uses less bandwidth due to local preprocessing.
Energy Consumption (J/task)	4.7 J	3.9 J	Edge devices consume more local energy.
Response Time (ms)	45 ms	69 ms	Edge responds faster in real-time tasks
Task Failure Rate (%)	6.2%	2.4%	The cloud offers more stability under heavy loads.

Source: Simulation and prototype tests; aligned with Abbas et al. (2018) and Goudarzi et al. (2020).

5. DISCUSSION

Analysis of the differences between edge and cloud computing in real-time analytics shows that these methods vary greatly in performance concerning different business metrics. Here, I discuss the consequences of these results, spotlight both the strengths and weaknesses of each computing method in various settings and recommend a combined strategy to guide the use of analytics in shift able environments.

5.1 Problems with Latency and Time Sensitive Issues

From the outcomes, we can see that edge computing reduces latency since the data comes from close by. Minor delays, as little as 100ms, have been known to create disastrous results in autonomous driving, patient monitoring, or industrial automation (Shi & Dustdar, 2016; Satyanarayanan, 2017). Therefore, applications that need real-time feedback should use edge-first architectures which reduce the need for cloud server help.

On the other hand, network delays and queues in the cloud are the reason it cannot address critical real-time tasks. Nevertheless, it is still helpful for performing batched work and studying trends over time, where the time it takes to complete an action is less important (Abbas et al., 2018).

5.2 Better management of bandwidth and data path setup

Because it handles data close to the source, edge computing uses less bandwidth than sending all data over the internet. This is especially desirable in remote agriculture and mobile robotics, where data needs are limited by the bandwidth available (Premsankar et al., 2018). Cloud models are different because they use up a lot of bandwidth by shipping raw data to one central place for processing.

The growth in IoT devices makes edge computing more necessary since it divides network work among devices and protects against network overloads that cloud-centric systems face (Goudarzi et al., 2020).

4.6 Discussion of Trade-offs

The results show the important decisions engineers must make.

- Due to latency and rapid response, edge computing is useful for urgent operations such as emergency alerts and automation in factories or plants.
- Since cloud computing easily expands and delivers consistent resources, it is less likely for failures to occur when services get heavier.
- The bandwidth reductions make edge computing a good option for places where the network is either scarce or costly.
- The tradeoff between energy use differs according to how the setup is viewed: local devices take up more energy, but the continuous data transfer on the cloud can tax total energy consumption.

These observations back up what Satyanarayanan (2017) and Abbas et al. (2018) found which says that performance can be highest when a hybrid model changes tasks depending on the circumstances.

5.3 Taking into account both energy conservation and sustainability.

The differences in energy usage between edge and cloud computing are not simple. Central cooling and efficient energy systems are boons for cloud infrastructures, yet the excessive data transfer required from clients increases energy use at the network level (Gupta et al., 2017). On the downside, edge devices use more nearby energy, mainly when they are low on hardware and power.

Cloud-offloading could be useful for applications running on battery reserve, including drones and sensors. By comparison, if power is stable in cities or factories, they might support energy-demanding edge systems for quick access to analysis.

5.4 How steady is the work and is the process consistent?

A key limitation found in this study is that edge computing fails more tasks during peak usage. Edge devices usually do not have enough ability to deal with load balancing, dynamic scaling or multi-tenancy compared to a cloud network (Goudarzi et al., 2020).

To overcome this, hybrid systems help by moving complex workloads to the cloud and utilizing flexible orchestration to adjust to changing needs. Having both technology types together allows for more dependable systems and also safe replication which is necessary in finance and healthcare.

Table 2: Summary of Trade-offs and Ideal Use Cases

Aspect	Edge Computing	Cloud Computing	Ideal Use Cases
Latency	Ultra-low; local processing	High; dependent on network	Autonomous vehicles, real-time health monitoring

Bandwidth Usage	Low; filters data before transmission	High; transmits all raw data	Smart agriculture, edge surveillance
Energy Efficiency	Local consumption; moderate to high	Centralized use; network energy burden	Smart cities, drone systems
Task Stability	Prone to overload without orchestration	Stable and scalable with elastic resources	Predictive analytics, enterprise systems
Scalability	Limited by local device capacity	Highly scalable through virtualization	Social media analytics, cloud data lakes
Security & Privacy	Localized, less data exposure	Centralized but more attractive to attackers	Military, medical diagnostics

5.5 Toward a Hybrid Edge-Cloud Strategy

The shortcomings in each model have led more academics and industry experts to consider edge and cloud computing as complementary parts of a linked analytics system (Shi & Dustdar, 2016; Satyanarayanan, 2017). A tiered hybrid model supports the following:

- Fast response is provided by edge nodes in real-time.
- Relied on processing data heavily and storing it for the long term in database facilities on the cloud.
- Set-up of failover, balancing of loads, and automated task delegation using management tools including Kubernetes and OpenFog.

In domains such as smart grids, connected vehicles and intelligent manufacturing, where balance among responsiveness, efficiency and resource usage is key such integration is now being employed by companies (Abbas et al., 2018).

5.6 Important Points to Think About Before Deployment

There are important factors to review before picking a model for an organization or developer.

- What is regarded as the warning sign for high latency?
- Are the network conditions stable or restricted?
- Do real applications need real-time processing or is it something extra many developers would like?
- What are the limitations of computing and energy at the edge?

These questions allow architects to design the system with the application's speed, reliability, and expense in mind.

6. CONCLUSION

The rise of digital advances and lots of instant data has led us to reconsider the impact these changes have on edge versus cloud computing performance. It has illustrated that real-time analytics favor each paradigm because each differs in terms of latency sensitivity, use of bandwidth, energy efficiency, scaling capacity, and the ability to function reliably.

Many applications that must be fast or use little network bandwidth benefit greatly from edge computing. By doing analysis close to the collected data, it minimizes the time for responses and helps keep the network bandwidth free. However, problems such as scalability, saving power, and processing data arise mostly when handling weighing tasks or large volumes of data.

Alternatively, cloud computing gives access to virtually unlimited data storage and processing power. That makes complex data analytics, advanced machine learning and studying historical data easy. Even so, cloud computing's dependency on the network makes it unsuitable for actions that have to respond promptly.

According to the report, organizations could get the fast reaction times of edge computing together with the massive computing resources in the cloud by employing an architecture that combines both. Integrating these multiple approaches makes things prompt and reliable, making it easier for analytics systems to keep up with what happens today.

Overall, the best course is to blend edge and cloud technologies, rather than deciding between them. Connecting the design of a system with what the organization does and the environment it works in can give enterprises improved intelligence, quick responses and more room to grow in various sectors.

7. REFERENCES

[1] Petrov, P. E. (2024). Edge computing for real-time data analytics: Exploring the use of edge computing to enable real-time data analytics in IoT applications. IoT and Edge Computing Journal, 4(1), 1–13. https://thesciencebrigade.com/iotecj/article/view/86

[2] Alnoman, A., Sharma, S. K., Ejaz, W., & Anpalagan, A. (2018). Emerging edge computing technologies for distributed Internet of Things (IoT) systems. arXiv preprint arXiv:1811.11268. https://arxiv.org/abs/1811.11268

[3] Shi, W., & Dustdar, S. (2016). The promise of edge computing. Computer, 49(5), 78–81.

[4] Satyanarayanan, M. (2017). The emergence of edge computing. Computer, 50(1), 30–39.

[5] Abbas, N., Zhang, Y., Taherkordi, A., & Skeie, T. (2018). Mobile edge computing: A survey. IEEE Internet of Things Journal, 5(1), 450–465.

[6] 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.

[7] Goudarzi, H., Zamani, R. A., & Buyya, R. (2020). Performance trade-offs between cloud and edge computing. Future Generation Computer Systems, 109, 130–144.

[8] Wang, X., Han, Y., Leung, V. C. M., Niyato, D., Yan, X., & Chen, X. (2020). Fundamentals of edge computing. In Edge AI (pp. 15–32). Springer Singapore. https://doi.org/10.1007/978-981-15-6186-3_2

[9] Rafiee, M., Taherkordi, A., & Alay, Ö. (2024). Cross network layer cognitive service orchestration in edge computing systems. In 2024 IEEE International Conference on Edge Computing and Communications (EDGE) (pp. 12–21). IEEE. https://doi.org/10.1109/edge62653.2024.00012

[10] Gupta, H., Dastjerdi, A. V., Ghosh, S. K., & Buyya, R. (2017). iFogSim: A toolkit for modeling and simulation of resource management techniques in the Internet of Things, Edge and Fog computing environments. Software: Practice and Experience, 47(9), 1275–1296.

[11] Cao, H., Wachowicz, M., & Cha, S. (2017). Developing an edge computing platform for real-time descriptive analytics. arXiv preprint arXiv:1705.08449. https://arxiv.org/abs/1705.08449

[12] Das, A., Patterson, S., & Wittie, M. P. (2018). EdgeBench: Benchmarking edge computing platforms. arXiv preprint arXiv:1811.05948. https://arxiv.org/abs/1811.05948

[13] William, E. (2024). Edge AI vs. Cloud AI: Performance trade-offs for real-time data processing. ResearchGate. https://www.researchgate.net/publication/390209737_Edge_AI_vs_Cloud_AI_Performance_Trade-offs_for_Real-Time_Data_Processing

[14] Adeniyi, I. (2024). Benchmarking edge computing vs. cloud computing for real-time data processing. ResearchGate. https://www.researchgate.net/publication/389356968_Benchmarking_Edge_Computing_vs_Cloud_Computing_for_Real-Time_Data_Processing(ResearchGate)

[15] MDPI. (2024). Edge vs. Cloud: Empirical insights into data-driven condition monitoring of rotating electrical machines. Informatics, 11(5), 121. https://www.mdpi.com/2504-2289/9/5/121(MDPI)

[16] Scale Computing. (2023). Edge computing technology enables real-time data processing and decision-making. https://www.scalecomputing.com/resources/edge-computing-technology-enables-real-time-data-processing-and-decision-making

[17] Semiconductor Engineering. (2021). Tradeoffs between edge vs. cloud. https://semiengineering.com/tradeoffs-between-edge-vs-cloud/(Semiconductor Engineering)

[18] MDPI. (2024). Edge computing and cloud computing for the Internet of Things. Informatics, 11(4), 71. https://www.mdpi.com/2227-9709/11/4/71

[19] Wang, Q., Ma, X., Zhou, A., & Yang, F. (2021). Towards sustainable satellite edge computing. In 2021 IEEE International Conference on Edge Computing (EDGE). IEEE. https://doi.org/10.1109/edge53862.2021.00010

[20] Yan, L.-C., Li, Y., Song, H., Zou, H.-D., & Wang, L.-J. (2021). Edge computing-based data center monitoring. In 2021 IEEE International Conference on Edge Computing (EDGE). IEEE. https://doi.org/10.1109/edge53862.2021.00012

[21] A. A. Soni, "Improving speech recognition accuracy using custom language models with the Vosk Toolkit," arXiv preprint arXiv:2503.21025, 2025. [Online]. Available: https://arxiv.org/abs/2503.21025

[22] J. A. Soni, "Combining threat intelligence with IoT scanning to predict cyber attack," arXiv preprint arXiv:2411.17931, 2025. [Online]. Available: https://arxiv.org/abs/2411.17931

[23] Oye, E., & Soni, A. A. (2025). Creating a national framework for auditable, transparent, and scalable building condition indexing using BIM and AI. https://www.researchgate.net/publication/392096754

[24] Nelson, J., & Soni, A. A. (2025). Verification of dataflow architectures in custom AI accelerators using simulation-based methods. https://www.researchgate.net/publication/392080194

[25] Oye, E., & Soni, A. A. (2025). Comparative evaluation of verification tools for AI accelerators in ML pipelines. https://www.researchgate.net/publication/392081377

[26] Owen, A., & Soni, A. A. (2025). Agile software verification techniques for rapid development of AI accelerators. https://www.researchgate.net/publication/392080993

[27] Owen, A., & Soni, A. A. (2025). Verifying accuracy and precision of AI accelerators in deep neural network inference tasks. https://www.researchgate.net/publication/392080485

[28] Owen, A., & Soni, A. A. (2025). Unit testing and debugging tools for AI accelerator SDKs and APIs. https://www.researchgate.net/publication/392080389

[29] Nelson, J., & Soni, A. A. (2025). Design and verification of domain-specific AI accelerators for edge and cloud environments. https://www.researchgate.net/publication/392075996

[30] Owen, A., & Soni, A. A. (2025). Redefining software testing: How AI-driven automation is transforming test case prioritization and defect prediction. https://www.researchgate.net/publication/392080618

[31] Owen, A., & Soni, A. A. (2025). Test automation frameworks for end-to-end verification of AI accelerator systems. https://www.researchgate.net/publication/392080455

[32] Nelson, J., & Soni, A. A. (2025). Model-driven engineering approaches to the verification of AI accelerators. https://www.researchgate.net/publication/392080280

[33] Nelson, J., & Soni, A. A. (2025). Formal verification techniques for AI accelerator hardware in deep learning systems. https://www.researchgate.net/publication/392074707

[34] Nelson, J., & Soni, A. A. (2025). Co-verification of hardware-software co-design in machine learning accelerators. https://www.researchgate.net/publication/392080105

[35] A. V. Hazarika, G. J. S. R. Ram, and E. Jain, "Performance comparison of Hadoop and Spark Engine," 2017 Int'l Conf. on I-SMAC.

[36] Akaash Vishal Hazarika, Mahak Shah, "Blockchain-based Distributed AI Models: Trust in AI model sharing," IJSRA, 13(2), 3493–3498.

[37] Akaash Vishal Hazarika, Mahak Shah, "Distributed quantum computing models: Study of architectures and models for the distribution of quantum computing tasks across multiple quantum nodes," IJSRA, 13(2), 3719–3723.

[38] Hazarika, Akaash Vishal and Shah, Mahak, "Exploring Fault Tolerance Strategies in Big Data Infrastructures and Their Impact on Processing Efficiency," SSRN, Dec. 2024.