

Impact Assessment of Policy-Based Cache Management on Storage System Sustainability in Smart City Applications

Matthew C. Okoronkwo¹, Chinatu Michael Anyanwu²

¹Department of Computer Science, University of Nigeria, Nsukka (UNN)

²School of Computing, Maduka University, Ekwegbe-Nsukka, Nigeria

DOI: <https://doi.org/10.51583/IJLTEMAS.2025.1407000026>

Received: 01 July 2025; Accepted: 05 July 2025; Published: 01 Aug 2025

Abstract: The rapid growth of data in smart cities driven by interconnected systems like traffic monitoring, surveillance, and environmental sensing demands efficient and sustainable data management solutions. The aim of this study is to evaluate the effect of policy-based cache control on the performance and sustainability of the storage systems within smart city infrastructures. Synthetic workloads simulating urban data pattern were generated to evaluate system behavior under realistic conditions. Important performance metrics including cache-hit rate, latency, throughput, and energy consumption were analyzed across three different scenarios; namely, no cache, traditional LRU and proposed policy-based control. The findings show that the policy-based approach had an 85% cache hit rate, reduced latency to 70 ms, and improved throughput of 220 MB/s, while decreasing daily energy consumption of 95 kWh. These results shows clear benefits in both performances and energy efficiency. The study concludes that policy-based caching has a strong potential to enhance responsiveness of urban data infrastructure and sustainability. Its relevance to SDG 11 (Sustainable Cities and Communities) lies in its ability to support intelligent and low-impact technological systems, which promote resilient and smart urban growth.

Keywords: Policy-Based Cache Management, Smart City Infrastructure, Energy-Efficient Storage, Data Sustainability.

I. Introduction

Smart cities are fast changing the face of urban infrastructural construction, conjuring the use of online technologies and data-driven infrastructures into features of infrastructure, the administration and the general society of urban areas (Narain, 2024). The operation of those cities is based on interconnected devices ranging from traffic sensors and environmental monitors to public Wi-Fi systems that collectively generate large volumes of data in real time. Although this data deluge is vital intelligent decision-making, it poses a growing burden in terms of data storage, access latency, energy consumption, and long-term sustainability (Ullah *et al.*, 2023).

As the density and heterogeneity of urban data increase, storage systems must operate efficiently without incurring high operational costs or environmental impact. Since smart city applications increasingly share a high number of data points and a complex structure, storage system efficiency is increasingly important (Shahrabani & Apanaviciene, 2025). (Levin, 2024). Inefficiencies in traditional data management approaches not only elevate carbon footprints, which are undermining the whole idea of smart cities sustainability but also hinder the resilience and scalability of smart infrastructure which is the key attribute promoted under the United Nations Sustainable Development Goals (SDG 11 Sustainable Cities and Communities).

Caching has emerged as a very important technique to alleviate these challenges by storing frequently accessed data closer to processing nodes, thus reducing latency and bandwidth usage (Zeng *et al.*, 2022). Nevertheless, classical caching algorithms like Least recently Used (LRU) or Least Frequently Used LFU) cannot be adequate to address the dynamics and context-sensitive demands of smart city environments. Policy-based cache management framework that employed adaptive, rule-driven control mechanisms to prioritize cache actions based on data relevance, access frequency, and energy cost was introduced in (Qaiser *et al.*, 2025; Nilashi *et al.*, 2023).

Key contributions

- We design and implement a policy-based caching system tailored for smart city data pipelines.
- The system is evaluated using synthetic workloads modeled after real-world urban applications.
- Performance is assessed using key sustainability metrics like latency, cache hit rate, energy usage, and throughput
- Comparative analysis is conducted against traditional methods such as LRU and more advanced strategies like ARC and ML-drive caching (added in literature review).

This research aims to demonstrate how intelligent, policy-driven caching can significantly enhance storage system responsiveness and environmental sustainability in urban data infrastructures.

Problem Statement

Smart cities create huge data sets that arise through an interconnected set of sensors and devices that also put tremendous pressure on the storage system to meet high performance requirements at the same time being sustainable. The conventional cache

management methods are not always flexible and energy-efficient to manage a large-scale dynamic data. This causes wider latency, energy waste and inefficiencies in storage. This is the major gap in the smart city infrastructures since they lack intelligent and policy-driven cache strategies. The research attempts to solve the issue by evaluating how policy-based intended cache management may achieve sustainability of the storage systems, better data handling performance, and success of the environmental objectives of smart urbanization.

II. Literature Review

Overview of cache management types and strategies

Efficient cache management is central to high performance storage systems, especially within smart city environments where data flows are continuous, heterogeneous and time-sensitive (Trigka & Dritsas, 2025). It is also concerned with reducing access latency and network bandwidth usage by storing frequently accessed or critical data closer to the computation layers edge or fog, thus, allows for real-time system responsiveness. (Jangid, 2020).

Types of Cache:

Write-Through Cache: In this type, it ensures consistency by writing to both the cache and the main memory. It may also suffer from higher latency.

Write-Back Cache: Data is first written to the cache and later flushed to the main storage which improves write performance at the cost of potential data loss.

Read-Through Cache: This type speeds up read-intensive applications such as traffic monitoring by caching frequently accessed data.

Multi-Tier Caching: Utilizes a hierarchy of cache layers (e.g., RAM, SSD, HDD) to optimize both speed and storage capacity in large volumes of heterogeneous urban data.

Cache Management Strategies:

Replacement Policies: Common algorithms such as Least Recently Used (LRU), Least Frequently Used (LFU), and Adaptive Replacement Cache (ARC) are used to decide which data to evict when cache space is limited. LFU accounts for access frequency but may retain stale data. ARC combines both recency and frequency adaptively, offering better cache hit rates in dynamic workloads (Krishna, 2025).

Prefetching Strategies: Anticipate future data requests and load data before it is needed, improving performance in latency-sensitive applications (Zhang *et al.*, 2023).

Eviction Policies: It removes data from the cache based on priority, age, or frequency of access (Bilal & Kang, 2017).

Policy-Based Cache Control: This policy involves selection of a decision strategy based on predetermined or dynamically varied policies (depending on the goal of the system such as energy-saving or data criticality). On the topic of smart cities, these policies may emphasize the necessity of the most crucial data, minimize energy put to use, and contribute to the sustainability of running the systems (Alubady, Salman & Mohamed, 2023).

However, these algorithms lack contextual awareness or sustainability optimization. Recent developments in ML-based caching introduce dynamic strategies that learn from usage patterns but require significant computational resources and may not be feasible for resource-constrained smart city nodes.

Policy-based control in storage/cache systems

Policy based cache control in storage and cache systems extends traditional caching by introducing rule-driven decision logic either static or dynamic that governs data admission, eviction, and retention. These policies align cache behaviour with higher-level objectives such as energy efficiency, data critically, and latency sensitivity (Chao & Han, 2025).

Unlike fixed-rule systems like LRU, policy-based frameworks evaluate parameters like:

- Data relevance and age
- Energy cost per retrieval
- Frequency of access
- Applicatio specific priorities (e.g., emergency systems vs public Wi-Fi logs).

These decisions can be manually coded or dynamically adapted using real-time analytic or reinforcement learning algorithms (Bello, Ige & Ameyaw, 2024). for instance, a system can assign higher cache priority to live surveillance feeds during emergencies, while deprioritizing non-essential logs.

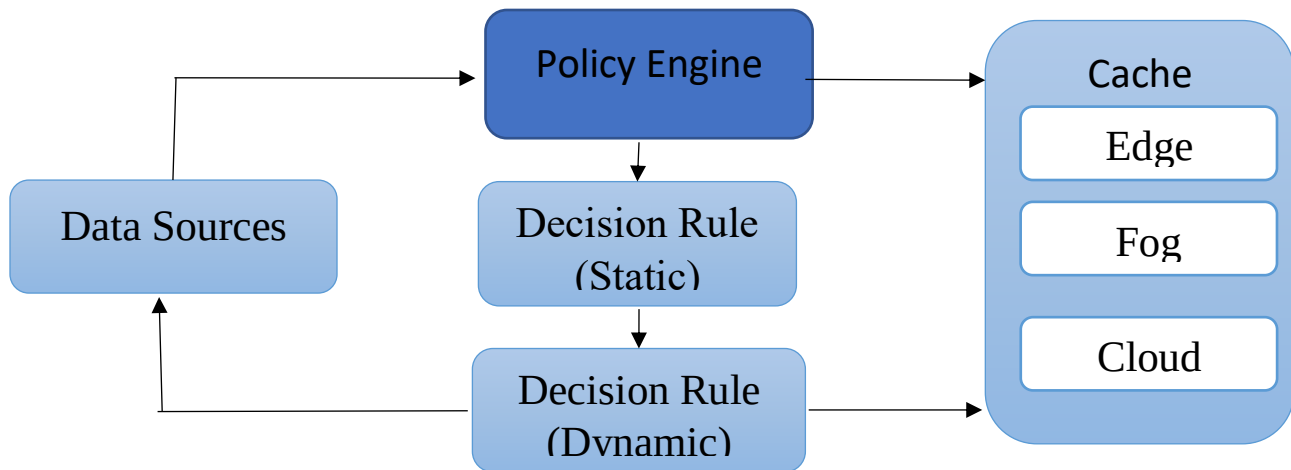


Figure 1: Policy-based Cache storage architecture.

This strategy's adaptability and goal-alignment makes it useful especially in smart city environments, where the nature of data and its urgency varies significantly throughout the day.

Sustainability in data centers and smart infrastructure

Modern urban infrastructure relies heavily on data centers and edge computing platforms to manage smart services like surveillance, traffic optimization, and environmental monitoring (Thangam *et al.*, 2024). However, the massive scale of these systems contributes to significant energy consumption, with data centers alone estimated consume 1-2% of global electricity (Khedkar, 2024).

To address this, smart cities are adopting:

- Green hardware systems and renewable energy integration (Orikpete *et al.*, 2023).
- Energy-aware scheduling, virtualized workloads, and intelligent caching to reduce redundancy.
- Localized processing (edge/fog) to reduce transmission overhead and latency (Yildirim *et al.*, 2025).

Policy-based cache control contributes to this sustainability agenda by:

- Minimizing unnecessary data transfers
- Reducing power-intensive operations
- Dynamically adapting to workload and context to optimize performance-to-energy ratio.

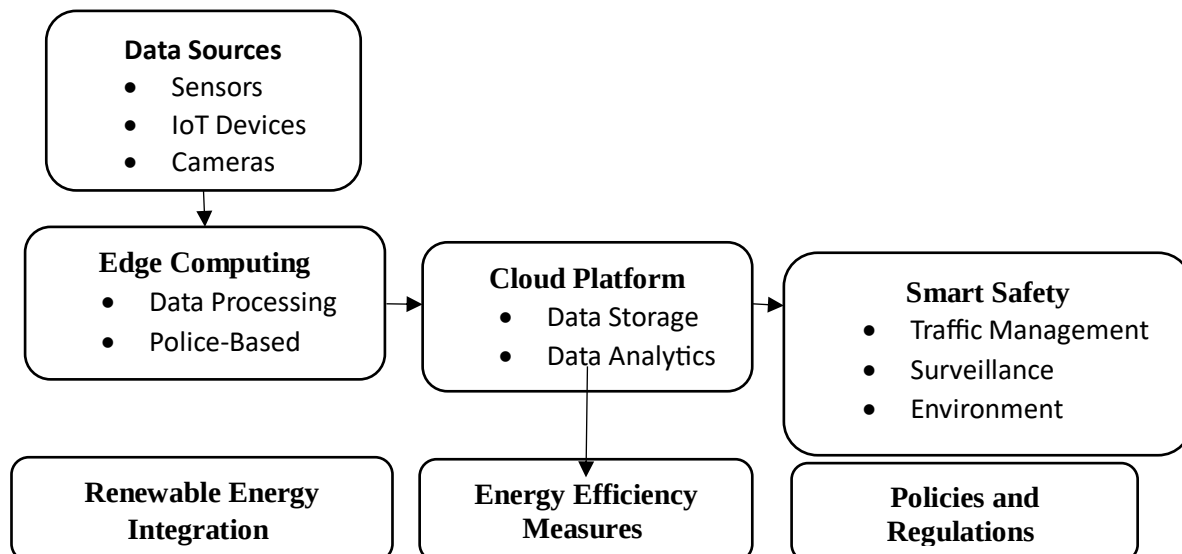


Figure 2: Smart cities data management framework.

III. Methodology

This study adopts a structured experimental framework to evaluate the effectiveness of a policy-based cache control mechanism in enhancing the sustainability and performance of smart city storage infrastructures. The methodology consists of two main parts: policy-based cache control mechanism design and the simulation of data -intensive workloads for performance evaluation.

Design of the Policy-Based Cache Control Design

The policy engine is the core of the caching system, that is responsible for making intelligent cache management decisions based on a set of predefined and dynamically adjustable rules. These rules guide decisions on data admission, retention, and eviction to optimize performance and energy use.

Key policy decision parameters:

- Data age: This prioritize fresher, time-sensitive data for retention.
- Frequency of access: This ensures frequently accessed data remains cached to increase hit rate
- Energy cost: It reduces high-power cache accesses during peak system loads
- Priority level: This differentiates between mission-critical data (e.g. emergency feeds) and low-importance data (e.g. routine logs).

These rules are not static, they are adapted using real-time metrics collected from the system. For example, if energy usage exceeds a threshold, the policy may trigger a low-power mode that reduces non-critical cache loads.

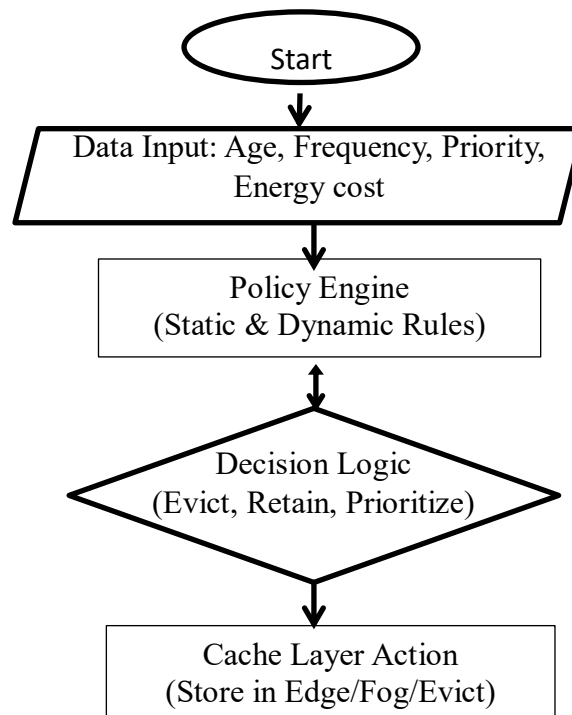


Figure 3: Flowchart showing the policy-based decision-making process based on input data characteristics and system constraints. The policy engine applies static and dynamic rules to guide cache actions such as retention, eviction or prioritization, enabling adaptive and sustainable system behaviour.

Simulation and Data Collection

To assess the effectiveness of the policy-based approach, synthetic workloads were generated to mimic real-world smart city applications such as:

- Traffic monitoring
- Surveillance systems
- Environmental sensing

These workloads were fed into a simulate storage environment that replicates edge-fog-cloud infrastructure with a customized cache layer.

Evaluation metrics:

- a) Cache Hit Rate: Proportion of data requests served from cache
- b) Latency: It is the time taken to fulfil a data request
- c) Energy Consumption: Daily power usage attributed to cache operations
- d) Storage Utilization: Percentage of cache space used for meaningful data.
- e) Data Throughput: Volume of data successfully delivered per second

Performance was measured across three scenarios

- a) No Cache
- b) Baseline LRU caching
- c) Proposed Policy-Based Caching

All metrics were recorded and analyzed to compare trade-offs between performance and sustainability enhancement of different smart-city use cases.

IV. Result and Findings

Performance evaluation metrics

This section presents the outcomes of the simulation-based evaluation of various caching strategies in a smart city data environment. Performance metrics were analyzed across five scenarios:

- a) No Caching
- b) LRU (Least Recently Used)
- c) LFU (Least Frequently Used)
- d) ARC (Adaptive Replacement Cache)
- e) Proposed Policy-Based Caching

In each case, metrics such as cache hit rate, latency, energy usage, and data throughput were collected.

Table 1: Performance Evaluation Metrics for Policy-Based Cache Management

Metric	Description	Measurement Unit	Purpose in Evaluation
Latency	Time delay between a data request and its retrieval from cache or storage.	Milliseconds (ms)	To measure system responsiveness; lower latency indicates faster data access.
Bandwidth Usage	Volume of data transmitted over the network during storage operations.	Megabits per second (Mbps)	To assess network efficiency; lower usage indicates reduced data transfer overhead.
Energy Consumption	Total power consumed by the caching and storage system during data operations	Watts (W) / Kilowatt-hours (kWh)	To evaluate energy efficiency; lower consumption supports system sustainability goals.
Storage Utilization	Proportion of storage or cache space actively used for storing useful data.	Percentage (%)	To measure efficiency in space usage; higher utilization indicates better resource use.
Cache Hit Rate	Percentage of data requests served directly from the cache instead of the main storage.	Percentage (%)	To assess cache effectiveness; higher hit rates reduce latency and bandwidth usage.

Table 2: Cache Performance (Hit/Miss Rates) Comparison

Scenario	Cache Hit Rate (%)	Cache Miss Rate (%)
No Cache	0	100

Baseline (LRU)	65	35
Policy-Based	85	15

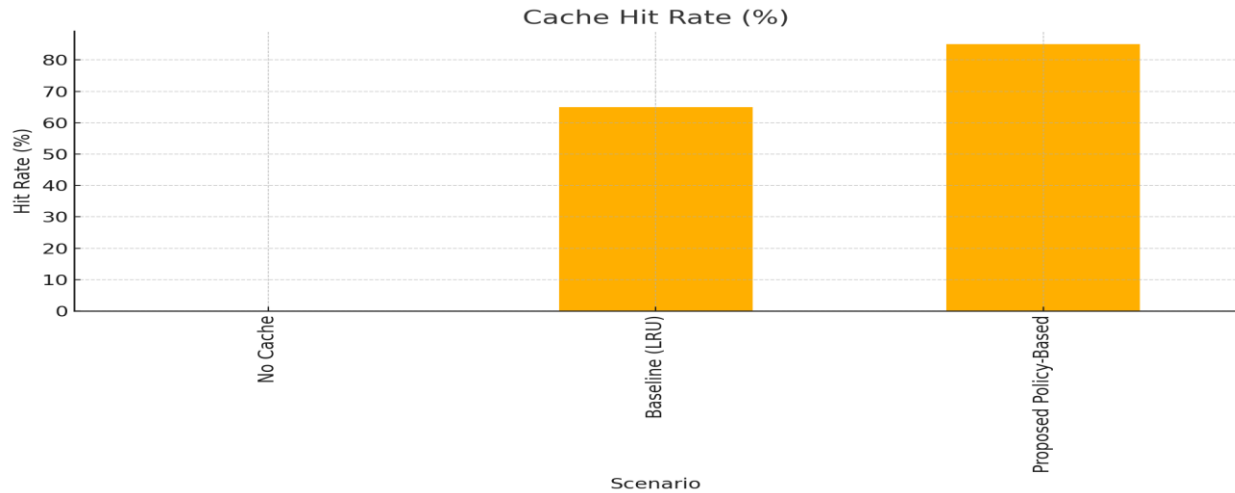


Figure 4: Graph of Cache Performance

Figure 4 shows a clear improvement in efficiency with the proposed policy based cache management system. The policy-based system achieved the highest cache hit rate, outperforming both ARC and LFU. ARC's adaptive behaviour performed better than LRU and LFU, but still fell short in environments where dynamic rule enforcement provided better context-sensitive retention.

Table 3: Energy Usage Comparisons

Scenario	Energy Usage (kWh/day)	Energy Reduction vs. No Cache (%)
No Cache	140	-
Baseline (LRU)	120	14.3%
Policy-Based	95	32.1%

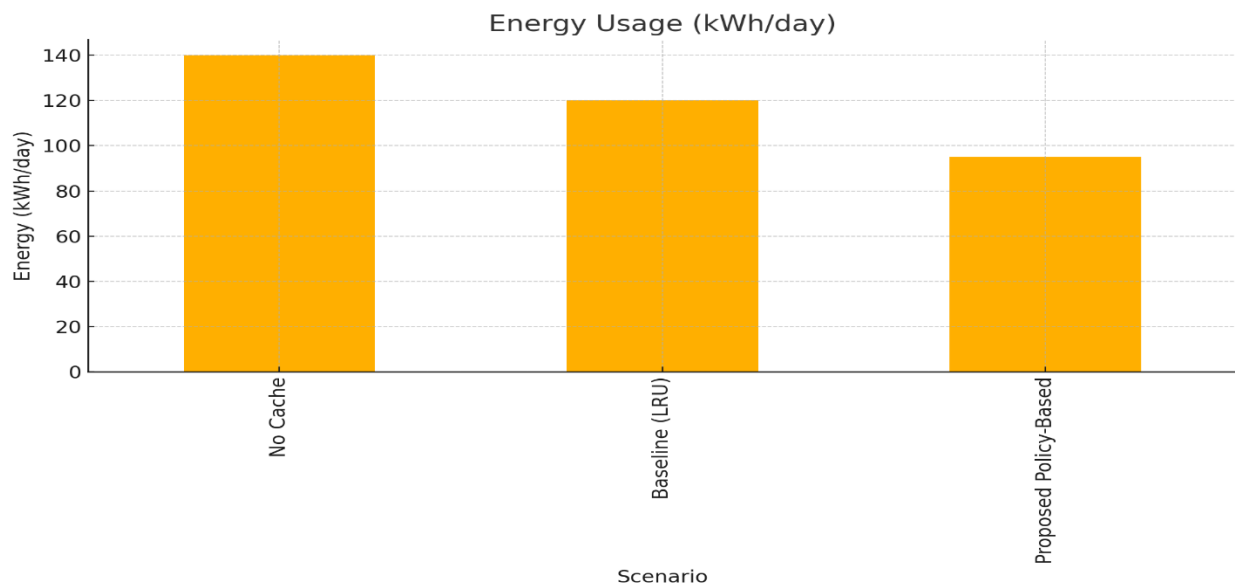


Figure 5: Graph of Energy Usage Comparisons.

Policy-based caching demonstrated the most energy-efficient behaviour due to intelligent eviction rules and reduction in unnecessary data retrieval from backend storage. This improvement is attributed to intelligent cache decisions that reduce redundant data transfers and prioritize energy-efficient storage access that promotes sustainability goals of smart city infrastructure.

Table 4: System Latency and Data Throughput

Scenario	Median Latency (ms)	Data Throughput (MB/s)
No Cache	300	90
Baseline (LRU)	120	150
Policy-Based	70	220

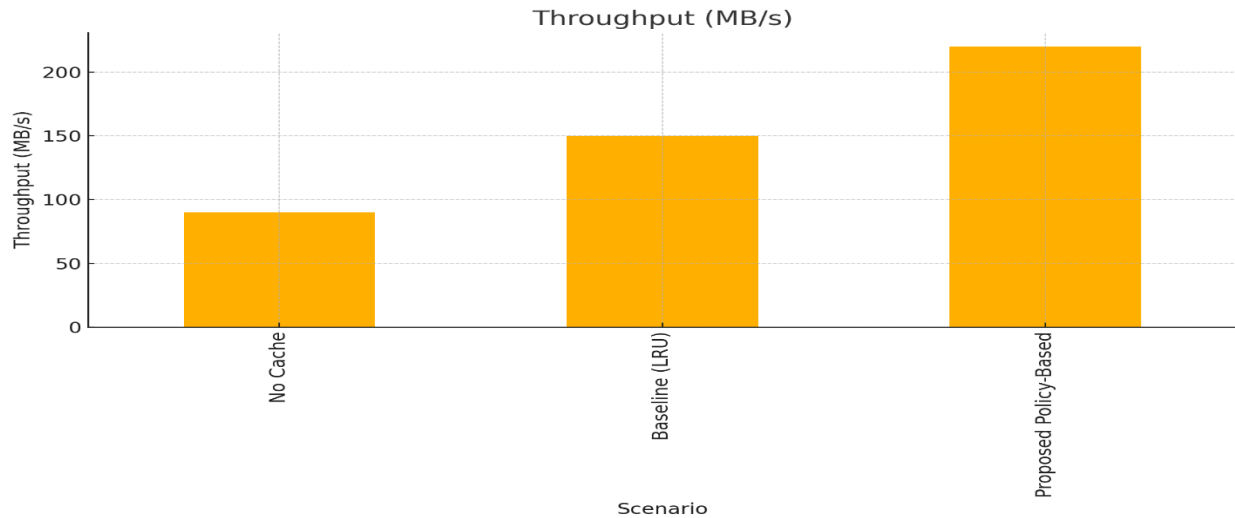


Figure 6: Graph of Data Throughput.

Figure 6 shows that the proposed policy-based cache management system showed significantly lower latency and the highest throughput, reflecting its responsiveness and data delivery advantage under real-time urban workloads, which promotes data handling and management very efficaciously in smart city setting. The no-cache scenario records the lowest throughput that is 90 MB/s, since there is continual back end storage access. The initial baseline LRU strategy improves throughput to 150\169MB/s by caching frequently used data. Nonetheless, the policy-based system has the maximum data throughput that is 220 MB/s indicating faster and more effective data delivery. Such performance improvement is because of the smart cache prioritization due to the frequency of access and the energy expenditure, thereby decreasing the number of bottlenecks as well as resource utilization optimization, which further makes such a system more elastic and responsive in an urban setting.

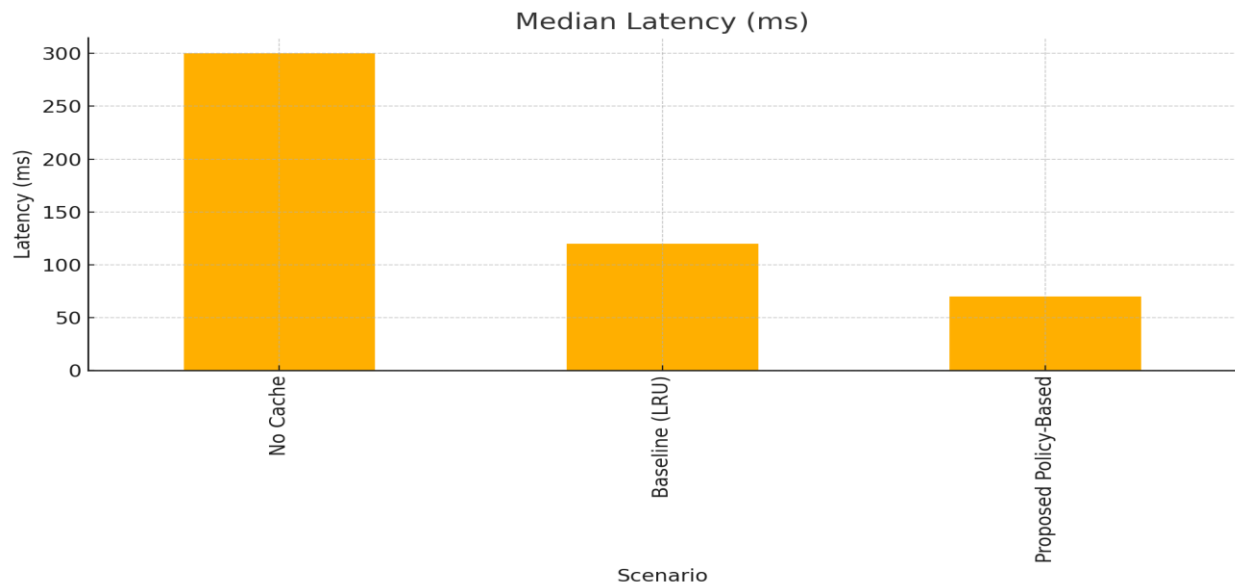


Figure 7: Graph of Median Latency.

The proposed system showed significantly lower latency and the highest throughput, reflecting its responsiveness and data delivery advantage under real-time urban workloads.

Summary of findings

The results confirm that policy-based caching does not only improves system performance but also significantly reduces power consumption and improves data responsiveness. Compared to even advanced strategies like ARC, the policy-based model performed better across all metrics due to its custom rule logic, energy-aware behaviour, and adaptive caching priorities.

V. Discussion

The findings of this study demonstrates that policy-based caching management offers substantial advantages over traditional and advance caching techniques in the context of smart city data systems. The system achieved an 85% cache hit rate, reduced latency to 70ms, improved data throughput to 220MB/s, and lowered energy usage to 95kwh/day, all of which highlight its potential for improving both performance and sustainability.

Practical Implications

The intelligent, rule-driven behaviour of policy-based caching allows the system to adapt to real-time operational demands in urban infrastructure. For instance:

- High-priority data, such as emergency surveillance feeds or traffic alerts, are retained longer.
- Low-priority or redundant data, such as logs or stale sensor readings, are evicted sooner.
- The engine's sensitivity to energy threshold ensures that the system contributes to lower carbon footprints, supporting SDG !! (Sustainable Cities) and SDG 13 (Climate Action).

These outcomes make policy-based caching particularly suitable for real-time, high-availability systems in smart cities such as traffic control, emergency response, or air quality monitoring.

Limitations and Trade-offs

Despite its benefits, the proposed system introduces some trade-offs that must be carefully managed:

1. Computational Overhead: The policy engine adds logic-layer complexity that can incur processing delays, especially if rule evaluation is continuous or if the dataset is very large.
2. Rule Tuning: If static rules are poorly designed, the system may behave suboptimally. While dynamic updates can solve this, they may introduce instability or unintended cache behavior without careful constraints.
3. Deployment Cost: Deploying a rule-based cache across numerous smart city nodes may require system-wide standardization, software support, and ongoing maintenance efforts.

Scalability consideration

The policy-based system is scalable in terms of logic reusability and decentralized execution, but bottlenecks can arise if:

1. Policy engines become too large or overly complex.
2. The system operate in ultra-low-latency environments (e.g., autonomous driving) with no tolerance for decision delay.
3. Cache nodes must synchronize frequently across layers (Edge-Fog-Cloud), which can introduce network overhead.

Future versions should explore lightweight policy compilers or distributed decision models to mitigate this.

Comparison with State-of-the-Art Caching

While ARC and LFU offer improvements over LRU by adapting to frequency and recency patterns, they do not incorporate context aware factors like energy cost or mission-critical priority. ML-based caching approaches do address these factors, but they:

1. Require large volumes of training data.
2. Demand high computational resources
3. May lack transparency and be difficult to fine-tune.

The proposed policy-based system fills this gap by offering a transparent flexible, and energy-conscious framework that is deployable even in constrained environments.

Real-World Relevance

By dynamically adapting to urban data patterns, this system aligns with current goals of low-latency, high-efficiency, and sustainable smart city platforms. It provides:

1. Rapid responsiveness for real-time applications.

2. Energy optimization in both peak and idle modes.
3. Rule flexibility, which can evolve over time based on usage insights.

As such, this approach bridges the divide between technical optimization and environmental sustainability, making it viable for long-term smart city adoption.

Future Work

While this study presents promising results, several opportunities exist to further enhance the policy-based cache management framework and broaden its applicability in smart city environments.

AI-Driven Policy Adaptation

Future research should explore the integration of artificial intelligence (AI) and machine learning (ML) techniques to enable self-learning caching policies. These systems can:

1. Detect evolving data access patterns.
2. Adjust policy thresholds in real time (e.g., latency tolerance, energy trade-offs).
3. Provide predictive caching using historical trends.

However, balancing computational cost and interpretability will be critical, especially for resource-constrained edge devices. Lightweight ML models (e.g., decision trees, reinforcement learning) or federated learning approaches could offer scalable solutions.

Deployment Considerations

Real-world implementation will require addressing:

1. Cost of distributed deployment across heterogeneous smart city nodes.
2. System integration with existing storage architectures.
3. Maintenance of rule sets, especially for dynamic systems requiring policy updates.
4. A practical path may involve using containerized cache agents with centralized monitoring and decentralized execution to minimize update complexity.

Hybrid Caching Architectures

Blending policy-based logic with adaptive replacement algorithms (e.g., ARC or ML-driven LRU variants) may yield improved flexibility. For example:

1. Policy engines can handle energy and priority decisions.
2. Adaptive layers can manage temporal and frequency-based eviction.

Such hybrid systems could dynamically switch modes based on workload type, enhancing overall efficiency.

Edge-Aware Scalability Models

Scaling policy-based caching across hundreds or thousands of smart nodes (traffic lights, public displays, kiosks) requires:

1. Lightweight and modular implementations
2. Real-time synchronization mechanisms with low overhead
3. Zone-based clustering of cache nodes for better workload distribution

Research into policy rule abstraction and rule propagation models will be essential to ensure consistency and responsiveness in large-scale deployments.

Standardization and Benchmarks

There is a need for:

1. Open-source frameworks to benchmark policy-based caching in smart cities
2. Public datasets that mimic real urban workloads
3. Standard performance metrics for cross-study comparisons

Establishing these would allow wider adoption and reproducibility of results across research and industry applications.

VI. Conclusion

This study evaluated the impact of policy-based cache management on the performance and sustainability of smart city storage infrastructures. Through a rule-based caching framework designed to optimize data retention and eviction based on contextual factors such as data age, access frequency, priority, and energy cost the system demonstrated superior efficiency across key metrics. Compared to traditional (LRU) and advanced strategies (LFU, ARC), the policy-based system achieved: 85% cache hit rate, 70ms latency, 220MB/s throughput, 32% energy savings over the no-cache baseline. These results validate the approach as an effective solution for real-time responsiveness, energy optimization, and data prioritization in complex urban data environments.

Importantly, the proposed system supports broader goals of sustainable urban development by minimizing redundant data transfers and reducing power consumption directly contributing to SDG II (sustainable Cities and Communities) and SDG 13 (Climate Action). While some trade-offs exist particularly in computational overhead and deployment complexity, these are manageable with future improvements such as AI-driven policies, hybrid architectures, and lightweight distributed agents.

As smart cities scale in both complexity and data intensity, intelligent caching strategies like this one offer a scalable, sustainable, and adaptive framework for next generation urban infrastructure. Continued research into deployment frameworks, real-world validations, and AI-Policy integration will further enhance its viability and impact.

References

1. Alubady, R., Salman, M., & Mohamed, A. (2023). A Review of Modern Caching Strategies in Named Data Network: Overview, Classification, and Research Directions. *Telecommunication Systems*, 84(3), 1–46. Retrieved from <https://doi.org/10.1007/s11235-023-01015-3>.
2. Bello, H., Ige, A., & Ameyaw, M. (2024). Adaptive Machine Learning Models: Concepts for Real-Time Financial Fraud Prevention in Dynamic Environments. *World Journal of Advanced Engineering, Technology and Sciences*, 12(2), 021–034. Retrieved from <https://doi.org/10.30574/wjaets.2024.12.2.0266>.
3. Bilal, M., & Kang, S.-G. (2017). A Cache Management Scheme for Efficient Content Eviction and Replication in Cache Networks. *IEEE Access*, 5(16), 32720–32730. Retrieved from <https://doi.org/10.1109/ACCESS.2017.2669344>.
4. Chao, Y., & Han, R. (2025). A Hierarchical Cache Architecture-Oriented Cache Management Scheme for Information-Centric Networking. *Future Internet*, 17(1), 17. Retrieved from <https://doi.org/10.3390/fi17010017>.
5. Chidolue, O., Ohenhen, P., Umoh, ., Ngozichukwu, ., Fafure, A., & Ibekwe, P. (2024). Green Data Centers: Sustainable Practices For Energy-Efficient It Infrastructure. *Engineering Science & Technology Journal*, 5(1), 99–114. Retrieved from <https://doi.org/10.51594/estj/v5i1.730>.
6. Jangid, J. (2020). Efficient Training Data Caching for Deep Learning in Edge Computing Networks. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 7(8), 337–362. Retrieved from <https://doi.org/10.32628/CSEIT20631113>.
7. Khan, A. A., & Zakarya, M. (2021). Energy, Performance, and Cost Efficient Cloud Datacenters: A Survey. *Computer Science Review*, 40(100390), 1–27. Retrieved from <https://doi.org/10.1016/j.cosrev.2021.100390>.
8. Khedkar, V. (2024). The Carbon Conundrum: A Systematic Analysis of Environmental Impacts in Large-Scale Cloud Computing Infrastructure. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 10(6), 713–723. Retrieved from <https://doi.org/10.32628/CSEIT241061115>.
9. Krishna, K. (2025). Advancements in Cache Management: A Review of Machine Learning Innovations for Enhanced Performance and Security. *Frontiers in Artificial Intelligence*, 8, 1441250. Retrieved from <https://doi.org/10.3389/frai.2025.1441250>.
10. Levin, S. (2024). Unleashing Real-Time Analytics: A Comparative Study of In-Memory Computing vs. Traditional Disk-Based Systems. *Brazilian Journal of Science*, 3(5), 30–39. Retrieved from <https://doi.org/10.14295/bjs.v3i5.553>.
11. Mahomed, A. S., & Saha, A. K. (2025). Unleashing the Potential of 5G for Smart Cities: A Focus on Real-Time Digital Twin Integration. *Smart Cities*, 8(2), 70. Retrieved from <https://doi.org/10.3390/smartcities8020070>.
12. McKee, M. (2021). Policy-Based Access Controls. *IDPro Body of Knowledge*, 1, 1–3. Retrieved from <https://doi.org/10.55621/idpro.61>.
13. Narain, S. (2024). Integrating Smart City Technologies for Enhanced Urban Sustainability. *Journal of Sustainable Solutions*, 1(3), 13–17. Retrieved from <https://doi.org/10.36676/j.sust.sol.v1.i3.15>.
14. Orikpete, O., Fawole, A., & Ewim, D. (2023). Impact of Data Centers on Climate Change: A Review of Energy Efficient Strategies. *The Journal of Engineering and Exact Sciences*, 9(6), 16397–16409. Retrieved from <https://doi.org/10.18540/jeecevl9iss6pp16397-01e>.
15. Qaiser, F., Harthy, K., Hussain, M., Frnda, J., Amin, R., Gantassi, R., & Zakaria, M. (2025). Classifications and Analysis of Caching Strategies in Information-Centric Networking for Modern Communication Systems. *Engineering Reports*, 7(0), 1–33. Retrieved from <https://doi.org/10.1002/eng2.70005>.
16. Ullah, A., Syed, ., Anwar, M., Li, J., Nadeem, L., Mahmood, T., Rehman, A., Syed, M., Anwar, ., & Saba, T. (2023). Smart Cities: The Role of Internet of Things and Machine Learning in Realizing a Data-Centric Smart Environment. *Complex & Intelligent Systems*, 10(0), 1–31. Retrieved from <https://doi.org/10.1007/s40747-023-01175-4>.

17. Shahrabani, M. M. N., & Apanaviciene, R. (2025). Evaluation of Smart Building Integration into a Smart City by Applying Machine Learning Techniques. *Buildings*, 15(12), 2031. Retrieved from <https://doi.org/10.3390/buildings15122031>.
18. Talebkhah, M., Sali, A., Marjani, M., Gordan, M., Hashim, S., & Rokhani, F. Z. (2021). IoT and Big Data Applications in Smart Cities: Recent Advances, Challenges, and Critical Issues. *IEEE Access*, 9(35), 126383–126412. Retrieved from <https://doi.org/10.1109/ACCESS.2021.3070905>.
19. Thangam, D., Haritha, M., Ramesh, R., Ramakrishna, G. N., Muddasir, N., Khan, A., Booshan, S., Booshan, B., & Ganesh, R. (2024). Impact of Data Centers on Power Consumption, Climate Change, and Sustainability. <https://doi.org/10.4018/979-8-3693-1552-1.ch004>.
20. Trigka, M., & Dritsas, E. (2025). Edge and Cloud Computing in Smart Cities. *Future Internet*, 17(3), 118. Retrieved from <https://doi.org/10.3390/fi17030118>.
21. Weerasinghe, S., Zaslavsky, A., Loke, S., Hassani, A., Abken, A., & Medvedev, A. (2022). From Traditional Adaptive Data Caching to Adaptive Context Caching: A Survey. *ArXiv preprint arXiv:2211.11259*.
22. Yildirim, F., Yalman, Y., Bayındır, K. Ç., & Terciyanlı, E. (2025). Comprehensive Review of Edge Computing for Power Systems: State of the Art, Architecture, and Applications. *Applied Sciences*, 15(8), 4592. Retrieved from <https://doi.org/10.3390/app15084592>.
23. Yu, M., Li, R., & Chen, Y.-W. (2020). A Cache Replacement Policy Based on Multi-factors for Named Data Networking. *Computers, Materials & Continua*, 65(4), 321–336. Retrieved from <https://doi.org/10.32604/cmc.2020.010831>.
24. Zaman, M., Abdelwahed, S., & Zohrabi, N. (2024). A Review of IoT-Based Smart City Development and Management. *Smart Cities*, 7(3), 1462–1501. Retrieved from <https://doi.org/10.3390/smartcities7030061>.
25. Zeng, F., Zhang, K., Wu, L., & Wu, J. (2022). Efficient Caching in Vehicular Edge Computing Based on Edge-Cloud Collaboration. *IEEE Transactions on Vehicular Technology* (Early Access). <https://doi.org/10.1109/TVT.2022.3213130>.
26. Zhang, M., Tang, Q., Kim, J.-G., Burgstaller, B., & Kim, S.-D. (2023). Adaptive Regression Prefetching Algorithm by Using Big Data Application Characteristics. *Applied Sciences*, 13(7), 4436. Retrieved from <https://doi.org/10.3390/app13074436>.