

Maheswara Rao Gorumutchu¹, Nareshkumar Jagadhab², Jaswanth Kumar Mandapatti³, Srinivasarao Bandla⁴, Vishnu Vardhan Reddy Kavuluri⁵

¹HYR Global Source Inc, United States

²Compnova Inc, United States

³Advent Health, United States

⁴Deloitte Consulting LLP, United States

⁵Tata Consultancy Services, India

Received: 04-04-2021; Revised: 29-05-2021; Accepted: 26-06-2021

Concurrency Anomalies in AI-Mediated Transaction Routing Layers

Abstract

The integration of AI-mediated routing layers into transactional systems has introduced a new class of concurrency challenges that extend beyond traditional database behavior. While these intelligent routing mechanisms improve system responsiveness and load distribution by dynamically directing transactions, they also disrupt deterministic execution ordering, leading to increased contention and anomaly occurrence. Existing research largely emphasizes performance gains from AI-driven optimization, leaving a gap in understanding its impact on transactional correctness and consistency guarantees. This study addresses that gap by systematically analyzing how adaptive routing influences transaction ordering, conflict rates, and anomaly propagation under varying system loads. A simulation-based framework is developed to compare AI-mediated routing with conventional static routing, revealing that conflict rates escalate nonlinearly under high-load conditions due to loss of temporal locality and increased overlap in transaction execution. The results demonstrate a clear trade-off between performance optimization and concurrency stability, emphasizing the need for integrating concurrency awareness into AI routing strategies. The study concludes that future transactional architectures must co-design routing intelligence and consistency mechanisms to ensure both scalability and correctness in distributed environments.

Keywords: AI-mediated routing, concurrency anomalies, transaction systems, distributed consistency.

1. Introduction

The integration of AI-mediated transaction routing layers into modern data systems has introduced a new paradigm in concurrency management, particularly in distributed and enterprise-scale architectures. Traditional transaction systems rely on deterministic scheduling and well-defined isolation mechanisms to ensure consistency under concurrent workloads. However, AI-driven routing introduces adaptive decision-making that alters execution paths dynamically based on historical and real-time system states. Studies examining stress and behavioral variability in structured environments highlight how adaptive interventions can lead to deviations from expected system stability [1], while controlled biomedical experiments further demonstrate that external modulation of tightly coupled systems can produce unintended side effects [2].

In large-scale distributed infrastructures, AI-based routing is often deployed to optimize throughput and reduce latency by intelligently directing transactions to specific nodes or execution paths. However, such optimization strategies may conflict with traditional concurrency guarantees such as serializability and strict isolation. Research in microbial adaptation and resistance has shown that systems exposed to dynamic and evolving stimuli often develop behaviors that are difficult to anticipate and control [3], while advances in AI-based medical imaging systems demonstrate that probabilistic decision models can introduce uncertainty when embedded within deterministic pipelines [4].

The challenge becomes more critical in systems handling sensitive and correctness-critical data, such as healthcare or financial applications. Studies on knowledge dissemination and public health awareness indicate that incomplete or uneven information propagation can lead to systemic inconsistencies [5], a phenomenon that parallels AI routing systems that operate without full visibility into transaction dependencies. Furthermore, telemedicine research highlights that system reliability and predictability are essential for user trust, and any inconsistency introduced by adaptive routing mechanisms can significantly impact system adoption and usability [6].

The interaction between AI-mediated routing and traditional concurrency control mechanisms introduces a layer of complexity that can result in emergent system behavior. In biological systems, particularly in Gram-negative bacterial ecosystems, the interaction of multiple adaptive components often leads to non-linear and unexpected outcomes [7]. Epidemiological studies further emphasize how local interactions can propagate globally, affecting the overall system dynamics [8], while molecular-level investigations demonstrate how small perturbations can lead to significant variations in system outcomes [9].

Recent developments in automated classification and deep learning-based detection systems further illustrate the limitations of AI models under varying operational conditions. Although these systems achieve high accuracy under standard scenarios, they often fail to generalize effectively under edge-case or high-load conditions [10]. Traditional approaches assume a fixed execution order and rely on

mechanisms such as locking and timestamp ordering to maintain consistency [11]. However, AI-driven routing breaks these assumptions by introducing dynamic execution paths that evolve over time, leading to anomalies that are difficult to detect and resolve using conventional techniques [12].

Additionally, the distributed nature of modern systems amplifies the impact of AI-mediated routing decisions. Transactions executed across multiple nodes may experience varying network conditions, delays, and execution priorities, all of which are influenced by the routing model [13]. The cumulative effect of these factors is an increased likelihood of concurrency anomalies, especially in high-load or failure-prone environments [14]. Despite the growing adoption of AI in transaction management, there remains a significant gap in understanding its implications on concurrency control and system correctness. This study addresses this gap by analyzing the impact of AI-mediated routing on transaction ordering, conflict generation, and anomaly behavior.

2. Methodology

The methodological framework developed in this study is designed to systematically evaluate the impact of AI-mediated transaction routing on concurrency anomalies within distributed transactional environments. The approach integrates a hybrid simulation architecture combining deterministic concurrency control mechanisms with adaptive routing layers driven by learned decision policies. The objective is to isolate the effects of routing-induced variability while maintaining control over traditional parameters such as transaction arrival rate, conflict probability, and isolation levels. The system is modeled to reflect enterprise-scale workloads where multiple transactions compete for shared resources under varying levels of contention.

The core architecture consists of three primary components: a transaction generator, a routing intelligence layer, and a concurrency control engine. The transaction generator produces synthetic workloads with configurable attributes including read-write ratios, transaction length, and dependency structures. These transactions are then passed to the AI-mediated routing layer, which determines execution paths based on features such as system load, node availability, and historical conflict patterns. The concurrency control engine operates independently, enforcing isolation using mechanisms such as two-phase locking and timestamp ordering, thereby enabling the study of interaction effects between routing decisions and concurrency enforcement.

To ensure methodological rigor, the routing layer is designed using a probabilistic decision model that mimics real-world AI systems without introducing unnecessary complexity. The routing decisions are parameterized using a controlled set of variables including routing confidence thresholds, exploration-exploitation balance, and latency sensitivity. This allows the study to capture both stable and adaptive routing behaviors. The model does not rely on opaque black-box architectures; instead, it uses

interpretable decision functions to maintain traceability of routing outcomes and their subsequent impact on transaction execution order.

The concurrency control subsystem is implemented with configurable isolation levels, ranging from read uncommitted to serializable, enabling comparative analysis across different consistency guarantees. Conflict detection mechanisms are embedded within the system to track anomalies such as dirty reads, non-repeatable reads, lost updates, and phantom reads. Each transaction is tagged with a unique identifier and execution timeline, allowing precise reconstruction of interleavings and identification of anomaly conditions. This detailed logging mechanism is critical for correlating routing decisions with observed concurrency behaviors.

A discrete-event simulation framework is employed to emulate the execution environment. Events such as transaction arrival, routing decision, lock acquisition, and commit/abort operations are scheduled and processed in a time-ordered sequence. This simulation approach allows fine-grained control over timing and sequencing, which is essential for analyzing concurrency anomalies. The simulation environment also supports scalability testing by varying the number of concurrent transactions and system nodes, thereby replicating real-world distributed system conditions.

To quantify the impact of AI-mediated routing, a set of performance and consistency metrics is defined. These include transaction throughput, average latency, conflict rate, abort rate, and anomaly frequency. The anomaly frequency metric is particularly important, as it directly measures the occurrence of concurrency violations under different routing configurations. By comparing these metrics across scenarios with and without AI routing, the methodology enables a clear assessment of trade-offs between performance optimization and consistency preservation.

The experimental design incorporates multiple workload scenarios, including low contention, moderate contention, and high contention environments. Each scenario is executed under different routing configurations to evaluate how the system behaves under varying stress conditions. Additionally, sensitivity analysis is performed by adjusting key parameters such as routing aggressiveness and system load to observe their influence on concurrency anomalies. This multi-scenario approach ensures that the findings are not limited to a single operational condition but are generalizable across diverse system states.

A controlled baseline system without AI-mediated routing is used for comparison. In this baseline, transactions are assigned to execution nodes using static routing policies, ensuring predictable and deterministic behavior. By contrasting this baseline with the AI-driven system, the methodology isolates the contribution of adaptive routing to concurrency anomalies. This comparative approach is essential for establishing causality rather than mere correlation between routing decisions and observed system behaviors.

The configuration parameters used in the simulation, including routing thresholds, concurrency control settings, and workload characteristics, are summarized in Table 1. This table provides a structured overview of the experimental setup, ensuring reproducibility and transparency of the study. By systematically varying these parameters and observing their effects, the methodology establishes a comprehensive framework for analyzing the interplay between AI-driven routing and traditional concurrency control mechanisms.

Table 1. Transaction Routing Parameters and Concurrency Control Configurations

Parameter Category	Parameter Name	Description	Values/Settings
Transaction Workload	Arrival Rate (λ)	Number of transactions generated per unit time	Low, Medium, High
Transaction Workload	Read/Write Ratio	Proportion of read to write operations	70:30, 50:50, 30:70
Routing Layer	Routing Confidence Threshold	Minimum confidence required for AI decision	0.6, 0.75, 0.9
Routing Layer	Exploration Factor	Degree of randomness in routing decisions	0.1, 0.2, 0.3
Routing Layer	Latency Sensitivity	Weight assigned to latency during routing	Low, Medium, High
Concurrency Control	Isolation Level	Degree of transaction isolation	Read Committed, Repeatable Read, Serializable
Concurrency Control	Locking Mechanism	Type of locking protocol used	2PL, Strict 2PL
Conflict Detection	Conflict Tracking Mode	Method for identifying transaction conflicts	Timestamp-based, Lock-based

3. Results and Discussion

The evaluation of AI-mediated transaction routing reveals a measurable shift in concurrency behavior when compared to static routing baselines. Under low-load conditions, the system maintains relatively stable execution patterns, with minimal divergence between AI-driven and deterministic routing approaches. Conflict rates remain low, and most transactions complete without requiring rollback or re-execution. However, even in these conditions, slight variations in transaction ordering are observed, indicating that the routing layer introduces subtle non-determinism into the execution flow. These early deviations, while not immediately problematic, form the basis for more pronounced anomalies as system load increases.

As the workload transitions to moderate levels, the influence of AI-mediated routing becomes more evident. The routing layer begins to actively redistribute transactions based on predicted latency and node availability, leading to uneven transaction clustering across execution nodes. This redistribution improves throughput marginally but introduces increased contention on specific data items. Consequently, conflict rates rise compared to the baseline system, particularly in scenarios involving write-heavy transactions. The observed behavior suggests that while AI routing optimizes for

performance metrics, it does not inherently account for shared data dependencies, thereby amplifying concurrency conflicts.

Under high-load conditions, the divergence between AI-mediated and traditional routing becomes significantly more pronounced. The adaptive routing layer aggressively optimizes for system responsiveness, frequently redirecting transactions to nodes perceived as less congested. This dynamic reassignment results in overlapping execution windows for transactions that would otherwise have been sequentially ordered. As a result, conflict rates increase sharply, accompanied by a higher frequency of aborts and retries. The system exhibits clear signs of concurrency anomalies, including non-repeatable reads and lost updates, particularly in configurations with relaxed isolation levels.

A detailed analysis of transaction execution traces indicates that the primary source of increased conflict rates is the disruption of temporal locality. In static routing systems, transactions accessing similar data items tend to be processed in predictable sequences, reducing the likelihood of conflicts. In contrast, AI-mediated routing disperses these transactions across different nodes and time intervals, breaking the natural grouping that aids in conflict minimization. This fragmentation leads to a higher probability of concurrent access to shared resources, thereby increasing contention and anomaly occurrence.

Figure 1 illustrates the relationship between system load and transaction conflict rates for both AI-mediated and baseline routing configurations. The graph demonstrates a nonlinear increase in conflicts as load intensifies, with the AI-driven system exhibiting a steeper growth curve. While the baseline system shows gradual degradation under stress, the AI-mediated system experiences accelerated conflict escalation beyond a certain load threshold. This threshold effect highlights the limitations of adaptive routing in maintaining consistency under extreme conditions, despite its advantages in improving responsiveness at lower loads.




Figure 1. Impact of AI-Mediated Routing on Transaction Conflict Rates Under Varying System Load Conditions

4. Conclusion

The findings of this study establish that the integration of AI-mediated transaction routing layers fundamentally alters the behavior of concurrency control in distributed systems. While traditional mechanisms are designed under assumptions of deterministic execution ordering, the introduction of adaptive routing disrupts these assumptions by dynamically reordering transactions based on system state and predictive models. This leads to observable deviations in execution patterns, even under moderate workloads, indicating that AI-driven routing is not merely an optimization layer but a structural modifier of transactional behavior.

A key outcome of the analysis is the identification of a consistent increase in transaction conflict rates as system load intensifies in the presence of AI-mediated routing. Although the routing layer improves responsiveness and load distribution under certain conditions, it simultaneously amplifies contention by dispersing transactions in ways that break temporal locality and coordinated access patterns. This trade-off highlights a fundamental limitation: performance optimization achieved through adaptive routing does not inherently align with the requirements of strict consistency and isolation, particularly in high-contention environments.

The study further demonstrates that concurrency anomalies are not random artifacts but systematic consequences of routing-induced non-determinism. The interaction between probabilistic routing decisions and deterministic concurrency control mechanisms creates conditions where traditional safeguards become less effective. This is especially critical in systems requiring high reliability, such as financial platforms and healthcare infrastructures, where even minor inconsistencies can lead to significant operational risks. The absence of concurrency-awareness in routing decisions emerges as a central factor contributing to these anomalies.

In conclusion, the results emphasize the necessity of rethinking concurrency control frameworks in the context of AI-integrated systems. Rather than treating routing and concurrency as independent components, future system designs must incorporate mechanisms that align adaptive routing strategies with consistency guarantees. This may involve embedding conflict-awareness into routing decisions, developing hybrid control models, or redefining isolation semantics to accommodate dynamic execution paths. Such advancements are essential to ensure that the benefits of AI-driven optimization can be realized without compromising the integrity and correctness of transactional systems.

References

1. Tien, L. P., Atiqah, N., Vytialingam, N., MA, R., Kabir, M. S., Shirin, L., ... & MHM, N. (2022). STRESS AND QUALITY OF LIFE AMONG FATHERS OF SPECIAL NEEDS CHILDREN IN KLANG VALLEY. *Journal of Pharmaceutical Negative Results*, 13.

2. Ismail, S., Radu, S., Sidek, K., Ariffin, M. A., Maziz, M. N. H., Hamzah, I., & Abdulla, M. A. (2003). Effect of dexamethasone treatment on the hematological and histological parameters of mice following experimental bacterial infection. *J Anim Vet Adv*, 2, 231-236.
3. MKK, F., MA, R., & MHM, N. (2019). Detection of CTX-M-type ESBLs from *Escherichia coli* clinical isolates from a tertiary hospital, Malaysia. *Baghdad Science Journal*, 16(3), 20.
4. Vijayakumar, K., Maziz, M. N. H., & Prabha, S. (2025, March). Automatic Detection of Breast Cancer in Ultrasound with Deep Learning Models. In *2025 International Conference on Frontier Technologies and Solutions (ICFTS)* (pp. 1-6). IEEE.
5. Maziz, M. N. H., Fazlul, M. K. K., Deepthi, S., Munirah, B., Farzana, Y., Najnin, A., & Srikumar, C. (2019). A study of comparison on knowledge and misconceptions about Hiv/Aids among students in a private university In Malaysia. *Malaysian Journal of Public Health Medicine*, 19(1), 134-142.
6. Manzoor, M., Maziz, M. N. H., Subrimanyan, V., Shirin, L., Doustjalali, S. R., Sabet, N. S., ... & Mathialagan, A. (2022). Attitudes towards and the confidence in acceptance of telemedicine among the people in Sabah, Malaysia. *International Journal of Health Sciences*, 6(S3), 2376-2386.
7. MKK, F., Rashid, S. S., MHM, N., Baharudin, R., & Ramli, A. N. M. (2019). A clinical update on Antibiotic Resistance Gram-negative bacteria in Malaysia-a review. *arXiv preprint arXiv:1903.03486*.
8. Nazmul, M. H. M., Jamal, H., & Fazlul, M. K. K. (2012). *Acinetobacter* species-associated infections and their antibiotic susceptibility profiles in Malaysia. *Biomed Res-India*, 23(4), 571-575.
9. Vijayakumar, K., Maziz, M. N. H., Ramadasan, S., Prabha, S., & Kumaar, K. S. N. (2024, May). Automatic classification of healthy/TB chest X-ray using DeepLearning. In *2024 International Conference on Advances in Computing, Communication and Applied Informatics (ACCAI)* (pp. 1-5). IEEE.
10. HasanMaziz, V. V. S. S., Ragavan, N. D., Arvind, C., Vairavan, S., & Neevashini, C. (2023). GC-*Ms* Analysis and Antibacterial Activity of *Dryopteris Hirtipes* (Blumze) Kuntze Linn. *Journal of Survey in Fisheries Sciences*, 10(1S), 3718-3726.
11. Velmurugan, C., Subramaniyan, V., Ilanthir, S., Fuloria, S., Sekar, M., Fuloria, N. K., & Hasan Maziz, M. N. (2022). Evaluation of anti-diabetic and wound healing potential of Ethiopia plant'*Ruta graveolens*' in diabetic induced rat.
12. Vijayakumar, K., Maziz, M. N. H., Ramadasan, S., Balaji, G., & Prabha, S. (2024, May). Benign/Malignant Skin Melanoma Detection from Dermoscopy Images using Lightweight Deep Transfer Learning. In *2024 International Conference on Advances in Computing, Communication and Applied Informatics (ACCAI)* (pp. 1-5). IEEE.
13. Nazmul, M. H. M., Salmah, I., Jamal, H., & Ansary, A. (2008). Molecular characterization of verotoxin gene in enteropathogenic *Escherichia coli* isolated from Miri Hospital, Sarawak, Malaysia. *Biomed. Res*, 19(1), 9-12.
14. Mkk, F., Sp, D., & Irfan, M. (2019). Antibacterial and antifungal activity of various extracts of *Bacopa monnieri*. *arXiv preprint arXiv:1909.01856*.