



Article

Self-healing IoT infrastructure for intelligent transportation systems: a multi-city comparative analysis

Si Liu, Midhun Chakkaravarthy*

School of AI Computing and Multimedia, Lincoln University College, Malaysia

ARTICLE INFO	ABSTRACT
<p><i>Article history:</i> Received 18 December 2025 Received in revised form 10 February 2026 Accepted 19 March 2026</p> <p>Keywords: Self-healing systems, IoT infrastructure, Intelligent transportation systems, Predictive maintenance, Distributed consensus</p> <p>*Corresponding author Email address: midhun@lincoln.edu.my</p> <p>DOI: 10.55670/fpjl.fdtai.2.1.2</p>	<p>Intelligent transportation systems critically depend on reliable real-time data from thousands of distributed IoT sensors, yet current management frameworks lack autonomous recovery mechanisms to maintain service continuity during device failures or network disruptions. This paper presents a self-healing infrastructure architecture specifically designed for the transportation domain, where service interruptions can cascade into traffic congestion and safety hazards. The architecture integrates three complementary mechanisms: predictive health monitoring, which uses time-series anomaly detection to identify failing devices before complete failure; preemptive workload migration, which redistributes critical sensing tasks to neighboring devices based on predicted failures; and distributed consensus protocols, which enable rapid reconfiguration without centralized coordination. Unlike existing reactive approaches that respond to failures after service degradation occurs, our proactive strategy maintains service quality by anticipating and preventing disruptions. A key innovation lies in the domain-specific failure prediction models trained on operational patterns unique to transportation applications, incorporating factors such as vehicle vibration exposure, environmental stress, and communication interference patterns. Comparative analysis of multiple metropolitan transportation networks with varying levels of infrastructure maturity reveals that self-healing capabilities significantly reduce the frequency of service interruptions and recovery time compared to operator-managed systems. The findings demonstrate that transportation authorities can deploy larger sensor networks with fewer maintenance personnel by leveraging autonomous resilience mechanisms. This work provides practical guidelines for deploying fault-tolerant IoT systems in mission-critical urban services where reliability directly impacts public safety and urban mobility.</p>

1. Introduction

The operation of intelligent transportation systems is highly dependent on the presence of numerous IoT devices deployed across the city's road infrastructure. These devices collect data in real time, perform vehicle detection and classification, monitor environmental conditions, and provide inputs to signal control systems, among other functions [1]. In the city-scale intelligent transportation system, roadside devices, in the form of inductive loop detectors, video surveillance, and vehicle-to-infrastructure terminals, combine to form a sensor network that has thousands, and in some instances, tens of thousands of devices. As the sensor network continues to grow, the effort required to sustain service continuity has increased nonlinearly, and managing these devices has eventually become a significant engineering challenge for the overall system's reliability [2]. The majority of existing management frameworks for IoT networks, especially those designed for transportation management, still rely on reactive workflows to handle device faults. Specifically, when detected on a device, an alert is sent; after that, the type of fault is diagnosed, and manual repair or replacement of the device is arranged. Although a workflow may be acceptable for other types of IoT networks, it is still a major concern for transportation management networks. For instance, when a fault is detected in

a traffic sensor, it is not an isolated incident, and it is unacceptable that the failure of a single device may trigger other effects, leading to desynchronization of local signals and the creation of blind spots that may result in traffic congestion [3]. Such public safety issues necessitate ensuring that fault recovery is performed substantially faster than in other IoT networks. Although autonomous recovery mechanisms are well researched for distributed networks, especially those involving software layers of cloud microservices and data center infrastructure, their application to physical IoT devices in transportation sensor networks remains limited. To address the aforementioned problem, this study proposes a self-healing infrastructure architecture designed for the transportation domain. This proposed architecture incorporates three mechanisms that operate sequentially [4]. First, the predictive health monitoring mechanism will use time-series anomaly detection to identify trends in performance degradation before the device fails. Second, the preemptive workload migration mechanism will migrate critical sensing operations from potentially failing devices to functionally compatible neighboring devices based on predictive health monitoring outcomes. Finally, the distributed consensus protocol will enable nodes in the affected region to quickly reach agreement on the network reconfiguration plan. Unlike other conventional reactive architectures, the key differentiating feature of this proposed architecture is its ability to maintain quality of service through anticipation and prevention, enabled by a domain-specific failure prediction model that incorporates transportation-specific stress factors, including vehicle vibration, environmental degradation, and radio-frequency interference.

2. Related works

Intelligent transportation infrastructure today utilizes a variety of heterogeneous sensing devices, such as inductive loop detectors, video cameras, dedicated short-range communication devices, and environmental sensors, deployed to enable real-time traffic perception and management. However, the fault management approaches employed by these infrastructures remain predominantly reactive, using Simple Network Management Protocol threshold-based alerting and Message Queuing Telemetry Transport heartbeat-based fault detection approaches to identify faults that have already occurred [5], along with traditional preventive maintenance approaches that involve maintenance at predetermined time intervals irrespective of the actual condition of the devices [6]. As a result, fault identification remains entirely dependent on symptoms that manifest after the fault has occurred, leading to mean time to recover values that often exceed several hours, which in turn adversely affects traffic safety. This requirement to reduce mean time to recover values has led to an investigation into approaches that utilize self-healing fault management, a domain within distributed systems. The paradigm of autonomic computing formalized self-management through the Monitor-Analyze-Plan-Execute loop, which allows software systems to detect anomalies, reason about their root causes, and execute appropriate countermeasures with minimal human intervention [7]. Container orchestration platforms have operationalized this concept through periodic health checks and auto-rescheduling of failed instances [8], and recent research has extended this capability to edge computing nodes to address software-layer failures, such as process crashes and memory leaks [9]. However, a major limitation of the above approaches is that they are limited to the software and service layers and do not account for the gradual degradation of physical devices in outdoor environments. If device failures could be predicted in advance, self-healing could shift from a traditional reactive paradigm to a proactive one.

Time-series-based anomaly detection offers a promising way forward. Deep learning-based approaches such as long short-term memory autoencoders and one-dimensional convolution neural networks have been found to be highly effective in identifying patterns of temporal degradation in IoT-based telemetry streams [10]. However, the applicability of this approach in transportation IoT networks is hampered by two main issues. The current models are designed for general IoT-based telemetry and do not account for transportation-specific stressors such as vehicle-induced vibration patterns, thermal degradation from radiation from asphalt-based road surfaces, and radio-frequency interference from vehicular communication. Even when accurate predictions are made, they do not account for autonomous network reconfiguration and still require human intervention during workload migration. This study aims to bridge this gap by developing network-specific prediction models that factor in transportation-specific stress factors and are designed for autonomous network reconfiguration.

3. The proposed architecture

3.1 System overview and workflow

In particular, the proposed self-healing architecture is based on a three-layered approach. The Sensing Layer comprises physical IoT devices typically installed along the road network, such as loop detectors, cameras, roadside units, and environmental sensors. Moreover, the Edge Processing Layer includes local computation and monitoring agents that are typically executed on resource-constrained hardware and are co-located with sensing nodes. Finally, the Self-Healing Coordination Layer manages the three fundamental mechanisms that collectively enable autonomous fault resilience. Three design principles are considered when developing the proposed architecture. First, decentralization distributes decision-making authority among local clusters of nodes, effectively eliminating the single-point-of-failure coordinator. Second, domain awareness is conducted to ensure that models and thresholds are calibrated against known operational patterns specific to transportation environments. Finally, proactive intervention shifts the maintenance paradigm from detect-report-repair to predict-prevent-adapt.

The three modules work together in a pipeline fashion. The predictive health monitoring module serves as the trigger: the device's condition is continuously monitored, and alerts are sent when degradation trends are identified. The preemptive workload migration module is the executor that develops a plan to migrate workloads to the affected nodes. Finally, the distributed consensus protocol acts as the coordinator, with consensus reached among the affected neighbors before the execution process is completed, while the whole process works independently, as depicted in [Figure 1](#).

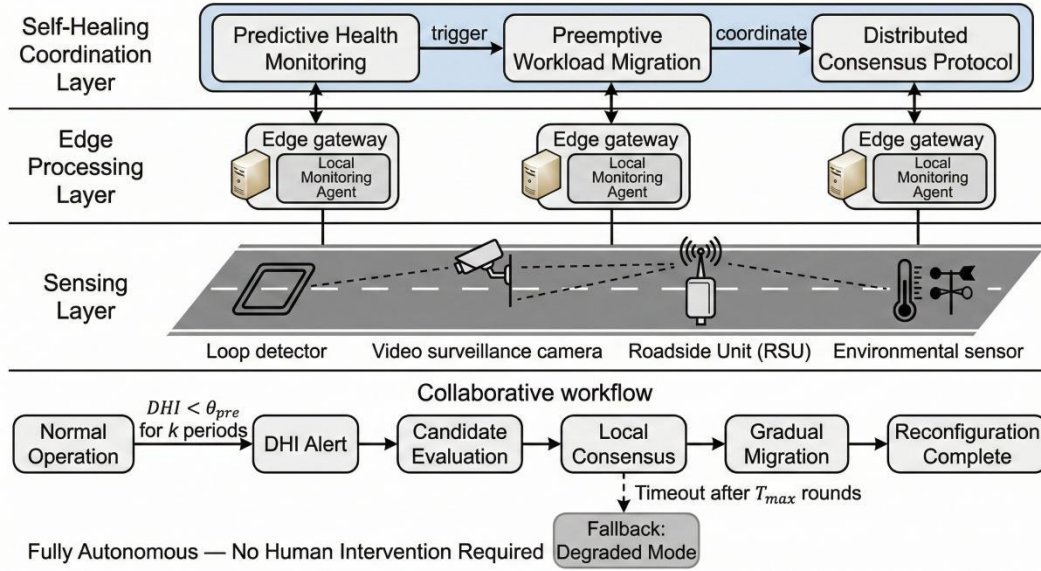


Figure 1. Self-healing architecture overview and autonomous workflow

3.2 Predictive health monitoring

The monitoring module receives three types of telemetry signals: operational, environmental, and communication, each reflecting distinct stress factors. The operational signals comprise the processor's temperature, memory, network latency, and supply voltage, while the environmental signals include the vibration intensity of road surfaces from moving vehicles and temperature and humidity changes caused by asphalt radiation. The communication signals encompass signal strength, packet loss, and retransmission rate, which are affected by radio-frequency interference among dense vehicular communication [11]. These heterogeneous inputs are fused into a Device Health Index mapped onto $[0,1]$ through a sliding time window. The prediction is formalized as:

$$DHI(t+\Delta) = f(X_{operational}(t), X_{environmental}(t), X_{communication}(t); \theta) \quad (1)$$

where θ denotes model parameters trained on historical ITS device failure logs and Δ represents the prediction horizon. A lightweight one-dimensional convolutional neural network serves as f , selected for its capacity to extract local temporal patterns from multivariate telemetry while remaining deployable on edge hardware with limited computational budgets [12].

To mitigate false positives, migration is not triggered immediately when DHI drops below θ_{pre} . The module instead enters a confirmation window requiring DHI to remain below θ_{pre} for k consecutive sampling periods, filtering transient fluctuations and ensuring only sustained degradation trends initiate migration.

3.3 Preemptive workload migration

The migration module activates once a sustained DHI decline has been confirmed through the observation window. Candidate receiver nodes are evaluated using a multi-objective scoring function.

$$S_{candidate} = \alpha \cdot Proximity + \beta \cdot Compatibility + \gamma \cdot (1 - Load) + \delta \cdot DHI_{candidate} \quad (2)$$

where $Proximity$ quantifies geographic closeness, $Compatibility$ measures functional overlap in sensing capabilities, $Load$ represents current resource utilization, and $DHI_{candidate}$ reflects the candidate's own health status. The constraint $\alpha + \beta + \gamma + \delta = 1$ is imposed, with optimal weights determined through grid search on a validation set derived from historical deployment data.

The redistribution process is a three-stage procedure. Specifically, the failing node is first decomposed into sub-tasks that are transferable. The workload is gradually transferred instead of a hard switch. Finally, coverage validation is performed to ensure that no blind spots are created in the monitoring process [13]. Two mechanisms are incorporated to prevent destabilization. A receiving node with utilization above 0.85 will reject incoming requests. A limit of three migrations per cluster prevents correlated failure scenarios, such as severe weather affecting multiple nodes, which could result in multiple requests being made to neighboring devices.

3.4 Distributed consensus protocol

Once the migration module has created its redistribution plan, agreement among all affected neighbors is achieved through the consensus protocol before execution. Sequential dependency ensures that unilateral reconfiguration does not interfere with ongoing sensing operations. The protocol is based on the leader election and log replication principles used in Raft. It has two adaptations for IoT resource constraints. The first modification is that consensus is limited to a neighborhood, not the entire network. The number of rounds scales as:

$O(\log N_{local})$

(3)

where N_{local} denotes participating nodes in the local neighborhood. If agreement is not achieved within T_{max} voting rounds, the protocol activates a pre-cached degraded-mode configuration maintaining basic sensing coverage at reduced resolution, ensuring service continuity even under network partitions or node unresponsiveness. The process of the proposed consensus protocol is illustrated in Figure 2.

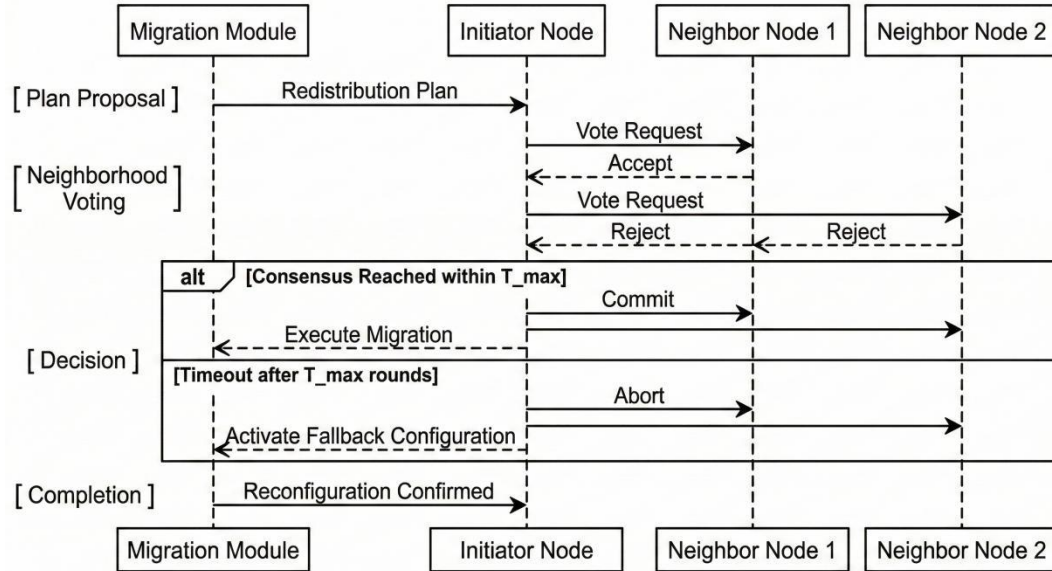


Figure 2. Consensus protocol sequence diagram

4. Experiment

4.1 Simulation environment

The experimental assessment is performed in a discrete-event simulation framework developed in Python and integrated with the NS-3 network simulator to model sensor communication devices [14]. The device resource profile is calibrated based on the specifications of the Raspberry Pi 4 and NVIDIA Jetson Nano platforms. Fault injection is performed using three statistical models derived from available ITS maintenance reports. Specifically, the three models used to generate hardware degradation faults include a Weibull distribution model that represents the increased probability of failure with device age. Disruptions in the network are simulated using a Poisson distribution model representing random communication outages. Extreme environmental conditions are simulated using a seasonal peak distribution model that represents climate conditions.

4.2 Multi-city configuration

There are three city infrastructures that signify three levels of maturity in the overall system, as shown in Table 1. City-A represents an advanced system with 2,000 sensor nodes, a fully meshed, redundant network, and around-the-clock maintenance staff. City-B indicates a mid-level maturity system with 800 nodes, a redundant network, and staff who operate on a scheduled shift basis. Lastly, City-C denotes an early-stage system with 300 nodes, a sparse topology, and a minimal tree topology, with staff available on a call basis. Besides the number of nodes and staff, the three cities differ in that the average age of devices and the climate intensity in City-C are the oldest and most extreme, with a hot, dusty environment. Simulation of each of the three cities is conducted over a 180-day period.

Table 1. Multi-City Deployment Parameters

Parameter	City-A (High)	City-B (Medium)	City-C (Low)
Sensor nodes	~2,000	~800	~300
Network redundancy	Full mesh	Partial	Sparse tree
Maintenance team	24/7 dedicated	Scheduled shifts	On-call only
Avg. device age	2.1 yrs	3.4 yrs	4.8 yrs
Climate stress	Moderate	High (humid)	Extreme (hot/dusty)

4.3 Baseline methods and selection rationale

The three baselines provide a progressive capability gradient from completely passive management to completely general-purpose autonomic management. Reactive-Manual (B1) is the current norm in most transport agencies, in which fault detection, diagnosis, and recovery rely entirely on manual operator intervention. Threshold-Auto (B2) adds the simplest form of automation by restarting devices when parameters exceed certain thresholds, but it lacks predictive capabilities. MAPE-K Autonomic (B3) adds analytical and planning capabilities through the Monitor-Analyze-Plan-Execute loop, providing rule-based self-management without the need for domain-specific prediction. The proposed architecture adds the missing transport-specific prediction, proactive migration, and consensus to the gradient, and the results will show the marginal contribution of each tier of the gradient capability.

4.4 Evaluation metrics

The four metrics represent different dimensions of the four strategies. Service Interruption Frequency (SIF) is the number of interruptions per 1,000 node days. Mean Time to Recovery (MTTR) is the average time in minutes from the occurrence of a fault until the service is restored. Service Availability Rate (SAR) is the percentage of time the sensors are operational and delivering data. Furthermore, the Maintenance Personnel Ratio (MPR) is the number of full-time personnel required per 100 nodes.

4.5 Implementation details

The predictive monitoring module is based on a three-layer one-dimensional convolutional neural network. The historical data is split into a set of training data, a set of validation data, and a set of test data in a ratio of 60:20:20. The Device Health Index threshold θ_{pre} is set to 0.35 by maximizing the F1-score using receiver operating characteristic curve analysis on the validation data set. The confirmation window k is set to 3 sampling periods. The weights in Equation 2 are determined through a grid search on the validation data set. The timeout parameter T_{max} is set to 5 rounds. All experiments are repeated 30 times with different random seeds. Statistical significance is determined using the Wilcoxon signed-rank test with a p -value less than 0.05.

5. Results and analysis

5.1 Overall multi-city performance

Table 2 shows the performance comparison of all four methods and three city configurations along the baseline capability gradient. From B1 to B2, threshold-triggered automation achieves 26-30% SIF reduction and 30-34% MTTR reduction. A reduction of 21-32% in SIF and 27-37% in MTTR is achieved by adding MAPE-K capabilities from B2 to B3. Finally, from B3 to our proposed method, the greatest improvements are observed in SIF reduction (58-64%) and MTTR reduction (69-70%). This holds across all maturity levels, validating the claim that domain-specific prediction with proactive migration and distributed consensus is the key differentiator. All improvements from B3 to the proposed method are statistically significant ($p < 0.05$).

Table 2. Performance comparison across methods and cities

Method	City	SIF (per 1k node-days)	MTTR (min)	SAR (%)	MPR
B1 Reactive-Manual	A	12.4	187.3	94.61	3.2
B2 Threshold-Auto	A	8.7	124.5	96.48	2.6
B3 MAPE-K	A	5.9	78.2	97.83	1.9
Proposed	A	2.1	23.6	99.47	0.8
B1 Reactive-Manual	B	18.6	246.8	91.37	4.1
B2 Threshold-Auto	B	13.2	168.4	93.85	3.3
B3 MAPE-K	B	9.4	109.7	95.92	2.4
Proposed	B	3.8	34.1	98.76	1.1
B1 Reactive-Manual	C	27.3	412.5	86.24	5.7
B2 Threshold-Auto	C	20.1	287.6	89.63	4.5
B3 MAPE-K	C	15.8	196.3	92.14	3.6
Proposed	C	6.5	58.7	97.31	1.6

5.2 Ablation analysis of self-healing mechanisms

An ablation study of City-B is performed to examine the contributions of individual components of the system, as depicted in Figure 3. There are three settings that are being compared. Prediction only enables health monitoring with warnings, but without automated migration. Prediction with Migration enables automated workload balancing, but without the consensus protocol. Finally, Full Pipeline enables all components of the system sequentially. Prediction achieves only a 41% reduction in MTTR with advanced warning, while SIF remains high because there is no automated redistribution to mitigate the interruption itself. The addition of the migration step yields the greatest reduction in SIF, reducing it by an additional 53% because workloads are migrated before any service degradation is realized. The Full Pipeline achieves an additional 38% reduction in the variance of SIF and MTTR compared to the pipeline without consensus, with this effect most pronounced during simulated correlated failures. Each module addresses a different facet of fault resilience, and the complete pipeline produces results that no single module can achieve on its own [15].




Figure 3. Ablation study results on City-B configuration

5.3 Impact of infrastructure maturity

To assess the effectiveness of self-healing across different deployment environments, the improvement in the architecture’s performance is normalized relative to the B1 baseline for each city, as illustrated in Figure 4.




Figure 4. Normalized performance gain of the proposed architecture relative to B1 across three city configurations

The highest percentage improvements in SIF (83.1%), MTTR (87.4%), and MPR (75.0%) are achieved by City-A, where full-mesh redundancy and high node coverage make it the richest migration candidate for the self-healing pipeline. However, the high starting point for City-A means that the absolute benefit of the managed state relative to the unmanaged state is the smallest among the three cities. The highest improvement in SAR is achieved by City-C, at 11.07%. This is because the starting point for the unmanaged state is low, leaving ample room for improvement. The reduction in MPR from 5.7 to 1.6 is the largest, implying that the self-healing pipeline has the most significant staffing implications and is a crucial aspect in resource-constrained settings where such skilled staff are scarce. City-B is in the middle in terms of the four metrics, but its high level of redundancy in the network makes it the most cost-effective option, since it does not require the high infrastructure that is needed in City-A to achieve the benefits of the self-healing pipeline, while at the same time not performing so badly in the unmanaged state that the benefits would be minimal. The non-linear relationship between self-healing effectiveness and infrastructure maturity means that the deployment of the self-healing pipeline should not be standardized across the transport landscape but should vary by city.

5.4 Domain-specific prediction vs. generic model

The domain-specific prediction model is compared with a generic anomaly detector that uses an LSTM architecture but lacks transportation-specific environmental features. The domain-specific model outperforms the generic model across all three cities, with F1-score differences of 6.3 percentage points in City-A, 9.7 percentage points in City-B, and 14.2 percentage points in City-C. The difference is greater under more extreme conditions due to the dominance of environmental stress factors. Amongst individual failure types, vibration-related degradations show the greatest improvement, by 18.5 percentage points, for City-C. These findings verify the importance of using transportation-specific features and confirm that domain awareness benefits are more significant under more extreme conditions.

6. Discussion

6.1 Practical implications

The multi-city results provide differentiated recommendations for transportation agencies operating at different levels of infrastructure maturity. Agencies that operate at a medium level of maturity are expected to achieve the highest return on investment from deploying the entire self-healing pipeline. This is because they already have redundancy in their infrastructures that can be leveraged to migrate workloads effectively. At the same time, they have substantial room for improvement over their current practices. Agencies operating at a low level of infrastructure maturity and with a shortage of skilled maintenance personnel are better off investing in the predictive monitoring module and simplified migration. This will allow them to realize the maximum return in terms of personnel savings. Agencies that already have a high level of infrastructure maturity and are benefiting from acceptable fault tolerance should invest in the consensus module. This will enable them to better handle correlated multi-node failure events. The modularity of the proposed architecture enables each component to be deployed independently. The use of MQTT and REST interfaces in the proposed edge-layer architecture enables easy integration with existing traffic management systems without replacing existing infrastructure.

6.2 Limitations and future work

Despite these contributions, certain limitations should be noted. First, the evaluation method is entirely simulation-based, but the gap between simulated fault patterns and operational scenarios underscores the need to conduct pilot tests on real transportation systems. Second, the proposed architecture does not consider adversarial scenarios in which attackers target multiple nodes to exhaust capacity, a key concern as the number of IoT devices in urban environments continues to grow. Third, the overhead of consensus operations in extremely dense networks, such as those with more than 5,000 nodes, has yet to be characterized, potentially requiring optimizations to the protocol itself. Three avenues for further research are proposed. Future work may explore collaborative learning of failure models without compromising actual telemetry data. Another promising direction is that digital twin support would enable validation of self-healing strategies. Finally, extending the architecture to support mobile vehicle-to-everything node scenarios, enabling its application across the entire spectrum of cooperative intelligent transportation systems.

7. Conclusion

This study proposed a self-healing infrastructure architecture for transportation IoT networks that combines predictive health monitoring, preemptive workload migration, and distributed consensus in a sequential pipeline without centralized coordination. A domain-specific failure prediction model that considered vehicle vibration exposure, environmental stresses, and communication interference patterns was developed to capture failure mechanisms specific to roadside sensor networks. Comparative evaluations were conducted across three metropolitan infrastructure configurations with high, medium, and low infrastructure maturity levels to demonstrate that the proposed architecture reduces the frequency of infrastructure service interruptions by 58-64% and the mean time to recover by 69-70% compared to general-purpose autonomous infrastructure solutions. Ablation studies were also performed to demonstrate that each component of the proposed architecture contributes an irreplaceable dimension to fault resilience. Finally, the nonlinear relationship between self-healing effectiveness and infrastructure maturity suggests that medium-maturity configurations derive the greatest overall benefit, while low-maturity infrastructure environments achieve the most significant staffing reduction. These results provide a basis for transportation authorities to transition urban infrastructure sensor networks from operator-centric to autonomous, self-sustaining modes of operation.

Ethical issue

The authors are aware of and comply with best practices in publication ethics, specifically regarding authorship (avoidance of guest authorship), dual submission, figure manipulation, competing interests, and compliance with research ethics policies. The author adheres to publication requirements that the submitted work is original and has not been published elsewhere in any language.

Data availability statement

The manuscript contains all the data. However, more data will be available upon request from the corresponding author.

Conflict of interest

The author declares no potential conflict of interest.

References

- [1] S. Guerrero-Ibáñez, J. Zeadally, and J. Contreras-Castillo, "Sensor technologies for intelligent transportation systems," *Sensors*, vol. 18, no. 4, pp. 1–24, 2018.
- [2] A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, "Internet of Things for smart cities," *IEEE Internet of Things Journal*, vol. 1, no. 1, pp. 22–32, 2014.
- [3] X. Zhang, E. Miller-Hooks, and K. Denny, "Assessing the role of network topology in transportation network resilience," *Journal of Transport Geography*, vol. 46, pp. 35–45, 2015.
- [4] G. Psailer and S. Dustdar, "A survey on self-healing systems: approaches and systems," *Computing*, vol. 91, no. 1, pp. 43–73, 2011.
- [5] P. P. Ray, "A survey on Internet of Things architectures," *Journal of King Saud University – Computer and Information Sciences*, vol. 30, no. 3, pp. 291–319, 2018.
- [6] R. Ahmad and S. Kamaruddin, "An overview of time-based and condition-based maintenance in industrial application," *Computers & Industrial Engineering*, vol. 63, no. 1, pp. 135–149, 2012.
- [7] J. O. Kephart and D. M. Chess, "The vision of autonomic computing," *Computer*, vol. 36, no. 1, pp. 41–50, 2003.
- [8] V. Medel, O. Rana, J. A. Bañares, and U. Arronategui, "Modelling performance and resource management in Kubernetes," in *Proc. 9th International Conference on Utility and Cloud Computing*, 2016, pp. 257–262.
- [9] A. Burns, B. R. Wills, and A. Barker, "Self-healing trans-cloud applications," *Computing*, vol. 104, no. 1, pp. 81–114, 2022.
- [10] T. Chen, X. Liu, B. Xia, W. Wang, and Y. Lai, "Unsupervised anomaly detection of industrial robots using sliding-window convolutional variational autoencoder," *IEEE Access*, vol. 8, pp. 47072–47081, 2020.
- [11] M. S. Mahdavinejad, M. Rezvan, M. Barekatin, P. Adibi, P. Barnaghi, and A. P. Sheth, "Machine learning for Internet of Things data analysis: a survey," *Digital Communications and Networks*, vol. 4, no. 3, pp. 161–175, 2018.
- [12] S. Kiranyaz, O. Avci, O. Abdeljaber, T. Ince, M. Gabbouj, and D. J. Inman, "1D convolutional neural networks and applications: a survey," *Mechanical Systems and Signal Processing*, vol. 151, pp. 1–21, 2021.
- [13] L. Liu, Z. Chang, X. Guo, S. Mao, and T. Ristaniemi, "Multi-objective optimization for computation offloading in fog computing," *IEEE Internet of Things Journal*, vol. 5, no. 1, pp. 283–294, 2018.
- [14] G. F. Riley and T. R. Henderson, "The ns-3 network simulator," in *Modeling and Tools for Network Simulation*, Berlin, Germany: Springer, 2010, pp. 15–34.
- [15] R. Meyes, M. Lu, C. W. de Puiseau, and T. Meisen, "Ablation studies in artificial neural networks," *arXiv preprint, arXiv:1901.08644*, 2019.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of Future Publishing LLC (Future) and/or the editor(s). Future and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.



This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).