Availability and Performance Evaluation of Vehicular Ad Hoc Networks

Luis Guilherme Silva[1], **Francisco Airton Silva**[1]

[1] Universidade Federal do Piauí – PI – Brasil

{luis.e,faps}@ufpi.edu.br

Abstract. *Urban mobility demands continuity of service, dynamic adaptation, and fault tolerance in distributed computing architectures. In this context, Vehicular Ad Hoc Networks (VANETs) integrated with edge computing support traffic monitoring and control applications under varying operational conditions, while resource constraints in Road Side Units (RSUs) make strategies based on physical replication unfeasible. Given this scenario, this dissertation develops analytical models based on Stochastic Petri Nets to evaluate the availability, performance, and sustainability of a VANET architecture composed of multiple RSUs connected to an edge server. The availability analysis considers physical, logical, and communication failures, evaluating operational continuity through virtual machine migration, with reduced recovery time even under simultaneous failures. The performance evaluation models horizontal autoscaling with reinstantiation, highlighting degradation of the average response time and an increase in the discard rate when scaling limits are inadequate, and better operational balance with dynamic policies. The sustainability analysis incorporates energy and environmental metrics, indicating a reduction in energy consumption and carbon emissions, while adjustments with experimental data from Pasid-Validator confirm the adherence between analytical and empirical results. In conclusion, the proposed models show that the combination of virtual machine migration and dynamic autoscaling allows maintaining service continuity, controlling performance under load variation, and reducing energy impact, supporting predictive planning of RSU-based VANET architectures.*

1. Introduction

The increase in vehicular traffic has intensified the dependence on distributed computing systems to support urban mobility. Vehicle Ad Hoc Networks (VANETs) emerge as an alternative to enable direct communication between vehicles and road infrastructure [Singh et al. 2022]. These networks allow for the continuous exchange of information related to traffic, safety, and urban management [Ni et al. 2021]. Edge computing shifts part of the processing to points close to the data origin. This proximity reduces latency and limits traffic towards central servers. Road Side Units (RSUs) concentrate communication and processing functions along the roads [Ferreira, W. 2023]. In these units, applications are executed in a distributed manner through virtual machines. The microservices organization favors functional isolation and independent management. The overall behavior of the system depends on the ability of the RSUs to respond to load variations and failure events [Verma 2023].

VANET architectures operate under recurring physical, logical, and communication failures [Shen et al. 2019]. The unavailability of a RSU compromises service execution and affects the operational continuity of the served segment. Isolated failures can cause load imbalance in adjacent units. Traditional fault tolerance strategies based on physical replication require permanent duplication of resources [Qiong et al. 2023]. This model increases operational costs and becomes unfeasible in edge infrastructures with limited capacity [Blessy A and S 2023]. The computational load on RSUs varies according to vehicular flow over time. The absence of automatic adaptation mechanisms leads to saturation or underutilization of available resources [Silva et al. 2024]. Much of the existing work analyzes availability and performance independently. The literature lacks models that represent failures, dynamic adaptation, and sustainability in an integrated way in vehicular environments [Ismail et al. 2022].

Given these challenges, this study develops analytical models based on Stochastic Petri Nets (SPNs) [Maciel 2023] to evaluate availability, performance, and sustainability in VANET architectures composed of multiple RSUs connected to an edge server. The proposal models the migration of virtual machines as a mechanism for operational continuity in the face of physical and logical failures. Horizontal autoscaling with reinstantiation represents the dynamic adaptation of processing capacity in response to load variation. The models allow for the analysis of the impact of these strategies on system response, disposal, and utilization metrics. The approach incorporates energy and environmental metrics associated with operational behavior. The proposed models provide an analytical basis for evaluating, in an integrated way, service continuity, operational performance, and energy impact in RSU-based VANET architectures, contributing to planning and decision-making in edge computing environments.

2. Research Questions

The objective of this research was to evaluate the availability, performance, and sustainability of VANET architectures composed of multiple RSUs interconnected to edge computing infrastructures. The study adopts analytical modeling based on SPN to represent failure events, workload variation, and dynamic resource management mechanisms, including virtual machine migration, horizontal autoscaling, and reinstantiation. By analyzing different operational scenarios under stochastic traffic demand, the research aims to quantify service continuity, response behavior, and energy impact, providing predictive support for the planning and configuration of RSU-based vehicular infrastructures.

- **RQ1:** How does the migration of virtual machines between RSUs contribute to service continuity and maintaining system availability in the face of physical, logical, and communication failures in VANET architectures?
- **RQ2:** How do horizontal autoscaling mechanisms with reinstantiation influence system performance behavior, considering variations in vehicle load, scheduling limits, and metrics such as average response time, discard probability, and resource utilization?
- **RQ3:** What is the impact of autoscaling and reinstantiation strategies on energy consumption, operational efficiency, and carbon emissions in edge computing-based VANET infrastructures, as investigated in the context of computational sustainability?

- **RQ4:** To what extent can the proposed SPN-based models consistently represent availability, performance, and sustainability metrics, enabling analytical validation and predictive assessment of different operational configurations in RSU-based VANETs?

Specifically, this study has employed Stochastic Petri Net modeling to evaluate availability, performance, and sustainability in RSU-based VANET architectures. Additionally, a Design of Experiments (DoE) approach is used to analyze the impact of factors related to virtual machine migration, autoscaling, and load variation on system behavior. By addressing these research questions, the study provides a structural understanding of the operational dynamics of vehicular infrastructures in edge computing and offers analytical support for resource planning and optimization in these distributed environments.

3. Main Publications and Contributions

The dissertation was structured based on a set of scientific publications developed throughout the research. In total, four articles make up the core of the results obtained, addressing complementary aspects of availability, performance, elasticity, and sustainability. Figure 1 shows the mapping between the research questions, the specific objectives, and the corresponding publications. The evaluation of the publications is also indicated, considering the classification of journals and events according to the Qualis Capes platform (https://qualis.capes.gov.br/).

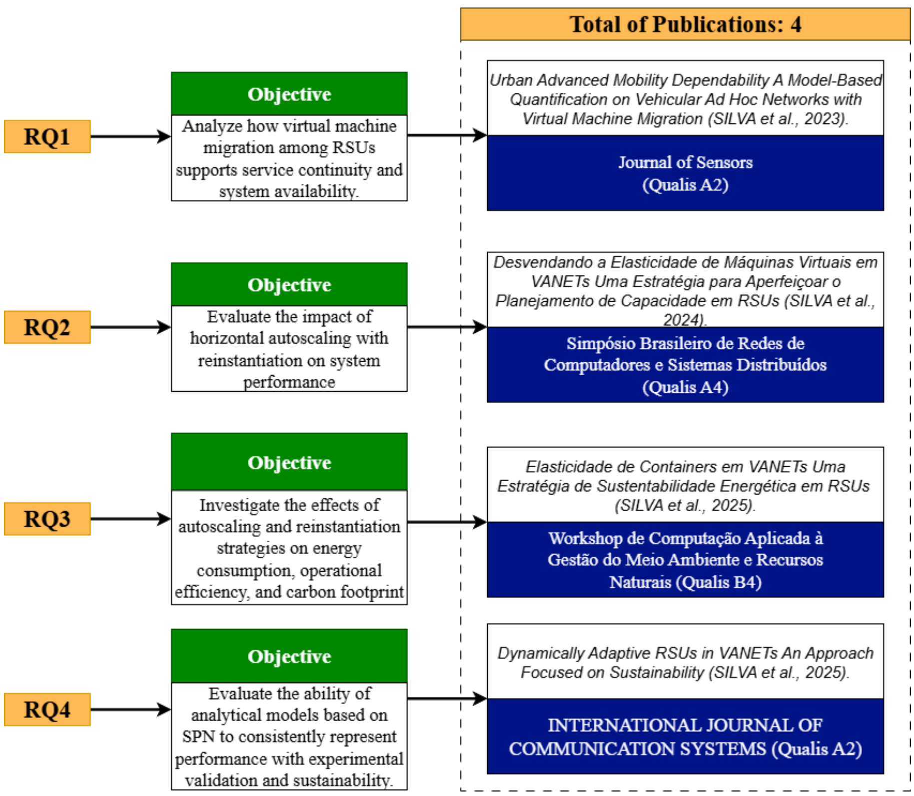

Figure 1. The research questions, highlighting the features employed and the published papers during the research.

3.1. Overview of Contributions

The first work establishes the basis for availability analysis by modeling VANET architectures with multiple RSUs subject to physical, logical, and communication failures. The analytical model, built with SPNs, represents the migration of virtual machines as a service continuity mechanism, allowing the quantification of the impact of detection time and recovery time on overall availability. The evaluation of scenarios with distributed failures shows that coordinated migration reduces downtime and limits the propagation of interruptions between adjacent units. These results underpin the operational continuity axis adopted throughout the dissertation.

Focusing on behavior under variable load, the second work analyzes the dynamic adaptation of computational capacity in RSUs. The model represents the processing of vehicle requests with horizontal autoscaling and automatic re-instantiation, incorporating failures in the instance creation process. Metrics of average response time, probability of discarding, utilization, and number of active instances are evaluated under different scaling limits. The analysis shows that restrictive configurations lead to saturation and performance degradation, while dynamic policies maintain operational balance in the face of temporal demand variation. This study consolidates the performance axis of the dissertation and guides capacity planning decisions.

The third publication expands the elasticity analysis by incorporating energy metrics into the analytical model. The system's behavior is now evaluated in terms of energy consumption and operational efficiency, considering the impact of autoscaling and container re-instantiation. The results show that excessive instance activation increases computational waste, while adaptive strategies reduce consumption without compromising response metrics. This analysis establishes the relationship between performance and energy impact, formally introducing the sustainability dimension into the scope of the dissertation.

The fourth paper integrates the axes of availability, performance, and sustainability into a unified analytical model, complemented by experimental validation. The model represents migration, autoscaling, re-instantiation, and energy metrics jointly, allowing for the evaluation of different operational configurations. The application of DoE identifies the parameters with the greatest influence on performance and energy efficiency. The comparison between analytical results and real data demonstrates statistical adherence, confirming the predictive capacity of the models. This stage consolidates the dissertation as an integrated approach for the evaluation and planning of VANET infrastructures based on RSUs.

3.2. Methodological Contributions

The main methodological contribution of this dissertation lies in the use of Stochastic Petri Nets as a unifying formalism for the integrated evaluation of availability, performance, and sustainability in VANET architectures based on RSUs. The modeling was conducted in a modular way, allowing the representation of concurrent failure, recovery, capacity adaptation, and stochastic load variation events in a single analytical structure. This choice makes it possible to capture temporal and structural dependencies that are not adequately represented by isolated approaches based on queues, Markov chains, or point simulations.

The methodology explicitly defines operational metrics, failure parameters, and adaptation policies, allowing the analysis of the average steady-state system behavior under different operating conditions. The incorporation of DoE-based sensitivity analysis complements the model by identifying parameters with the greatest impact on critical metrics, providing quantitative support for the planning and configuration of vehicular infrastructures. Finally, experimental validation with real data reinforces the consistency of the method by comparing analytical and empirical results, ensuring that the proposed models not only describe the system formally, but also reflect its observed behavior.

3.3. Technical Contributions

The technical contributions of this dissertation stem from the progressive construction of analytical models capable of representing, in an integrated way, the operational dynamics of VANET architectures based on RSUs. The formal modeling adopts SPNs to describe the system's behavior in the face of failures, load fluctuations, and limitations inherent to edge computing. From this basis, the work goes on to explore how adaptation mechanisms influence service continuity and operational behavior under variable conditions.

The representation of virtual machine migration introduces the ability to analyze service continuity dynamically, considering the time required to detect failures and restore operation in adjacent units. By explicitly modeling failure, recovery, and reallocation states, the system ceases to be treated as a set of isolated components and begins to be observed as a cooperative structure. This perspective allows us to understand how recovery decisions affect overall availability, especially in scenarios with distributed failures. The focus shifts from the simple occurrence of the failure to the temporal impact of its mitigation.

With the introduction of horizontal autoscaling, the analysis advances to the behavior of the system under stochastic variation of vehicle load. The model begins to capture queue formation, capacity limits, and the effects of automatic re-instantiation, revealing how adaptation decisions directly influence response, disposal, and utilization metrics. The interaction between load, capacity, and failures during the adaptation process shows that performance depends not only on the availability of resources, but also on how these resources are adjusted over time. Adaptation ceases to be an auxiliary mechanism and assumes a central role in controlling operational behavior.

The incorporation of energy metrics broadens the scope of the analysis by relating the adaptive behavior of the system to the consumption of computational resources. The number of active instances, defined by scheduling policies, is now also observed from the perspective of operational efficiency. This extension allows us to understand that decisions focused exclusively on latency reduction can generate energy waste, while more balanced policies tend to reduce consumption without compromising service. The model now reflects not only the system's operation but also the consequences of its operational choices.

The contributions of this dissertation result from the consistent articulation between formal modeling, sensitivity analysis, and experimental validation, applied to the study of VANET architectures based on solid waste containers. Throughout the work, recovery mechanisms, dynamic capacity adaptation, and energy efficiency are progressively incorporated into the same analytical model, allowing for the evaluation of the system's behavior under different operational conditions. This integrated approach enables the joint analysis of availability, performance, and sustainability, highlighting relationships that are not observed when these aspects are treated in isolation. As a result, the dissertation offers a set of models capable of supporting planning and configuration decisions for vehicular infrastructures in edge computing environments, providing analytical references for defining operational policies compatible with resource constraints, demand variation, and service continuity requirements.

4. Related Work

This section organizes related studies according to their evaluation focus, since this choice directly influences how VANET infrastructures are planned, dimensioned, and managed under stochastic demand. Table 1 consolidates the literature by summarizing the main performance and sustainability metrics, the modeling or evaluation methods employed, and whether traffic variation, multiple RSUs, and energy-related aspects are explicitly considered. This classification enables a structured comparison of existing approaches and clarifies the methodological gaps that motivate the analytical framework proposed in this dissertation.

A significant portion of the literature concentrates on *communication performance*, adopting analytical formulations or optimization models to characterize delay, transmission probability, and coverage under varying traffic conditions. [Tang et al. 2019] explored centralized routing with mobility prediction to improve transmission probability and reduce average delay. [Cumbal et al. 2019] addressed RSU deployment through an ILP formulation focused on coverage and communication rates. [Wu et al. 2020] investigated task assignment in vehicular fog computing environments with emphasis on total time delay, while [Siddiqi et al. 2020] applied queuing theory to evaluate response time, resource cost, and service quality in cloud-assisted vehicular processing. Although these works provide relevant insights into latency and efficiency, they typically abstract elasticity mechanisms and do not incorporate fault handling or coordinated resource management across multiple RSUs.

Other studies shift the analysis toward *availability* or *energy-aware operation*. [Silva et al. 2023] employed Stochastic Petri Net models to quantify reliability and availability in urban mobility scenarios, explicitly considering the interaction among multiple RSUs. In parallel, energy-oriented works such as [Mehta and Mahato 2023], [Wu et al. 2022], and [Agarwal et al. 2016] evaluated energy consumption, load balancing efficiency, and network load using simulation or hybrid modeling approaches. These contributions advance the understanding of resilience and sustainability in vehicular infrastructures; however, they generally treat performance, availability, and energy efficiency as separate concerns and do not integrate adaptive mechanisms such as autoscaling or service reinstantiation into a unified analytical model.

Overall, the literature reveals a fragmentation in how VANET infrastructures are evaluated. Traffic variation, multi-RSU operation, elasticity mechanisms, and sustainability metrics rarely coexist within the same formal framework. These studies demonstrate that existing approaches analyze communication performance, availability, and energy efficiency in isolation, without a unified formal model capable of capturing elasticity, fault recovery, and sustainability under stochastic traffic in environments with multiple RSUs.

This study addresses these limitations by proposing an integrated SPN-based model that jointly evaluates performance, availability, and sustainability in RSU-based VANETs. The model explicitly represents stochastic traffic profiles, distributed RSU operation, virtual machine migration, dynamic autoscaling, and energy-related metrics, with analytical results validated through experimental data. This integration enables a more realistic and consistent assessment of adaptive vehicular edge infrastructures than those provided by existing approaches.

Table 1. Related work.

Work	Metrics	Method	Traffic variation	Multiple RSUs	Sustainability
[Tang et al. 2019]	Transmission probability, Average delay	Mathematical Model	✓	✗	✗
[Silva et al. 2023]	Reliability, Availability	SPN Model	✗	✓	✗
[Cumbal et al. 2019]	Coverage rate, Communication rate	ILP Model	✓	✗	✗
[Wu et al. 2020]	Total Time Delay	VFC Model	✓	✗	✗
[Siddiqi et al. 2020]	Response Time, Resource Cost, Service Quality	Queuing Model	✓	✗	✗
[Mehta and Mahato 2023]	Energy consumption, queue wait time, load balancing efficiency	Simulation	✗	✗	✓
[Wu et al. 2022]	Task response time, vehicle energy, load balancing	Simulation, modeling	✗	✗	✓
[Agarwal et al. 2016]	Energy consumption, network load, avg. packet delay	Simulation, modeling	✗	✗	✓
This work	NVMU, DP, TP, MRT, U, Availability, Energy consumption, Carbon footprint, Energy efficiency	SPN Model + Experimental validation	✓	✓	✓

5. Discussion

This section discusses the main challenges and the relevance of the results obtained from the integrated evaluation of availability, performance, and sustainability in VANET architectures based on RSUs and edge computing. The analysis considers virtual machine migration, horizontal autoscaling, and reinstantiation mechanisms modeled through Stochastic Petri Nets.

5.1. Challenges

The evaluation of distributed vehicular infrastructures introduces challenges that arise from the simultaneous variability of workload, failures, and resource constraints at the edge. A primary challenge concerns the **trade-off between service continuity and efficient resource usage**. Virtual machine migration enables service preservation after physical or logical failures in RSUs, but it introduces operational overhead and depends on detection latency and recovery times. The absence of migration mechanisms reduces architectural complexity while exposing the system to extended service interruptions under failure conditions.

Another challenge is related to elasticity under stochastic traffic variation. Vehicular demand fluctuates over time and presents pronounced peaks during specific intervals. Without autoscaling, RSUs tend to saturate, increasing smean response time and discard probability. Conversely, excessive activation of instances leads to underutilization and increased energy consumption, which is critical in resource-constrained edge environments. The interdependence between performance and sustainability represents an additional challenge. Strategies optimized exclusively for latency often maintain a high number of active instances, negatively affecting energy efficiency. More aggressive energy-saving configurations, in turn, may degrade operational metrics. This interaction requires joint evaluation across multiple scenarios, considering distributed operation with multiple RSUs and realistic scaling limits.

5.2. Application of Stochastic Petri Net Models

To address these challenges, this dissertation employed SPN-based models as a formal tool to represent the dynamic behavior of VANET infrastructures. **The models capture concurrency, timed events, failures, recovery actions, and capacity adaptation** within a unified analytical framework. The **availability models** incorporate physical, logical, and communication failures, as well as virtual machine migration across RSUs, enabling the quantification of availability and reliability under different configurations. This approach supports the evaluation of recovery behavior without relying on static redundancy strategies.

The performance models represent vehicular request processing with horizontal autoscaling and reinstantiation, allowing the evaluation of metrics such as mean response time, discard probability, utilization, and number of virtual machines in use. The application of DoE enabled the identification of parameters with the greatest impact on performance and energy consumption, including the initial number of instances, reserved capacity, and failure-related weights. This methodology supports structured exploration of the parameter space while preserving statistical consistency.

5.3. Importance of the Results

The **results provide direct contributions to the planning of vehicular infrastructures based on edge computing**. The availability analysis shows that virtual machine migration reduces recovery time and preserves service continuity even under simultaneous RSU failures. Performance evaluation demonstrates that horizontal autoscaling improves operational balance under variable workloads when scaling limits are adequately configured. From a sustainability perspective, the results indicate that dynamic scaling and reinstantiation policies reduce energy consumption and carbon footprint without degrading response-related metrics. The integrated evaluation reveals that isolated optimization strategies focused solely on performance or energy savings tend to produce adverse effects when analyzed systemically. By modeling migration, elasticity, and sustainability together, the proposed approach captures interactions that are not visible in isolated analyses.

The results support architectural decision-making for designers of vehicular and edge-based systems to software engineers. The models allow software architects to anticipate saturation points, define scaling thresholds, and evaluate recovery strategies before deployment. This capability reduces dependence on costly field experimentation and supports the specification of adaptive operational policies aligned with performance and sustainability requirements in distributed RSU-based environments. Overall, this dissertation demonstrates that the combined modeling of migration, elasticity, and energy-aware operation, evaluated under stochastic traffic variation, provides a consistent analytical basis for the design and operation of VANET infrastructures with multiple RSUs.

6. Conclusion

This dissertation presented an integrated analysis of VANET architectures based on RSUs, focusing on the evaluation of availability, performance, and sustainability through analytical modeling with SPNs and experimental validation. By addressing resource limitations at the edge, vehicle load variation, and the occurrence of failures, the study contributes to understanding the impact of virtual machine migration mechanisms, autoscaling, and dynamic adaptation on service continuity and operational efficiency in vehicular environments. The research investigated four research questions (QP1–QP4), each associated with a publication and a specific analytical axis. The first analyzed the migration of virtual machines between RSUs as a recovery strategy, highlighting a reduction in recovery time and maintenance of availability under physical, logical, and communication failures. The second evaluated horizontal autoscaling mechanisms with reinstantiation, showing that inadequate limits result in degradation of the average response time and an increase in the discard rate, while dynamic policies promote better operational balance. The third incorporated the sustainability dimension, indicating that adaptive strategies reduce energy consumption and carbon footprint without compromising performance. The fourth confirmed that the proposed models consistently represent performance, availability, and energy metrics in different operational scenarios. The results of this dissertation contribute to the field of VANETs by providing analytical models capable of supporting planning and configuration decisions for vehicular infrastructures based on edge computing. The articulation between formal modeling, sensitivity analysis, and experimental validation ensures that the contributions presented in the publications derived from this work are applicable to real-world scenarios, allowing for the early assessment of operational limitations and guiding the efficient use of computational and energy resources. By providing an analytically grounded and experimentally validated framework, this work supports predictive capacity planning and energy-aware policy definition for future intelligent transportation systems based on distributed RSU infrastructures.

References

Agarwal, S., Das, A., and Das, N. (2016). An efficient approach for load balancing in vehicular ad-hoc networks. In *2016 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*, pages 1–6.

Blessy A, M. C. and S, B. (2023). Maximizing vanet performance in cluster head selection using intelligent fuzzy bald eagle optimization. *Vehicular Communications*, page 100660.

Cumbal, R., Gutiérrez, S., Guerrero, C., Hincapié, R., and Arévalo, G. (2019). Optimal resources allocation from vanet infrastructures in dynamic mobile environments. In *2019 IEEE Latin-American Conference on Communications (LATINCOM)*, pages 1–5. IEEE.

Ferreira, W. (2023). Ibm alerta sobre pressão de dados dos carros conectados. Disponível em: https://www.telesintese.com.br/ibm-alerta-sobre-pressao-dos-carros-conectados-nas-redes-de-telecom/. Acesso em: 09 de agosto 2023.

Ismail, N. et al. (2022). Enhanced congestion control model based on message prioritization and scheduling mechanism in vehicle-to-infrastructure (v2i). In *Journal of Physics: Conference Series*, volume 2312, page 012087. IOP Publishing.

Maciel, P. R. M. (2023). *Performance, reliability, and availability evaluation of computational systems, volume I: performance and background*. CRC Press.

Mehta, T. and Mahato, D. P. (2023). Effective scheduling and nature inspired hybrid load balancing in vanets. In *8th International Conference on Computing in Engineering and Technology (ICCET 2023)*, volume 2023, pages 266–273.

Ni, Y., Zhao, C., and Cai, L. (2021). Hybrid rsu management in cybertwin-iov for temporal and spatial service coverage. *IEEE Transactions on Vehicular Technology*, 71(5):4596–4606.

Qiong, W., Shuai, S., Ziyang, W., Qiang, F., Pingyi, F., and Cui, Z. (2023). Towards v2i age-aware fairness access: a dqn based intelligent vehicular node training and test method. *Chinese Journal of Electronics*, 32(6):1230–1244.

Shen, J., Liu, D., Chen, X., Li, J., Kumar, N., and Vijayakumar, P. (2019). Secure real-time traffic data aggregation with batch verification for vehicular cloud in vanets. *IEEE Transactions on Vehicular Technology*, 69(1):807–817.

Siddiqi, M. H., Alruwaili, M., Ali, A., Haider, S. F., Ali, F., and Iqbal, M. (2020). Dynamic priority-based efficient resource allocation and computing framework for vehicular multimedia cloud computing. *IEEE access*, 8:81080–81089.

Silva, L. G., Brito, C., Cardoso, I., Sabino, A., Lima, L. N., Gonçalves, G., Rocha Filho, G. P., Fé, I., and Silva, F. A. (2024). Desvendando a elasticidade de máquinas virtuais em vanets: Uma estratégia para aperfeiçoar o planejamento de capacidade em rsus. In *Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos (SBRC)*, pages 169–182. SBC.

Silva, L. G., Cardoso, I., Brito, C., Barbosa, V., Nogueira, B., Choi, E., Nguyen, T. A., Min, D., Lee, J. W., and Silva, F. A. (2023). Urban advanced mobility dependability: A model-based quantification on vehicular ad hoc networks with virtual machine migration. *Sensors*, 23(23):9485.

Singh, G. D., Prateek, M., Kumar, S., Verma, M., Singh, D., and Lee, H.-N. (2022). Hybrid genetic firefly algorithm-based routing protocol for vanets. *IEEE Access*, 10:9142–9151.

Tang, Y., Cheng, N., Wu, W., Wang, M., Dai, Y., and Shen, X. (2019). Delay-minimization routing for heterogeneous vanets with machine learning based mobility prediction. *IEEE Transactions on Vehicular Technology*, 68(4):3967–3979.

Verma, R. (2023). An efficient secure vanet communication using multi authenticate homomorphic signature algorithm. In *2023 International Conference on Distributed Computing and Electrical Circuits and Electronics (ICDCECE)*, pages 1–5.

Wu, X., Zhao, S., Zhang, R., and Yang, L. (2020). Mobility prediction-based joint task assignment and resource allocation in vehicular fog computing. In *2020 IEEE Wireless Communications and Networking Conference (WCNC)*, pages 1–6. IEEE.

Wu, Y., Wu, J., Chen, L., Yan, J., and Han, Y. (2022). Load balance guaranteed vehicle-to-vehicle computation offloading for min-max fairness in vanets. *IEEE Transactions on Intelligent Transportation Systems*, 23(8):11994–12013.