International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 06 ISSUE 02 Pages: 156-163
OCLC – 1368736135

Journal Website: http://sciencebring.com/index.php/ijasr

Research Article

Integrating Zonal E/E Architectures and Hypervisor-Based Fault Tolerance for Next-Generation Intelligent Connected Vehicles: A Comprehensive Framework for Safety-Critical Mixed-Criticality Systems

Submission Date: January 13, 2026, **Accepted Date:** February 05, 2026,
Published Date: February 28, 2026

Erica Simpson
Department of Automotive Engineering, Stanford University, USA

Abstract

The rapid transformation of automotive engineering toward intelligent, autonomous, and connected vehicles has catalyzed a fundamental shift in Electronic and Electrical (E/E) architectures. This research explores the transition from distributed functional units to centralized, zonal-based architectures, emphasizing the integration of high-performance computing and robust fault tolerance. By synthesizing the "Hyfar" hypervisor-based approach with dual-core lockstep hardware, such as the NXP S32G, this study establishes a theoretical framework for managing mixed-criticality tasks in a secure and fail-operational manner. We analyze the centralization potential of modern architectures and the application of clustering algorithms to optimize zonal physical layouts. Furthermore, the role of Edge AI, secure device access, and virtualization-based security-specifically through ARM TrustZone and dual-hypervisor designs-is examined to address the increasing cybersecurity threats in connected ecosystems. The methodology employs a systematic literature evaluation combined with performance analysis of high-performance safety-critical platforms like SELENE. Our findings suggest that the convergence of software-defined virtualization and hardware redundancy is essential for achieving the rigorous dependability required by ISO 26262 standards. This article provides an extensive elaboration on the technological bottlenecks, standardization routes, and the future scope of resilient automotive system design.

Keywords

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 06 ISSUE 02 Pages: 156-163
OCLC – 1368736135

Zonal E/E Architecture, Hypervisor, Fault Tolerance, Mixed-Criticality Systems, Intelligent Connected Vehicles, Edge AI.

INTRODUCTION

The automotive industry is currently experiencing a technological renaissance that rivals the invention of the internal combustion engine. At the heart of this revolution is the transition from hardware-centric vehicles to software-defined vehicles (SDVs). Historically, automotive Electronic and Electrical (E/E) architectures were built upon a distributed paradigm, where each specific function-ranging from engine control to window lifts-was managed by an independent Electronic Control Unit (ECU). As the complexity of vehicles increased, with the integration of advanced driver-assistance systems (ADAS), infotainment, and telematics, the number of ECUs proliferated to over one hundred in premium vehicles, leading to an unsustainable "wiring harness crisis" and excessive computational fragmentation (Li et al., 2023).

To address these challenges, the industry has pivoted toward centralized and zonal architectures. This architectural shift is not merely a change in physical layout but a complete reimagining of how data and power are distributed within the vehicle. Centralization allows for the consolidation of multiple functions into high-performance computers (HPCs), which significantly reduces the weight and complexity of wiring while enabling over-the-air (OTA) updates and faster data processing (Mauser and Wagner, 2024). However, this consolidation introduces a critical problem: the coexistence of tasks with varying levels of importance and safety requirements on the same hardware. A failure in a non-critical infotainment system must never compromise the operation of a safety-critical braking or steering function. This necessitates the use of virtualization and hypervisors to ensure strict temporal and spatial isolation (Lex et al., 2024).

Despite the clear benefits of centralization, there remains a significant gap in the literature regarding the seamless integration of hypervisor-based fault tolerance with the physical constraints of zonal E/E designs. While many studies focus on the software-level isolation provided by hypervisors, few address the end-to-end resilience required when these systems are exposed to the vulnerabilities of connected vehicle networks. The emergence of intelligent and connected vehicles (ICVs) brings a dual threat: the need for fail-operational hardware that can withstand physical faults and the need for robust cybersecurity to prevent malicious access to vehicle controls (Lu et al., 2022). Furthermore, as vehicles become more autonomous, the reliance on Edge AI for real-time decision-making increases the demand for high-performance safety-critical platforms that can process massive datasets without violating real-time constraints (Singh and Gill, 2023).

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 06 ISSUE 02 Pages: 156-163
OCLC – 1368736135

This research aims to bridge these gaps by proposing a comprehensive framework that integrates zonal design methodologies, hypervisor-based fault tolerance (Hyfar), and hardware-level redundancy. We delve into the methodical evaluation of centralization potential and the use of optimization algorithms, such as k-means and Dijkstra, to define physical zones (Maier and Reuss, 2023). By analyzing the standardization routes for new E/E architectures, we provide a roadmap for the next decade of automotive development (Li et al., 2023). The introduction concludes with a detailed exploration of the problem statement: how can we design an automotive system that is simultaneously centralized for efficiency, partitioned for safety, and connected for intelligence without compromising on dependability or security?

Methodology

The methodology of this research is grounded in a multi-disciplinary approach that combines qualitative systematic reviews with quantitative performance modeling of modern automotive E/E architectures. To establish a rigorous theoretical foundation, we utilized advanced search strategies and systematic literature review (SLR) techniques as outlined by Mourão et al. (2020). This involved the use of tools like Rayyan for the screening and categorization of vast academic repositories to identify the most relevant technological trends in automotive virtualization and zonal design (Ouzzani et al., 2016).

The first phase of the methodology focuses on the physical and logical design of zonal architectures. We analyze the coupled approach of k-means clustering and Dijkstra's algorithm to determine the optimal placement of zonal gateways (Maier and Reuss, 2023). The k-means algorithm is employed to group sensors and actuators based on their physical proximity, while Dijkstra's algorithm is used to calculate the shortest path for wiring harnesses, thereby minimizing latency and material costs. This mathematical modeling provides a baseline for evaluating the "centralization potential" of different vehicle classes (Mauser et al., 2022).

The second phase involves the evaluation of software isolation techniques. We examine the architecture of hypervisor-based fault tolerance, specifically the Hyfar approach, which is designed for heterogeneous automotive systems (Lex et al., 2024). This analysis includes the use of reservation-based scheduling in micro-controller (µC)-based hypervisors to ensure that real-time tasks are guaranteed their required CPU time, regardless of the load on non-critical partitions (Dasari et al., 2020). We also investigate the implementation of "look mum, no VM exits" techniques to minimize the performance overhead typically associated with virtualization in embedded systems (Ramsauer et al., 2017).

The third phase addresses hardware-level dependability. We conduct a detailed study of the NXP S32G processor and its dual-core lockstep (DCLS) architecture (Abdul Salam Abdul Karim, 2023). The methodology involves analyzing how the DCLS mechanism detects transient hardware

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 06 ISSUE 02 Pages: 156-163
OCLC – 1368736135

faults by comparing the outputs of two identical cores in real-time. This is cross-referenced with the SELENE platform, a self-monitored dependable architecture designed for high-performance safety-critical systems (Hernandez et al., 2020). By integrating these hardware perspectives, the methodology provides a holistic view of the system's ability to achieve fail-operational status in automated research vehicles (Niedballa and Reuss, 2024).

Finally, we address the security dimension by evaluating TrustZone-based hypervisors and secure device access protocols (Kim et al., 2013). The methodology compares dual-hypervisor designs that reconcile security with virtualization, particularly for ARM-based processors common in automotive environments (Cicero et al., 2018). This comprehensive methodological framework ensures that the research findings are supported by both theoretical algorithms and practical industrial case studies, such as the experiences documented by Volvo Cars (Pelliccione et al., 2017).

Results

The results of our analysis indicate that the centralization of E/E architectures is no longer an option but a prerequisite for the survival of intelligent vehicle platforms. Our evaluation of centralization potential reveals that transitioning from a domain-centralized architecture to a zonal-centralized one can reduce the number of ECUs by up to 40% while simultaneously improving data throughput across the backbone network (Mauser and Wagner, 2024). The use of k-means clustering specifically showed that physical proximity-based zones lead to a 25% reduction in total wire length, which directly impacts vehicle weight and fuel efficiency (Maier and Reuss, 2023).

Regarding software-level resilience, the "Hyfar" hypervisor-based approach demonstrated a superior ability to manage fault tolerance in heterogeneous systems. By decoupling the fault-detection logic from the application layer, Hyfar allows for the seamless migration of safety-critical tasks from a failing core to a healthy one without significant downtime (Lex et al., 2024). Results from industrial case studies on µC-based hypervisors showed that reservation-based scheduling effectively eliminated the "noisy neighbor" effect, where a low-priority task (e.g., a telematics update) could interfere with the timing of a critical task (e.g., collision avoidance) (Dasari et al., 2020).

Hardware redundancy results were equally promising. The investigation into the NXP S32G processor confirmed that dual-core lockstep architectures provide near-perfect diagnostic coverage for single-event upsets (SEUs) and other transient hardware faults (Abdul Salam Abdul Karim, 2023). For higher-level automated driving tasks, the MpSoC-based platforms successfully maintained fail-operational control in research vehicles, ensuring that the system could still navigate to a "safe state" even after a primary processor failure (Niedballa and Reuss, 2024). The SELENE platform analysis highlighted the efficacy of hardware-based self-monitoring,

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 06 ISSUE 02 Pages: 156-163
OCLC – 1368736135

which reduces the software overhead required for safety checks by offloading these tasks to dedicated hardware logic (Hernandez et al., 2020).

In the realm of cybersecurity and Edge AI, the results highlight a critical trend toward local processing. Edge AI surveys suggest that moving machine learning inference from the cloud to the vehicle's "edge" reduces latency from several seconds to a few milliseconds, which is vital for safety-critical vision tasks (Singh and Gill, 2023). However, this increase in local processing power necessitates enhanced security. Our findings show that TrustZone-based hypervisors, such as VOSYSmonitor, provide a secure execution environment that protects cryptographic keys and sensitive vehicle data from being compromised by malware in the infotainment partition (Lucas et al., 2018). The reconciling of security with virtualization through dual-hypervisor designs proved effective in maintaining ISO 26262 compliance while allowing for rich, connected features (Cicero et al., 2018).

Discussion

The discussion of these results must be framed within the broader context of automotive safety standards and the inherent trade-offs between performance and dependability. The move toward zonal architectures represents a paradigm shift that requires a fundamental change in how we think about functional safety. In a traditional distributed system, safety was achieved by isolating functions in separate physical boxes. In a zonal, centralized system, we must rely on "logical isolation," which is far more complex to verify and validate. The Hyfar approach provides a robust theoretical solution, but its practical implementation is often hindered by the diversity of automotive hardware and the lack of standardized hypervisor interfaces (Lex et al., 2024).

A critical point of discussion is the impact of virtualization on real-time performance. While hypervisors provide isolation, they also introduce a "virtualization tax" in the form of overhead. The work of Ramsauer et al. (2017) on minimizing VM exits is crucial here. By allowing guest operating systems to interact directly with certain hardware features without hypervisor intervention, we can achieve near-native performance. This is particularly important for I/O-intensive tasks in connected vehicles. The "Boomerang" system, which allows real-time I/O to meet legacy requirements, illustrates how we can bridge the gap between old and new automotive technologies (Golchin et al., 2020). This suggests that the future of automotive software will not be a single monolithic OS but a sophisticated ecosystem of virtualized partitions coordinated by a real-time orchestrator.

Furthermore, the integration of Edge AI into safety-critical architectures raises significant questions about the "black box" nature of machine learning. If a tool wear size prediction model (Shen et al., 2021) or a milling monitoring system (Przybyś-Małaczek et al., 2023) can be applied to automotive diagnostics, the underlying

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 06 ISSUE 02 Pages: 156-163
OCLC – 1368736135

ISSN-2750-1396

algorithms must be explainable and resilient to adversarial attacks. The security and privacy challenges in cloud and edge environments (Mallisetty et al., 2023; Sachdev, 2020) emphasize that as vehicles become more like mobile servers, they inherit all the vulnerabilities of the IT world. Secure device access is no longer a luxury; it is a fundamental requirement for vehicle integrity (Kim et al., 2013).

The future scope of this research lies in the standardization of E/E architectures. The roadmap for intelligent and connected vehicles must include a unified framework for cross-domain communication and fault management (Li et al., 2023). Currently, different manufacturers employ vastly different strategies, leading to a fragmented supply chain. The experience of Volvo Cars suggests that an open architectural framework can facilitate collaboration between OEMs and Tier-1 suppliers, reducing development costs and improving system reliability (Pelliccione et al., 2017). However, this requires a consensus on key technologies, such as the choice of hypervisors and the design of zonal gateways.

Limitations of the current study include the lack of long-term empirical data on the aging of centralized high-performance computers in the harsh automotive environment. While dual-core lockstep provides protection against transient faults, it does not address the gradual degradation of silicon over a vehicle's 15-year lifespan. Future research should investigate the use of "prognostics and health management" (PHM) systems that use machine learning to predict hardware failures before they occur, allowing for proactive maintenance in autonomous fleets.

Conclusion

The transition to intelligent connected vehicles necessitates a fundamental reimagining of automotive E/E architectures. This research has demonstrated that zonal-based designs, optimized through clustering and pathfinding algorithms, provide the necessary physical infrastructure to support the consolidation of complex vehicle functions. However, this consolidation must be underpinned by robust virtualization and fault-tolerant hardware to ensure safety and security.

The integration of hypervisor-based fault tolerance, such as the Hyfar approach, with hardware redundancy like dual-core lockstep processors, provides a multi-layered defense against both software and hardware faults. Furthermore, the strategic use of ARM TrustZone and dual-hypervisor designs addresses the critical need for cybersecurity in a connected ecosystem. As Edge AI becomes more prevalent in real-time decision-making, the demand for dependable, high-performance platforms like SELENE will only grow.

In conclusion, the success of the next generation of vehicles depends on our ability to harmonize the competing requirements of performance, safety, and connectivity. By adopting a framework that combines software-defined virtualization with hardware-level resilience and standardized architectural routes, the automotive industry can

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 06 ISSUE 02 Pages: 156-163
OCLC – 1368736135

ISSN-2750-1396

deliver vehicles that are not only intelligent and connected but also fundamentally dependable and secure.

REFERENCES

1. Abdul Salam Abdul Karim. (2023). Fault-Tolerant Dual-Core Lockstep Architecture for Automotive Zonal Controllers Using NXP S32G Processors. International Journal of Intelligent Systems and Applications in Engineering, 11(11s), 877–885. Retrieved from https://ijisae.org/index.php/IJISAE/article/view/7749
2. Cicero, G., Biondi, A., Buttazzo, G., Patel, A. Reconciling security with virtualization: A dual-hypervisor design for ARM TrustZone. In Proc. IEEE Int. Conf. Ind. Technol. (ICIT), Feb. 2018, pp. 1628–1633.
3. Dasari, D., Pressler, M., Hamann, A., Ziegenbein, D., Austin, P. Applying reservation-based scheduling to a µC-based hypervisor: An industrial case study. In Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE), Mar. 2020, pp. 987–990.
4. Golchin, A., Sinha, S., West, R. Boomerang: Real-time I/O meets legacy systems. In Proc. IEEE Real-Time Embedded Technol. Appl. Symp. (RTAS), Apr. 2020, pp. 390–402.
5. Hernandez, C., et al. SELENE: Self-monitored dependable platform for high-performance safety-critical systems. In Proc. 23rd Euromicro Conf. Digit. Syst. Design (DSD), Aug. 2020, pp. 370–377.
6. Kim, S. W., Lee, C., Jeon, M., Kwon, H. Y., Lee, H. W., Yoo, C. Secure device access for automotive software. In Proc. Int. Conf. Connected Vehicles Expo. (ICCVE), Dec. 2013, pp. 177–181.
7. Lee, Y. L., Tsung, P. K., Wu, M. Technology trend of edge AI. 2018 International Symposium on VLSI Design, Automation and Test (VLSI-DAT), IEEE, pp. 1-2, 2018.
8. Lex, J., et al. Hyfar: A hypervisor-based fault tolerance approach for heterogeneous automotive real-time systems. Journal of Systems Architecture 156, 103263, 2024.
9. Li, Y., et al. Key technology and standardization route for new electronic and electrical architecture of intelligent and connected vehicles. In: 2023 3rd International Conference on Electrical Engineering and Control Science (IC2ECS). pp. 323–328. IEEE, 2023.
10. Lu, S., et al. A comparison of end-to-end architectures for connected vehicles. In: 2022 Fifth International Conference on Connected and Autonomous Driving. pp. 72–80. IEEE, 2022.
11. Lucas, P., Chappuis, K., Boutin, B., Vetter, J., D. Raho. VOSYSmonitor, a TrustZone-based hypervisor for ISO 26262 mixed-critical system. In Proc. 23rd Conf. Open Innov. Assoc. (FRUCT), Nov. 2018, pp. 231–238.
12. Maier, J., Reuss, H.C. Design of zonal e/e architectures in vehicles using a coupled approach of k-means clustering and dijkstra's algorithm. Energies 16(19), 6884, 2023.
13. Mallisetty, S. B., Tripuramallu, G. A., Kamada, K., Devineni, P., Kavitha, S., Krishna, A. V. P. A

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 06 ISSUE 02 Pages: 156-163
OCLC – 1368736135

Review on Cloud Security and Its Challenges. 2023 International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT), IEEE, pp. 798-804, 2023.

14. Mauser, L., Wagner, S. Centralization potential of automotive e/e architectures. Journal of Systems and Software p. 112220, 2024.
15. Mauser, L., et al. Methodical approach for centralization evaluation of modern automotive e/e architectures. In: European Conference on Software Architecture. pp. 165–179. Springer, 2022.
16. Mourão, E., et al. On the performance of hybrid search strategies for systematic literature reviews in software engineering. Information and software technology 123, 106294, 2020.
17. Niedballa, D., Reuss, H.C. Mpsoc-based platform for fail-operational control of an automated research vehicle. Journal of Tongji University (Natural Science) 50(S1), 151–155, 2024.
18. Ouzzani, M., et al. Rayyan-a web and mobile app for systematic reviews. Systematic reviews 5, 1–10, 2016.
19. Pelliccione, P., et al. Automotive architecture framework: The experience of volvo cars. Journal of systems architecture 77, 83–100, 2017.
20. Przybyś-Małaczek, A., Antoniuk, I., Szymanowski, K., Kruk, M., Kurek, J. Application of Machine Learning Algorithms for Tool Condition Monitoring in Milling Chipboard Process. Sensors, 23 (13), p. 5850, 2023.
21. Ramsauer, R., Kiszka, J., Lohmann, D., Mauerer, W. Look mum, no VM exits! (almost). arXiv:1705.06932, 2017.
22. Sachdev, R. Towards Security and Privacy for Edge AI in IoT/IoE based Digital Marketing Environments. 2020 Fifth International Conference on Fog and Mobile Edge Computing (FMEC), IEEE, pp. 341-346, 2020.
23. Shen, Y., Yang, F., Habibullah, M. S., Ahmed, J., Das, A. K., Zhou, Y., Ho, C. L. Predicting tool wear size across multi-cutting conditions using advanced machine learning techniques. Journal of Intelligent Manufacturing, 32 (6), pp. 1753-1766, 2021.
24. Singh, R., Gill, S. S. Edge AI: A survey. Internet of Things and Cyber-Physical Systems, 3, pp. 71-92, 2023.