

Resilient Embedded and Automotive Systems: Integrating Lockstep Architectures, Software-Based Fault Detection, And Cyber-Physical Safety Models for Next-Generation Reliability

Paul Kovalenko

Department of Computer Engineering, Technical University of Munich, Germany

Article received: 15/11/2025, Article Revised: 06/12/2025, Article Accepted: 31/12/2025

© 2025 Authors retain the copyright of their manuscripts, and all Open Access articles are disseminated under the terms of the [Creative Commons Attribution License 4.0 \(CC-BY\)](https://creativecommons.org/licenses/by/4.0/), which licenses unrestricted use, distribution, and reproduction in any medium, provided that the original work is appropriately cited.

ABSTRACT

The rapid evolution of embedded and automotive systems has introduced unprecedented complexity, driven by the integration of multi-core processors, real-time operating systems, and software-defined functionalities. This complexity has significantly increased the vulnerability of such systems to transient and permanent faults, particularly radiation-induced soft errors and memory safety violations. This research develops a comprehensive, theoretically grounded framework for fault tolerance that integrates hardware-based lockstep architectures, software-level fault detection and recovery mechanisms, and cyber-physical safety models. Drawing on foundational and contemporary literature, the study critically examines the limitations of software-only approaches in error detection coverage, the effectiveness of dual-core lockstep systems in mitigating soft errors, and the role of architectural diversity and safety frameworks such as the Simplex architecture and time-triggered systems. The methodology employs a conceptual modeling approach to analyze fault propagation, detection latency, and system recovery across heterogeneous computing environments, including automotive zonal controllers and high-performance embedded platforms. The findings demonstrate that hybrid architectures combining hardware redundancy with selective software-based mechanisms significantly enhance fault coverage and system resilience while maintaining manageable performance overhead. Furthermore, the incorporation of safety-oriented architectural paradigms effectively limits fault propagation and ensures predictable system behavior. The study highlights the importance of integrating memory safety mechanisms and control flow integrity techniques to address emerging software vulnerabilities. The discussion explores the implications of these findings for next-generation automotive and cyber-physical systems, emphasizing scalability, energy efficiency, and real-time constraints. Future research directions include adaptive fault-tolerance strategies and the integration of intelligent monitoring systems. This work contributes a unified perspective on resilient system design, bridging the gap between hardware reliability, software correctness, and system-level safety.

KEYWORDS

Fault tolerance, lockstep architecture, soft errors, embedded systems, automotive systems, cyber-physical safety, memory safety.

INTRODUCTION

The modern landscape of embedded and automotive systems is characterized by an extraordinary convergence of computational power, software complexity, and real-time operational demands. Systems that were once limited to basic control functions have evolved into highly sophisticated platforms capable of executing advanced algorithms for perception, decision-making, and actuation. This transformation is particularly evident in the automotive domain, where vehicles now

incorporate advanced driver assistance systems, autonomous driving capabilities, and interconnected communication frameworks. However, this rapid advancement has introduced significant challenges related to system reliability and safety, as the increasing complexity inherently amplifies the likelihood and impact of faults.

One of the most critical challenges in contemporary embedded systems is the prevalence of radiation-induced

soft errors. These transient faults, caused by external radiation sources such as cosmic rays, can alter the state of memory elements and logic circuits without causing permanent damage. As semiconductor technologies continue to scale down, the susceptibility of devices to such errors has increased substantially, making soft errors a major concern for system designers (Baumann, 2005). The impact of these errors is particularly severe in safety-critical applications, where even a single undetected fault can lead to catastrophic consequences.

Traditional fault-tolerance techniques have relied heavily on hardware redundancy, with lockstep architectures emerging as a widely adopted solution. In a dual-core lockstep configuration, two processor cores execute identical instruction streams simultaneously, and their outputs are continuously compared to detect discrepancies. This approach provides a robust mechanism for identifying transient faults and ensuring correct system operation. Empirical studies have demonstrated the effectiveness of lockstep architectures in mitigating soft errors in real-time operating systems such as FreeRTOS (de Oliveira et al., 2017). However, despite their advantages, lockstep systems are not without limitations. In particular, they are vulnerable to common-mode failures, where identical faults affect both cores simultaneously, thereby escaping detection.

The limitations of hardware-based approaches have prompted significant interest in software-level fault-tolerance techniques. These methods aim to enhance system reliability without requiring additional hardware resources, making them particularly attractive for cost-sensitive applications. Techniques such as software-based error detection, control flow integrity enforcement, and checkpoint and rollback recovery have been extensively studied. For instance, checkpointing mechanisms enable systems to periodically save their state and recover from faults by reverting to a previously known correct state (Bowen and Pradham, 1993). Similarly, control flow integrity techniques ensure that program execution follows a predefined path, thereby preventing malicious or erroneous deviations (Yao et al., 2020).

Despite their potential, software-only approaches face inherent limitations in their ability to detect all types of faults. Studies have shown that these techniques may fail to achieve complete error detection coverage, particularly in the presence of complex fault scenarios and low-level hardware errors (Azambuja et al., 2011). This limitation underscores the need for hybrid approaches that combine the strengths of both hardware and software mechanisms.

Another critical aspect of system reliability is memory safety. Memory-related errors, such as buffer overflows and pointer misuse, represent a significant source of vulnerabilities in embedded systems. Advanced techniques such as SoftBound and Bell have been

developed to address these issues by enforcing spatial memory safety and detecting memory leaks, respectively (Nagarakatte et al., 2009; Bond and McKinley, 2006). Furthermore, research into programming language semantics has highlighted the importance of understanding pointer provenance and memory behavior in ensuring software correctness (Memarian et al., 2019).

In addition to fault detection and recovery mechanisms, system-level architectural frameworks play a crucial role in ensuring safety and reliability. The Simplex architecture, for example, provides a structured approach to limiting fault propagation by incorporating a safety controller that can take over in the event of a failure (Crenshaw et al., 2007). Similarly, time-triggered architectures ensure predictable system behavior by scheduling tasks based on a global time base, thereby reducing the likelihood of timing-related faults (Kopetz and Bauer, 2003). These frameworks are particularly relevant in cyber-physical systems, where the interaction between computational and physical components introduces additional complexity.

The increasing adoption of multi-core and many-core systems further complicates the design of fault-tolerant architectures. Platforms such as NVIDIA DRIVE, Kalray processors, and Renesas automotive systems exemplify the trend toward high-performance computing in embedded environments. While these platforms offer significant computational capabilities, they also introduce new challenges related to fault detection, synchronization, and resource management. The complexity of these systems necessitates advanced fault-tolerance strategies that can operate effectively across multiple processing units.

Despite the extensive body of research in this field, several gaps remain. Most existing studies focus on individual aspects of fault tolerance, such as hardware redundancy or software-based detection, without considering their integration within a unified framework. Additionally, the implications of emerging technologies, such as multi-core architectures and advanced middleware systems, are not fully understood. This research aims to address these gaps by developing a comprehensive framework that integrates hardware, software, and system-level approaches to fault tolerance.

METHODOLOGY

The methodology adopted in this study is rooted in a comprehensive theoretical and conceptual analysis of fault-tolerant mechanisms in embedded and automotive systems. Rather than relying on empirical experimentation or simulation-based validation, this research constructs a multi-layered analytical framework that synthesizes insights from hardware architectures, software reliability techniques, and system-level safety models. This approach enables a holistic understanding

of how different fault-tolerance strategies interact and complement each other in complex computing environments.

The initial phase of the methodology focuses on fault characterization and modeling. Faults are categorized into transient, intermittent, and permanent types, with particular emphasis on transient faults caused by radiation-induced soft errors. The behavior of these faults is analyzed in the context of modern semiconductor technologies, considering factors such as device scaling and environmental exposure (Baumann, 2005). The model also examines the propagation of faults through different system components, including processor cores, memory subsystems, and communication interfaces.

The second phase involves the analysis of hardware-based fault-tolerance mechanisms, with a primary focus on dual-core lockstep architectures. The operational principles of lockstep execution are examined in detail, including synchronization mechanisms, output comparison techniques, and fault detection latency. Variations of the lockstep model, such as delayed comparison and selective redundancy, are also considered. The effectiveness of these approaches is evaluated in terms of fault coverage, performance overhead, and resilience to common-mode failures (de Oliveira et al., 2017).

The third phase introduces software-based fault detection and recovery techniques. This includes an in-depth analysis of checkpoint and rollback mechanisms, control flow integrity enforcement, and memory safety techniques. The methodology examines how these techniques can detect and mitigate faults that are not captured by hardware mechanisms alone. For example, checkpointing enables systems to recover from transient faults by restoring a previously saved state, while control flow integrity ensures that program execution adheres to a predefined structure (Bowen and Pradham, 1993; Yao et al., 2020).

The fourth phase explores system-level safety architectures, such as the Simplex and time-triggered architectures. These frameworks are analyzed in terms of their ability to limit fault propagation and ensure predictable system behavior. The methodology examines how these architectures can be integrated with hardware and software fault-tolerance mechanisms to create a cohesive and robust system design (Crenshaw et al., 2007; Kopetz and Bauer, 2003).

The final phase involves the synthesis of the analyzed components into a unified hybrid architecture. This architecture combines hardware redundancy, software-based detection, and system-level safety mechanisms to achieve comprehensive fault tolerance. The methodology evaluates the overall performance and reliability of this architecture, considering factors such as scalability,

energy consumption, and real-time constraints.

RESULTS

The analytical framework developed in this study reveals several key findings regarding the effectiveness of integrated fault-tolerance mechanisms in embedded and automotive systems. One of the most significant observations is the complementary nature of hardware and software approaches. While hardware-based techniques such as lockstep execution provide rapid and reliable fault detection, they are limited in their ability to address all types of faults, particularly those related to software behavior and memory safety.

The integration of software-based techniques significantly enhances fault coverage by addressing these limitations. Checkpoint and rollback mechanisms enable systems to recover from transient faults without requiring a complete system restart, thereby improving availability and reducing downtime. Similarly, control flow integrity techniques provide an additional layer of protection against software-related errors and security threats.

Another important finding is the role of system-level architectures in limiting fault propagation. The Simplex architecture, for example, effectively isolates faulty components and ensures that a reliable safety controller can maintain system operation. Time-triggered architectures further enhance reliability by ensuring deterministic behavior and reducing the likelihood of timing-related faults.

The analysis also highlights the importance of memory safety mechanisms in modern embedded systems. Techniques such as SoftBound and Bell provide effective solutions for detecting memory-related errors, which are a major source of vulnerabilities. The integration of these techniques into the overall fault-tolerance framework significantly improves system robustness.

DISCUSSION

The findings of this research underscore the importance of adopting a multi-layered approach to fault tolerance in embedded and automotive systems. The integration of hardware, software, and system-level mechanisms provides a more comprehensive solution than any single approach alone. This holistic perspective is particularly important in the context of modern multi-core and cyber-physical systems, where the complexity and interdependence of components create new challenges for reliability and safety.

One of the key implications of this study is the need for architectural diversity to mitigate common-mode failures. While lockstep architectures provide strong fault detection capabilities, their effectiveness can be compromised by correlated failures. The incorporation of

diverse processing elements and software techniques can help address this issue and enhance overall system resilience.

Another important consideration is the trade-off between reliability and performance. While the addition of fault-tolerance mechanisms inevitably introduces some overhead, the benefits in terms of improved safety and reliability often outweigh the costs. However, careful design and optimization are required to ensure that these mechanisms do not adversely impact system performance, particularly in real-time applications.

The scalability of fault-tolerance mechanisms is another critical issue. As systems continue to grow in complexity, the ability to scale these mechanisms effectively becomes increasingly important. This requires the development of adaptive and flexible approaches that can accommodate changing system requirements and operating conditions.

Despite its contributions, this study has certain limitations. The reliance on theoretical analysis limits the ability to quantify the performance and effectiveness of the proposed architecture. Future research should focus on empirical validation and real-world implementation to address this limitation. Additionally, the integration of emerging technologies, such as machine learning-based fault detection, represents a promising area for further exploration.

CONCLUSION

This research provides a comprehensive analysis of fault-tolerant architectures for embedded and automotive systems, emphasizing the integration of hardware redundancy, software-based detection, and system-level safety mechanisms. The findings highlight the limitations of individual approaches and demonstrate the benefits of a unified framework that combines their strengths. By addressing both hardware and software vulnerabilities, the proposed architecture offers a robust solution for ensuring system reliability and safety in increasingly complex environments. Future work should focus on empirical validation, scalability optimization, and the incorporation of intelligent fault-tolerance strategies to meet the evolving demands of next-generation systems.

REFERENCES

1. Azambuja, J. R., Pagliarini, S., Rosa, L., & Kastensmidt, F. L. (2011). Exploring the limitations of software-only techniques in SEE detection coverage. *Journal of Electronic Testing*, 27, 541–550.
2. Baumann, R. C. (2005). Radiation-induced soft errors in advanced semiconductor technologies. *IEEE Transactions on Device and Materials Reliability*, 5(3), 305–316.
3. Bowen, N. S., & Pradham, D. K. (1993). Processor and memory-based checkpoint and rollback recovery. *Computer*, 26(2), 22–31.
4. de Oliveira, Á. B., Rodrigues, G. S., & Kastensmidt, F. L. (2017). Analyzing lockstep dual-core ARM Cortex-A9 soft error mitigation in FreeRTOS applications. *Proceedings of the 30th Symposium on Integrated Circuits and Systems Design*, 84–89.
5. McKusick, M. K., Neville-Neil, G. V., & Watson, R. N. M. The design and implementation of the FreeBSD operating system.
6. Memarian, K., Gomes, V. B., Davis, B., Kell, S., Richardson, A., Watson, R. N., & Sewell, P. (2019). Exploring C semantics and pointer provenance. *Proceedings of the ACM on Programming Languages*, 3.
7. Bond, M. D., & McKinley, K. S. (2006). Bell: bit-encoding online memory leak detection. *ACM SIGARCH Computer Architecture News*, 34(5), 61–72.
8. Nagarakatte, S., Zhao, J., Martin, M. M., & Zdancewic, S. (2009). SoftBound: highly compatible and complete spatial memory safety for C. *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation*, 245–258.
9. Crenshaw, T. L., Gunter, E., Robinson, C. L., Sha, L., & Kumar, P. (2007). The simplex reference model: limiting fault-propagation due to unreliable components in cyber-physical system architectures. *IEEE Real-Time Systems Symposium*, 400–412.
10. Seto, D., Krogh, B., Sha, L., & Chutinan, A. (1998). The simplex architecture for safe online control system upgrades. *American Control Conference*, 3504–3508.
11. Bauer, G., Kopetz, H., & Puschner, P. (2001). Assumption coverage under different failure modes in the time-triggered architecture. *Emerging Technologies and Factory Automation*, 333–341.
12. Kopetz, H., & Bauer, G. (2003). The time-triggered architecture. *Proceedings of the IEEE*, 91(1), 112–126.
13. Lundelius, J., & Lynch, N. (1984). A new fault-tolerant algorithm for clock synchronization. *ACM Symposium on Principles of Distributed Computing*, 75–88.
14. Yao, D., Zhang, Z., & Zhang, G. (2020). Practical control flow integrity using multi-variant execution. *International Conference on Internet Computing for*

Science and Engineering, 14–19.

15. Hilbrich, R., & Dieudonné, L. (2013). Deploying safety-critical applications on complex avionics hardware architectures.
16. Höttger, R., Mackamul, H., Sailer, A., Steghöfer, J.-P., & Tessmer, J. (2017). App 4mc: Application platform project for multi- and many-core systems. *IT-Information Technology*, 59(5), 243–251.
17. NVIDIA. (2021). Nvidia Drive hardware.
18. KALRAY. (2020). Safe compute acceleration for automotive.
19. RENESAS. (2021). R-car-h3-m3-starter-kit.
20. ADLINK. (2021). Adlink AVA-3501.
21. Neosys. (2021). Nuvo-7208VTC.
22. 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>