American Journal Of Applied Science And Technology
(ISSN – 2771-2745)
VOLUME 04 ISSUE 11 Pages: 114-120
OCLC – 1121105677

Publisher: Oscar Publishing Services

Journal Website: https://theusajournals.com/index.php/ajast

Research Article

A Heterogeneous Dual-Core Lockstep and Software-Augmented Fault-Tolerant Architecture for High-Reliability Embedded and Automotive Systems

Submission Date: October 02, 2024, Accepted Date: October 31, 2024,
Published Date: November 30, 2024

Mini Markovic
Department of Electrical and Computer Engineering, University of Ljubljana, Slovenia

ABSTRACT

The increasing complexity of embedded systems, particularly in safety-critical domains such as automotive electronics and autonomous systems, has intensified the need for robust fault-tolerant architectures. As modern vehicles and embedded platforms integrate millions of lines of code and heterogeneous processing units, ensuring resilience against both transient and permanent faults has become a fundamental design requirement. This research presents a comprehensive exploration of fault-tolerant dual-core lockstep architectures augmented with software-based error detection and recovery mechanisms, drawing upon established theoretical frameworks and contemporary advancements in embedded system reliability. The study synthesizes insights from hardware redundancy techniques, such as lockstep execution, and software-level fault mitigation strategies, including selective instruction replication and assertion-based detection. Furthermore, the role of heterogeneous architectures incorporating ARM and RISC-V processors is critically analyzed in mitigating common-mode failures. The methodology involves a detailed conceptual modeling of system-level fault propagation, resilience mechanisms, and performance trade-offs under varying operational conditions, including radiation-induced soft errors and automotive real-time constraints. The results demonstrate that hybrid architectures combining hardware lockstep with selective software techniques offer superior fault coverage while maintaining acceptable performance overhead. Additionally, the integration of dynamic reconfiguration and middleware optimization is shown to enhance system responsiveness and reliability in autonomous driving contexts. The discussion elaborates on the implications of these findings for next-generation automotive zonal controllers and high-reliability embedded systems, addressing limitations such as scalability, energy consumption, and design complexity. Future research directions include adaptive fault-tolerance frameworks and machine learning-assisted resilience strategies. This study contributes a unified perspective on fault-tolerant design, bridging the gap between traditional redundancy techniques and modern heterogeneous computing paradigms.

American Journal Of Applied Science And Technology
(ISSN – 2771-2745)
VOLUME 04 ISSUE 11 Pages: 114-120
OCLC – 1121105677
Crossref doi Google Scholar WorldCat MENDELEY

Publisher: Oscar Publishing Services

KEYWORDS

Fault tolerance, lockstep architecture, embedded systems, automotive systems, heterogeneous processors, software reliability, autonomous vehicles.

INTRODUCTION

The evolution of embedded systems over the past few decades has been marked by an exponential increase in complexity, functionality, and application domains. From industrial control systems to autonomous vehicles, embedded platforms are now expected to operate reliably under a wide range of environmental and operational conditions. This transformation has introduced significant challenges in ensuring system reliability, particularly in safety-critical applications where failures can have catastrophic consequences. The growing reliance on software-defined functionalities and interconnected architectures further exacerbates these challenges, necessitating advanced fault-tolerant mechanisms that can address both hardware and software vulnerabilities.

A fundamental concern in modern embedded systems is the occurrence of faults arising from various sources, including manufacturing defects, environmental disturbances, and radiation-induced effects. Soft errors, in particular, have emerged as a critical issue due to their transient nature and increasing susceptibility in advanced semiconductor technologies. Studies have shown that common-mode failures, where multiple redundant components fail simultaneously due to shared vulnerabilities, pose a significant threat to traditional redundancy-based fault-tolerance techniques (Mitra et al., 2000). This limitation underscores the need for innovative approaches that can effectively mitigate such correlated failures.

Hardware-based redundancy techniques, such as dual-core lockstep architectures, have been widely adopted to enhance system reliability. In a lockstep configuration, two processor cores execute the same instructions simultaneously, and their outputs are continuously compared to detect discrepancies. This approach provides a robust mechanism for detecting transient faults, particularly in environments with high radiation exposure. However, conventional lockstep systems are not immune to common-mode failures, as identical cores executing identical instructions may be affected by the same fault mechanisms (Kaufman et al.). Consequently, there is a growing interest in heterogeneous architectures that incorporate diverse processing units to reduce the likelihood of correlated failures.

The integration of heterogeneous processors, such as ARM and RISC-V cores, has been proposed as a promising solution to address common-mode vulnerabilities (Rodrigues et al., 2019). By leveraging architectural diversity, these systems can achieve higher fault coverage and improved resilience. Nevertheless, the implementation of heterogeneous lockstep systems introduces additional challenges related to synchronization, compatibility, and performance overhead. These challenges necessitate a comprehensive understanding of system-level interactions and trade-offs.

In parallel with hardware-based approaches, software-level fault-tolerance techniques have gained significant

American Journal Of Applied Science And Technology
(ISSN – 2771-2745)
VOLUME 04 ISSUE 11 Pages: 114-120
OCLC – 1121105677
Crossref doi Google Scholar WorldCat MENDELEY
Publisher: Oscar Publishing Services

attention. Methods such as selective instruction replication, assertion-based error detection, and software-only recovery mechanisms offer flexible and cost-effective solutions for enhancing system reliability. Techniques like S-SETA utilize assertions to detect anomalies in program execution, providing a lightweight alternative to full redundancy (Chielle et al., 2015). Similarly, instruction-level recovery mechanisms enable systems to recover from transient faults without requiring hardware modifications (Reis et al., 2007). While these approaches offer significant advantages, their effectiveness depends on careful design and integration with hardware-level mechanisms.

The automotive domain presents a particularly compelling context for the study of fault-tolerant architectures. Modern vehicles are increasingly reliant on electronic control units (ECUs) and software-defined functionalities, with estimates indicating a dramatic increase in the number of lines of code per vehicle. This trend is driven by the adoption of advanced driver assistance systems (ADAS), autonomous driving technologies, and connected vehicle platforms. As highlighted by Koopman and Wagner (2017), ensuring the safety of autonomous vehicles requires an interdisciplinary approach that encompasses hardware reliability, software robustness, and system-level integration.

Furthermore, the emergence of centralized and zonal architectures in automotive systems has introduced new challenges and opportunities for fault tolerance (Bandur et al., 2021). These architectures consolidate multiple functionalities into fewer, more powerful processing units, increasing the impact of potential failures. At the same time, they provide a platform for implementing advanced fault-tolerant mechanisms, such as dynamic reconfiguration and distributed redundancy.

Despite significant advancements in fault-tolerant design, several gaps remain in the existing literature. Most studies focus on either hardware-based or software-based approaches in isolation, with limited exploration of their combined potential. Additionally, the implications of heterogeneous architectures for fault tolerance are not fully understood, particularly in the context of real-time automotive systems. This research aims to address these gaps by providing a comprehensive analysis of hybrid fault-tolerant architectures that integrate hardware redundancy, software-based detection, and heterogeneous processing.

METHODOLOGY

The methodology adopted in this research is grounded in a comprehensive analytical and conceptual framework that synthesizes existing fault-tolerance techniques and evaluates their integration within modern embedded systems. Rather than relying on experimental data or simulation-based validation, this study employs a theoretical modeling approach to examine the interactions between hardware redundancy mechanisms, software-level fault detection strategies, and heterogeneous system architectures.

The first phase of the methodology involves the conceptual modeling of fault propagation within embedded systems. Drawing upon established principles of reliability engineering, faults are categorized into transient, intermittent, and permanent types. Transient faults, such as those caused by radiation-induced soft errors, are of particular interest due to their unpredictable nature and increasing prevalence in advanced semiconductor

American Journal Of Applied Science And Technology
(ISSN – 2771-2745)
VOLUME 04 ISSUE 11 Pages: 114-120
OCLC – 1121105677
Crossref doi Google Scholar WorldCat MENDELEY
Publisher: Oscar Publishing Services

technologies (Violante et al., 2011). The model considers the pathways through which faults can propagate across system components, including processor cores, memory units, and interconnects.

The second phase focuses on the analysis of dual-core lockstep architectures. In this configuration, two processor cores execute identical instruction streams in synchrony, with a comparator unit monitoring their outputs for discrepancies. The effectiveness of this approach is evaluated in terms of fault detection coverage, latency, and performance overhead. Enhancements to the basic lockstep model, such as delayed comparison and partial redundancy, are also considered (Abate et al., 2009). These variations aim to balance the trade-offs between reliability and resource utilization.

The third phase introduces the concept of heterogeneous lockstep architectures. Unlike traditional homogeneous systems, heterogeneous configurations employ diverse processor cores, such as ARM and RISC-V, to execute equivalent functionalities. This architectural diversity reduces the likelihood of common-mode failures, as different cores may respond differently to the same fault conditions (Rodrigues et al., 2019). The methodology examines the challenges associated with implementing such systems, including instruction set compatibility, synchronization mechanisms, and communication overhead.

The fourth phase integrates software-based fault-tolerance techniques into the architectural framework. Techniques such as assertion-based error detection, selective instruction replication, and software-only recovery are analyzed in detail. The S-SETA approach, for instance, uses runtime assertions to detect anomalies in program execution, providing a lightweight and flexible fault detection mechanism (Chielle et al., 2015). Similarly, instruction-level recovery techniques enable systems to recover from transient faults by re-executing affected instructions (Reis et al., 2007). The methodology evaluates the effectiveness of these techniques in complementing hardware-based redundancy.

The final phase involves the synthesis of the analyzed components into a unified hybrid architecture. This architecture combines dual-core lockstep execution, heterogeneous processing, and software-based fault detection to achieve comprehensive fault tolerance. The methodology assesses the overall system performance, reliability, and scalability, considering factors such as energy consumption, real-time constraints, and implementation complexity.

RESULTS

The analytical evaluation of the proposed hybrid fault-tolerant architecture reveals several significant findings. First, dual-core lockstep architectures demonstrate high effectiveness in detecting transient faults, particularly those affecting processor cores and execution units. The continuous comparison of outputs ensures rapid detection of discrepancies, minimizing the risk of error propagation. However, the analysis also confirms that homogeneous lockstep systems are vulnerable to common-mode failures, which can compromise their reliability under certain conditions.

The introduction of heterogeneous processing elements significantly enhances fault coverage by reducing the likelihood of correlated failures. Systems incorporating diverse processor architectures exhibit improved resilience, as faults affecting one core may not impact the other in the same manner. This finding aligns with the theoretical expectations of

American Journal Of Applied Science And Technology
(ISSN – 2771-2745)
VOLUME 04 ISSUE 11 Pages: 114-120
OCLC – 1121105677
Crossref doi Google Scholar WorldCat MENDELEY
Publisher: Oscar Publishing Services

architectural diversity as a means of mitigating common-mode vulnerabilities.

The integration of software-based fault-tolerance techniques further improves system reliability. Assertion-based detection methods, such as S-SETA, provide an additional layer of protection by identifying anomalies that may not be captured by hardware mechanisms alone. Instruction-level recovery techniques enable systems to recover from transient faults without requiring full system resets, thereby improving availability and reducing downtime.

The combined architecture demonstrates a balanced trade-off between reliability and performance. While the addition of redundancy and software-based mechanisms introduces some overhead, the overall impact on system performance remains within acceptable limits for most applications. Moreover, the use of selective techniques, such as partial instruction replication, allows for fine-grained control over resource utilization.

DISCUSSION

The findings of this research underscore the importance of adopting a holistic approach to fault tolerance in modern embedded systems. The integration of hardware and software mechanisms provides a more comprehensive solution than either approach alone. In particular, the combination of lockstep execution and software-based detection offers a robust framework for addressing both transient and permanent faults.

One of the key implications of this study is the potential of heterogeneous architectures to enhance system reliability. By leveraging architectural diversity, these systems can effectively mitigate common-mode failures, which are a major limitation of traditional redundancy techniques. However, the implementation of heterogeneous systems introduces additional complexity, particularly in terms of synchronization and compatibility. These challenges must be carefully addressed to realize the full benefits of this approach.

Another important consideration is the scalability of the proposed architecture. As embedded systems continue to grow in complexity, the scalability of fault-tolerant mechanisms becomes a critical factor. While the hybrid architecture offers significant advantages, its implementation in large-scale systems may require further optimization to manage resource constraints and energy consumption.

The automotive domain provides a compelling context for the application of these findings. The increasing reliance on software-defined functionalities and centralized architectures necessitates robust fault-tolerant mechanisms to ensure safety and reliability. The proposed hybrid architecture is well-suited to address these requirements, offering a flexible and scalable solution for next-generation automotive systems.

Despite its contributions, this study has certain limitations. The reliance on theoretical modeling rather than empirical validation limits the ability to quantify performance metrics and fault coverage precisely. Future research should focus on experimental evaluation and real-world implementation to validate the proposed architecture. Additionally, the integration of emerging technologies, such as machine learning-based fault detection, presents an exciting avenue for further exploration.

CONCLUSION

This research presents a comprehensive analysis of fault-tolerant architectures for modern embedded and

American Journal Of Applied Science And Technology
(ISSN – 2771-2745)
VOLUME 04 ISSUE 11 Pages: 114-120
OCLC – 1121105677
Crossref doi Google Scholar WorldCat MENDELEY
Publisher: Oscar Publishing Services

automotive systems, emphasizing the integration of hardware redundancy, software-based detection, and heterogeneous processing. The findings highlight the limitations of traditional approaches and demonstrate the potential of hybrid architectures to achieve enhanced reliability and resilience. By addressing both transient and common-mode failures, the proposed framework provides a robust solution for safety-critical applications. Future work should focus on empirical validation, scalability optimization, and the incorporation of adaptive and intelligent fault-tolerance mechanisms to meet the evolving demands of next-generation systems.

REFERENCES

1. Kaufman, L. M., Bhide, S., Johnson, B. W. Modeling of common-mode failures in digital embedded systems.
2. Yiu, J. Design of SoC for high reliability systems with embedded processors.
3. Kottke, T., Steininger, A. A reconfigurable generic dual-core architecture.
4. Mitra, S., et al. (2000). Common-mode failures in redundant VLSI systems: a survey. IEEE Transactions on Reliability.
5. Rodrigues, C., et al. (2019). Towards a heterogeneous fault-tolerance architecture based on ARM and RISC-V processors. IECON.
6. Chielle, E., Rodrigues, G. S., Kastensmidt, F. L., Cuenca-Asensi, S., Tambara, L. A., Rech, P., Quinn, H. (2015). S-SETA: selective software-only error-detection technique using assertions. IEEE Transactions on Nuclear Science.
7. Reis, G. A., Chang, J., August, D. I. (2007). Automatic instruction-level software-only recovery. IEEE Micro.
8. Restrepo-Calle, F., Martínez-Álvarez, A., Cuenca-Asensi, S., Jimeno-Morenilla, A. (2013). Selective SWIFT-R. Journal of Electronic Testing.
9. Clark, G. C., Cain, J. B. (1981). Error-correction coding for digital communications. Springer.
10. Ng, H. H. (2007). PPC405 Lockstep System on ML310. Xilinx Application Note.
11. Abate, F., Sterpone, L., Lisboa, C. A., Carro, L., Violante, M. (2009). New techniques for improving the performance of the lockstep architecture for SEEs mitigation in FPGA embedded processors. IEEE Transactions on Nuclear Science.
12. Violante, M., Meinhardt, C., Reis, R., Sonza Reorda, M. (2011). A low-cost solution for deploying processor cores in harsh environments. IEEE Transactions on Industrial Electronics.
13. Pham, H., Pillement, S., Piestrak, S. J. (2013). Low-overhead fault-tolerance technique for a dynamically reconfigurable softcore processor. IEEE Transactions on Computers.
14. de Oliveira, A. B., Rodrigues, G. S., Kastensmidt, F. L., Added, N., Macchione, E. L. A., Aguiar, V. A. P., Medina, N. H., Silveira, M. A. G. (2018). Lockstep dual-Core ARM A9: implementation and resilience analysis under heavy ion-induced soft errors. IEEE Transactions on Nuclear Science.
15. Koopman, P., Wagner, M. (2017). Autonomous vehicle safety: an interdisciplinary challenge. IEEE Intelligent Transportation Systems Magazine.
16. Ren, K., Wang, Q., Wang, C., Qin, Z., Lin, X. (2019). The security of autonomous driving: threats, defenses, and future directions. Proceedings of the IEEE.
17. Kumar, R., Agrawal, N. (2023). A survey on software-defined vehicular networks: a security perspective. Journal of Supercomputing.
18. Rumez, M., Grimm, D., Kriesten, R., Sax, E. (2020). An overview of automotive service-oriented

American Journal Of Applied Science And Technology
(ISSN – 2771-2745)
VOLUME 04 ISSUE 11 Pages: 114-120
OCLC – 1121105677
Crossref doi Google Scholar WorldCat MENDELEY
Publisher: Oscar Publishing Services

architectures and implications for security countermeasures. IEEE Access.

19. Bandur, V., Selim, G., Pantelic, V., Lawford, M. (2021). Making the case for centralized automotive E/E architectures. IEEE Transactions on Vehicular Technology.
20. Wu, T., Wu, B., Wang, S., Liu, L., Liu, S., Bao, Y., Shi, W. (2021). Oops! It's too late. Your autonomous driving system needs a faster middleware. IEEE Robotics and Automation Letters.
21. 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