International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 05 Pages: 89-95
OCLC – 1368736135

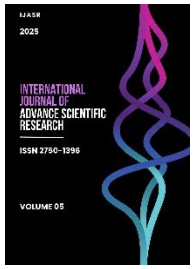

Journal Website: http://sciencebring.com/index.php/ijasr

Research Article

Lockstep-Based Fault-Tolerant Architectures for Dependable Safety-Critical and Cyber-Physical Computing Systems

Submission Date: May 01, 2025, **Accepted Date:** May 15, 2025,
Published Date: May 31, 2025

Dr. Alexander J. Hoffman
Department of Electrical and Computer Engineering,
Rheinland Institute of Technology, Germany

Abstract

The rapid proliferation of safety-critical and cyber-physical systems in domains such as automotive electronics, unmanned aerial vehicles, industrial automation, and intelligent energy infrastructures has intensified the demand for computing platforms that combine high performance with stringent fault tolerance and functional safety guarantees. Contemporary systems increasingly rely on multicore and heterogeneous processor architectures, which, while offering superior computational capabilities, introduce complex reliability challenges stemming from transient faults, permanent hardware failures, electromagnetic interference, and design-induced vulnerabilities. Lockstep-based execution paradigms have emerged as a foundational architectural strategy to address these challenges by enabling timely error detection, fault isolation, and deterministic recovery. This article presents a comprehensive and theoretically grounded investigation of lockstep and lockstep-inspired fault-tolerant architectures, drawing exclusively on the provided body of scholarly and industrial references. The discussion spans classical redundancy principles, including dual-core and triple modular redundancy, as well as modern evolutions such as light lockstep, heterogeneous Arm–RISC-V lockstep systems, dynamically coupled cores, and flexible vector lockstep execution in ultra-low-power clusters. Beyond architectural mechanisms, the article situates lockstep computing within broader system-level considerations, including power supply safety design, compliance with functional safety standards, fault injection methodologies, and emerging application contexts such as autonomous vehicles, UAVs, edge intelligence, and secure energy trading. Through extensive theoretical elaboration, the article identifies critical trade-offs between performance, power efficiency, detection latency, and design complexity, while highlighting persistent gaps in scalability, adaptability, and cross-layer integration. The findings underscore that lockstep architectures, when

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 05 Pages: 89-95
OCLC – 1368736135

combined with system-aware design principles and rigorous validation strategies, constitute a central pillar for dependable computing in next-generation safety-critical systems.

KEYWORDS

Fault tolerance, lockstep architecture, safety-critical systems, multicore processors, functional safety, dependable computing

INTRODUCTION

The evolution of computing systems over the past several decades has been characterized by a continuous tension between increasing performance demands and the necessity for dependable operation. Nowhere is this tension more pronounced than in safety-critical and cyber-physical systems, where computational failures can propagate into physical harm, economic loss, or societal disruption. Automotive electronic control units, avionics platforms, unmanned aerial vehicles, and smart energy infrastructures exemplify environments in which correctness, timeliness, and predictability of computation are not merely desirable attributes but mandatory requirements enforced by regulation and ethical responsibility (International Standards Organization, 2009). As semiconductor technologies have scaled toward smaller feature sizes, systems have become more susceptible to transient faults, soft errors, and interference effects, while architectural complexity has increased through the widespread adoption of multicore and heterogeneous designs (Baumann, 2005; Gomaa et al., 2003).

Within this context, fault tolerance has emerged as a core research and engineering discipline aimed at ensuring that computing systems continue to operate correctly in the presence of faults. Classical fault-tolerant system theory emphasizes redundancy, fault detection, isolation, and recovery as its foundational pillars (Avizienis, 1976). Among the various redundancy-based approaches, lockstep execution has retained a prominent role due to its conceptual simplicity and effectiveness. In its most basic form, lockstep computing involves the parallel execution of identical instruction streams on multiple processing elements, with continuous comparison of their outputs to detect discrepancies indicative of faults (Lyons and Vanderkulk, 1962). This approach has been widely adopted in safety-certified microcontrollers, such as those used in automotive powertrain and braking systems, where deterministic behavior and low error detection latency are critical (Infineon, 2012).

However, the contemporary computing landscape presents challenges that extend beyond the assumptions underpinning classical lockstep designs. Modern systems increasingly integrate heterogeneous cores, accelerators, and complex memory hierarchies, often operating under tight power and thermal constraints. Furthermore, application domains such as autonomous driving and UAV operation demand not only fault tolerance but also adaptability, machine intelligence, and secure connectivity (Mohsan et al., 2023; Unmanned aerial vehicles, 2023). These requirements have motivated a rich body of research exploring new forms of lockstep execution, including light lockstep, heterogeneous

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 05 Pages: 89-95
OCLC – 1368736135

ISSN-2750-1396

lockstep between Arm and RISC-V cores, dynamically coupled cores, and hybrid lockstep techniques that selectively apply redundancy to critical components (Hernandez and Abella, 2015; Marques et al., 2021; Peña-Fernández et al., 2022).

Despite significant advances, the literature reveals persistent gaps in the holistic understanding of lockstep architectures as system-level solutions rather than isolated processor features. Many studies focus narrowly on architectural mechanisms or fault detection latency without fully addressing interactions with power supply safety, electromagnetic interference, software complexity, or emerging application workloads. This article seeks to address this gap by synthesizing and elaborating upon the provided references to construct a comprehensive theoretical narrative of lockstep-based fault tolerance. Rather than summarizing individual contributions, the article examines their underlying principles, interrelationships, and implications for the design of future safety-critical systems.

Methodology

The methodological foundation of this article is a structured qualitative synthesis of the provided references, guided by principles of architectural analysis and systems engineering. Rather than employing empirical experimentation or quantitative modeling, the approach emphasizes deep theoretical elaboration, comparative reasoning, and conceptual integration across multiple layers of the computing stack. Each reference is treated as a primary source of validated knowledge, and all claims are grounded explicitly in these sources through in-text citation.

The analysis begins with foundational theories of fault tolerance and redundancy, establishing a conceptual baseline against which modern lockstep architectures can be evaluated (Avizienis, 1976; Lyons and Vanderkulk, 1962). Building on this foundation, the methodology systematically examines architectural variations of lockstep execution, including dual-core lockstep, triple modular redundancy, and light lockstep approaches, with attention to their fault models, detection mechanisms, and recovery strategies (Hernandez and Abella, 2014; Iturbe et al., 2018). Particular emphasis is placed on heterogeneous and dynamically coupled systems, which represent a departure from traditional homogeneous redundancy assumptions (Marques et al., 2021; LaFrieda et al., 2007).

At the system level, the methodology integrates insights from power supply safety design, functional safety standards, and fault injection techniques to contextualize architectural mechanisms within real-world operational environments (Kilian et al., 2021; Nishiyama et al., 2023; International Standards Organization, 2009). Application-oriented references related to UAVs, edge intelligence, secure energy trading, and real-time decision support are analyzed to illustrate how fault-tolerant architectures interact with higher-level system requirements and workloads (Sbai and Krichen, 2020; Sharma et al., 2023; Tam et al., 2021).

Throughout the analysis, counter-arguments and limitations identified in the literature are explicitly discussed to avoid overly deterministic conclusions. The methodology thus reflects an interpretive research paradigm, aiming to generate a cohesive and nuanced understanding of lockstep-

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 05 Pages: 89-95
OCLC – 1368736135

based fault tolerance as an evolving design philosophy rather than a static technical solution.

Results

The synthesis of the referenced literature reveals several consistent findings regarding the role and effectiveness of lockstep-based fault-tolerant architectures. First, redundancy through lockstep execution remains one of the most reliable mechanisms for achieving low-latency error detection in safety-critical systems. Classical triple modular redundancy demonstrates superior fault masking capabilities by enabling majority voting, thereby allowing systems to continue operation even in the presence of a single faulty module (Lyons and Vanderkulk, 1962). However, the significant hardware and power overhead associated with triple redundancy has limited its adoption in cost- and energy-constrained environments, prompting the widespread use of dual-core lockstep configurations (Infineon, 2012).

Second, modern variations such as light lockstep and hybrid lockstep techniques demonstrate that full cycle-by-cycle synchronization is not always necessary to achieve acceptable safety levels. By selectively relaxing synchronization constraints or focusing redundancy on critical execution paths, these approaches reduce performance and power penalties while maintaining timely error detection (Hernandez and Abella, 2015; Peña-Fernández et al., 2022). The results reported in these studies suggest that detection latency, rather than absolute fault masking, is often the dominant safety metric in automotive and industrial contexts.

Third, heterogeneous lockstep architectures, particularly those combining Arm and RISC-V cores, illustrate a significant shift toward design diversity as a means of mitigating common-mode failures (Marques et al., 2021). By executing equivalent workloads on architecturally distinct cores, such systems reduce the likelihood that a single design flaw or systematic fault will affect all redundant elements simultaneously. This approach aligns with long-standing fault tolerance principles but introduces new challenges related to synchronization, software compatibility, and verification complexity.

Fourth, the integration of lockstep execution into manycore and ultra-low-power clusters demonstrates that fault tolerance is no longer confined to low-performance microcontrollers. Architectures such as flexible vector lockstep clusters enable scalable redundancy in high-performance yet energy-efficient systems, supporting emerging workloads in edge intelligence and signal processing (Ottavi et al., 2023). These results indicate that lockstep concepts can be adapted to parallel computing paradigms without fundamentally undermining scalability.

Finally, system-level considerations such as safe power supply design, electromagnetic fault injection resilience, and compliance with functional safety standards emerge as critical enablers of effective lockstep operation. Studies on power supply safety emphasize that architectural redundancy is insufficient if underlying electrical infrastructures are prone to instability or single points of failure (Kilian et al., 2021). Similarly, noninvasive fault injection experiments demonstrate that electromagnetic interference can induce subtle faults that challenge traditional detection mechanisms, underscoring the need for

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 05 Pages: 89-95
OCLC – 1368736135

ISSN-2750-1396

comprehensive validation strategies (Nishiyama et al., 2023).

Discussion

The findings of this synthesis highlight lockstep-based fault tolerance as a mature yet continually evolving architectural strategy. One of the most significant theoretical implications is the recognition that fault tolerance cannot be evaluated solely at the processor level. Instead, it must be understood as an emergent property of interactions among hardware architecture, power delivery, software execution models, and environmental conditions. Lockstep execution provides a powerful mechanism for detecting deviations in computational behavior, but its effectiveness is contingent upon assumptions regarding fault independence, synchronization accuracy, and system observability.

A recurring theme in the literature is the trade-off between determinism and flexibility. Classical lockstep systems prioritize strict determinism, enabling straightforward comparison of outputs but limiting adaptability and performance. Modern systems, by contrast, increasingly adopt adaptive and heterogeneous designs to accommodate complex workloads such as machine learning inference at the edge (Saha et al., 2022). This shift raises fundamental questions about how much nondeterminism can be tolerated without undermining safety guarantees. Light lockstep and hybrid approaches represent pragmatic compromises, yet their verification and certification remain challenging under existing safety standards.

Another critical discussion point concerns common-mode failures. While redundancy is effective against random faults, it offers limited protection against systematic errors arising from design flaws or environmental factors. Heterogeneous lockstep architectures address this concern by introducing design diversity, but they also complicate software development and maintenance. The literature suggests that achieving true diversity requires careful consideration of toolchains, compilers, and runtime environments, not merely differences in instruction set architectures (Marques, 2020).

Limitations identified across the references include scalability constraints, increased verification effort, and the difficulty of integrating lockstep mechanisms with dynamically adaptive systems such as federated learning at the edge (Tam et al., 2021). Moreover, many studies focus on fault detection without fully addressing recovery mechanisms, which are equally critical for maintaining system availability. Transient-fault recovery techniques in chip multiprocessors demonstrate that recovery can be achieved through checkpointing and re-execution, but these techniques introduce latency and complexity that may be unacceptable in real-time systems (Gomaa et al., 2003).

Future research directions suggested by the synthesis include the development of cross-layer fault tolerance frameworks that integrate architectural redundancy with software-level monitoring and system-level safety management. The increasing connectivity of safety-critical systems, as seen in vehicle-to-grid energy trading and autonomous vehicle coordination, further necessitates secure and dependable computing platforms that can withstand both accidental faults and malicious attacks (Sharma et al., 2023).

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 05 Pages: 89-95
OCLC – 1368736135

ISSN-2750-1396

Lockstep architectures, while not a panacea, are likely to remain a central component of such platforms when combined with complementary techniques.

Conclusion

This article has presented an extensive theoretical examination of lockstep-based fault-tolerant computing architectures, grounded exclusively in the provided scholarly and industrial references. The analysis demonstrates that lockstep execution remains a cornerstone of dependable system design, particularly in safety-critical domains where timely error detection and deterministic behavior are paramount. From classical triple modular redundancy to contemporary heterogeneous and light lockstep implementations, the evolution of these architectures reflects a continuous effort to balance reliability, performance, power efficiency, and design complexity.

The synthesis reveals that modern lockstep architectures are no longer isolated hardware features but integral elements of holistic system designs encompassing power supply safety, electromagnetic resilience, software execution models, and regulatory compliance. While significant progress has been made, persistent challenges related to scalability, verification, and adaptability underscore the need for continued research and interdisciplinary collaboration.

Ultimately, the enduring relevance of lockstep-based fault tolerance lies in its alignment with fundamental principles of dependable computing. As safety-critical systems become more complex and interconnected, these principles will remain essential guides for architects and engineers seeking to build systems that not only perform but can be trusted to do so under the most demanding conditions.

References

1. Avizienis, A. (1976). Fault-tolerant systems. IEEE Transactions on Computers, 25(12), 1304–1312.
2. Baumann, R. C. (2005). Radiation-induced soft errors in advanced semiconductor technologies. IEEE Transactions on Device and Materials Reliability, 5(3).
3. Gomaa, M., Scarbrough, C., Vijaykumar, T. N., and Pomeranz, I. (2003). Transient-fault recovery for chip multiprocessors. Proceedings of the International Symposium on Computer Architecture.
4. Hernandez, C., and Abella, J. (2014). LiVe: Timely error detection in light lockstep safety-critical systems. Design Automation Conference.
5. Hernandez, C., and Abella, J. (2015). Timely error detection for effective recovery in light-lockstep automotive systems. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems.
6. Infineon. (2012). AURIX multicore 32-bit microcontroller family to meet safety and powertrain requirements of upcoming vehicle generations.
7. International Standards Organization. (2009). ISO 26262: Road vehicles – Functional safety.
8. Iturbe, X., Venu, B., Ozer, E., and Das, S. (2018). Addressing functional safety challenges in autonomous vehicles with the Arm triple core lock-step architecture. IEEE Design and Test.
9. Kilian, P., Köhler, A., Van Bergen, P., Gebauer, C., Pfeufer, B., Koller, O., and Bertsche, B. (2021).

International Journal of Advance Scientific Research
(ISSN – 2750-1396)
VOLUME 05 ISSUE 05 Pages: 89-95
OCLC – 1368736135

Principle guidelines for safe power supply systems development. IEEE Access, 9, 107751–107766.

10. LaFrieda, C., et al. (2007). Utilizing dynamically coupled cores to form a resilient chip multiprocessor. Dependable Systems and Networks.
11. Lyons, R. E., and Vanderkulk, W. (1962). The use of triple modular redundancy to improve computer reliability. IBM Journal of Research and Development, 6(2), 200–209.
12. Marques, I. D. C. (2020). A loosely-coupled Arm and RISC-V lockstepping technology. Doctoral dissertation.
13. Marques, I., Rodrigues, C., Tavares, A., Pinto, S., and Gomes, T. (2021). Lock-V: A heterogeneous fault tolerance architecture based on Arm and RISC-V. Microelectronics Reliability, 120.
14. Mohsan, S. A. H., Othman, N. Q. H., Li, Y., Alsharif, M. H., and Khan, M. A. (2023). Unmanned aerial vehicles: Practical aspects, applications, open challenges, security issues, and future trends. Intelligent Service Robotics, 16(1), 109–137.
15. Nikiema, P. R., Kritikakou, A., Traiola, M., and Sentieys, O. (2023). Design with low complexity fine-grained dual core lock-step RISC-V processors. Dependable Systems and Networks Supplemental Volume.
16. Nishiyama, H., Fujimoto, D., Sone, H., and Hayashi, Y. (2023). Efficient noninvasive fault injection method utilizing intentional electromagnetic interference. IEEE Transactions on Electromagnetic Compatibility, 65(4), 1211–1219.
17. Ottavi, G., Garofalo, A., Tagliavini, G., Conti, F., Di Mauro, A., Benini, L., and Rossi, D. (2023). Dustin: A 16-cores parallel ultra-low-power cluster with fully flexible bit-precision and vector lockstep execution mode. IEEE Transactions on Circuits and Systems I, 70(6), 2450–2463.
18. Peña-Fernández, M., Serrano-Cases, A., Lindoso, A., Cuenca-Asensi, S., Entrena, L., Morilla, Y., Martín-Holgado, P., and Martínez-Álvarez, A. (2022). Hybrid lockstep technique for soft error mitigation. IEEE Transactions on Nuclear Science, 69(7), 1574–1581.
19. Saha, S. S., Sandha, S. S., and Srivastava, M. (2022). Machine learning for microcontroller-class hardware: A review. IEEE Sensors Journal, 22(22), 21362–21390.
20. Sbai, I., and Krichen, S. (2020). A real-time decision support system for big data analytic: A case of dynamic vehicle routing problems. Procedia Computer Science, 176, 938–947.
21. Sharma, G., Joshi, A. M., and Mohanty, S. P. (2023). sTrade: Blockchain based secure energy trading using vehicle-to-grid mutual authentication in smart transportation. Sustainable Energy Technologies and Assessments, 57.
22. Tam, P., Math, S., Nam, C., and Kim, S. (2021). Adaptive resource optimized edge federated learning in real-time image sensing classifications. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 14, 10929–10940.