



Holistic Resilience in Modern Embedded Architectures: Soft-Error Mitigation, Traceability, and Fault-Tolerant Design

Ananya R. Mehra

Department of Computer Engineering, Westbridge Institute of Technology

OPEN ACCESS

SUBMITTED 08 December 2023

ACCEPTED 16 December 2023

PUBLISHED 30 December 2023

VOLUME Vol.05 Issue 12 2023

COPYRIGHT

© 2023 Original content from this work may be used under the terms of the creative common's attributes 4.0 License.

Abstract: Modern safety-critical domains — including automotive zonal controllers, aerospace avionics, and industrial control systems — demand embedded computing platforms that deliver high performance while guaranteeing reliability under transient faults and malicious disturbances. Existing literature highlights techniques spanning hardware redundancy, software assertions, trace-based observability, and controller-level mitigation, yet integrating these approaches into cohesive, deployable architectures remains challenging (Entrena et al., 2012; Arifeen et al., 2020).

Objective: This paper synthesizes established and emerging strategies to propose a unified conceptual architecture for resilient heterogeneous embedded systems. The work aims to reconcile the competing objectives of performance, observability, and safety certification by combining selective software-only detection, hardware soft-error controllers, and advanced trace/monitoring infrastructures.

Methods: We perform an in-depth theoretical synthesis of published methods — including selective software assertions (Chielle et al., 2015), soft-error mitigation controllers (Xilinx Inc., 2014), CoreSight trace architectures (ARM Ltd., 2009; 2011), and multilevel emulation-based fault injection (Entrena et al., 2012). We develop a narrative method that maps failure modes to countermeasures and elaborates design patterns for co-design, emphasizing component-level contracts, traceability, and staged mitigation.

Results: The proposed architecture layers lightweight software assertions with hardware error-detection and correction at memory and interconnect levels, integrates a traceable CoreSight-style program-flow telemetry fabric, and places a configurable soft-error mitigation controller at critical fault domains. Analytical

reasoning demonstrates that this composition provides graceful degradation, improved diagnosability, and a path to satisfy stringent safety standards while bounding performance overheads.

Conclusions: A co-design strategy that explicitly couples observability (trace), selective detection (assertions), and hardware mitigation (SEM) yields a pragmatic path toward certifiable, high-performance embedded platforms. Future work should validate the architecture through targeted fault-injection experiments and quantify trade-offs in representative automotive and aerospace workloads.

Keywords: fault tolerance, soft errors, hardware-software co-design, CoreSight trace, assertions, automotive zonal controllers.

Introduction

Safety-critical embedded systems are undergoing rapid change due to increased computational demands, heterogeneous multicore processors, and tighter integration of complex functions in limited physical footprints. Automotive zonal controllers and avionics systems now host workloads that once required separate units, raising the stakes for reliability when a single platform must cope with sensor fusion, control loops, and connectivity. The trend toward higher performance within constrained thermal and power envelopes inherently increases susceptibility to transient faults such as single-event effects (SEEs), neutron-induced latchup, and other radiation-induced phenomena, especially as semiconductor process nodes shrink (Dodd et al., 2003; Baumann, 2005). Concomitantly, adversarial disturbances such as electromagnetic interference are gaining attention as realistic sources of system-level faults (Beckers et al., 2022). These realities require a holistic rethinking of resilience: one that unites hardware-level mitigation, software-level detection, and system-level observability into a coherent engineering methodology.

Prior research offers many effective mechanisms in isolation. Hardware soft-error mitigation controllers and error-correcting codes address bit-flips in memories and configuration bitstreams (Xilinx Inc., 2014), while fault-tolerant topologies such as lockstep cores provide deterministic redundancy for safety-critical control tasks (Abdul Karim, 2023). Software-only techniques, including selective assertion frameworks, can detect faulty program states without full hardware redundancy (Chielle et al., 2015). Trace and telemetry

technologies (CoreSight) add invaluable observability for post-fault analysis and run-time monitoring (ARM Ltd., 2009; 2011). Emulation-based fault injection has been used to determine microprocessor sensitivity to faults and to validate mitigation strategies (Entrena et al., 2012). Yet, gaps remain in how to combine these elements to meet the stringent, often conflicting constraints of performance, cost, certification, and timely fault detection.

This article addresses that gap by presenting an integrated architectural framework that couples selective software assertions, hardware soft-error mitigation, and an observability fabric inspired by CoreSight program flow trace. The framework is motivated by real-world safety requirements and is anchored in the literature's key findings: the efficacy of selective software assertions for detection with limited overhead (Chielle et al., 2015), the practical deployment of SEM controllers for FPGA-based designs (Xilinx Inc., 2014), the necessity of traceability for debugging and certification (ARM Ltd., 2009; 2011), and quantified fault sensitivities from fault-injection studies (Entrena et al., 2012). Through methodical synthesis, we elucidate design patterns, trade-offs, and practical deployment strategies that maintain performance while increasing diagnosability and resilience.

Methodology

The methodological approach adopted here is a structured, theory-driven synthesis of the referenced literature, combined with prescriptive architectural design patterns tailored for embedded safety-critical systems. Given the user's instruction to base the article strictly on the supplied references, the method deliberately draws only on those works and closely allied theoretical inference. The methodology comprises four complementary activities: taxonomy construction, failure-mode mapping, architecture derivation, and analytical evaluation. Each is described below.

Taxonomy Construction: We begin by extracting and classifying the core resilience mechanisms described across the references. Categories include hardware mitigation (e.g., Soft Error Mitigation Controllers, ECC), redundancy topologies (e.g., dual-core lockstep, triple modular redundancy), software detection techniques (e.g., selective assertions), and observability/trace fabrics (CoreSight program flow trace). This classification follows established fault-tolerance

taxonomies and allows a systematic mapping from fault domains to potential countermeasures (Dubrova, 2008; Arifeen et al., 2020).

Failure-Mode Mapping: For each taxonomy element, we map relevant failure modes (temporary bit-flips, latchup, transient control-flow deviations, and adversarial disturbances) to recommended mitigation strategies found in the literature. For example, neutron-induced latchup in SRAMs motivates hardware-level mitigation and shielding strategies (Dodd et al., 2003), while control-flow deviations are amenable to software assertions and program-flow trace for detection and diagnosis (Chielle et al., 2015; ARM Ltd., 2011).

Architecture Derivation: Drawing on the mapping, we synthesize a layered architecture that places detection and mitigation in complementary layers: instrumentation and observability, selective software detection, hardware error correction and mitigation, and system-level decision/containment mechanisms. The CoreSight program flow trace serves as the blueprint for the observability layer, providing non-intrusive telemetry for run-time monitoring and offline forensic analysis (ARM Ltd., 2009; 2011). The Soft Error Mitigation (SEM) controller and ECC modules provide memory and configuration-level correction responsibilities (Xilinx Inc., 2014).

Analytical Evaluation: Without experimental data, the evaluation focuses on analytical reasoning supported by the cited findings. We reason about detection latency, system overhead, diagnosability, and certification traceability. We extrapolate from empirical observations in the literature — such as the measured sensitivity of microprocessors to injected faults (Entrena et al., 2012) and the overheads of assertion-based detection (Chielle et al., 2015) — to bound expected impacts of the integrated design. The evaluation also considers adversarial disturbance vectors as described by Beckers et al. (2022) and discusses how the layered approach can mitigate such threats.

Throughout, we adhere to the constraint of no experimental additions beyond the references; when speculation is necessary to bridge concepts, we explicitly frame it as reasoning or inference grounded in the cited works.

Results

The synthesis yields a conceptual architecture and a

set of concrete design propositions. Here we describe the architecture in detail, explain the expected benefits, and analyze trade-offs and constraints as inferred from the literature.

Architecture Overview

The recommended architecture is a layered, heterogeneous platform that purposefully blends hardware and software resilience mechanisms. At the bottom lies the physical silicon and device-level protections (e.g., layout hardening, process-selection guidance), which are acknowledged but not the central focus here. Above this, we define four principal layers:

1. **Observability Layer (Trace Fabric):** A non-intrusive program-flow and event trace fabric modeled after ARM CoreSight specifications enables continuous telemetry of control-flow events, exceptions, and selected register snapshots (ARM Ltd., 2009; 2011). This layer is the system's "sensorium" and provides both real-time monitoring hooks and post-event forensic capability.
2. **Detection Layer (Selective Software Assertions):** Strategic insertion of software assertions — chosen by criticality and based on program-slicing and fault-impact analysis — provides run-time detection of semantic errors without full replicated execution. The S-seta approach demonstrates that selective assertions, when targeted at high-risk program variables and control points, can catch many salient faults with modest overhead (Chielle et al., 2015).
3. **Hardware Mitigation Layer (SEM, ECC, Redundancy):** For memory arrays, configuration bitstreams, and critical datapaths, hardware-level error-correction mechanisms such as ECC and dedicated Soft Error Mitigation controllers (SEM) provide automated detection and correction. For reconfigurable platforms, Xilinx's SEM controller provides automatic scrubbing and error handling for configuration memory (Xilinx Inc., 2014). For general-purpose cores, combinations of lockstep and selective redundancy can be used for the highest-integrity tasks (Abdul Karim, 2023).
4. **Containment and Response Layer (System Decision Logic):** This topmost layer aggregates

alerts from the trace fabric, assertion violations, and hardware fault counters to perform fail-over, graceful degradation, or controlled restart. Policies for containment must be qualified by safety requirements and real-time constraints.

Expected Benefits and Justification

Observability enables fast detection and rich post-mortem analysis: CoreSight-style tracing provides low-intrusion collection of program counters, branch packets, and exception events, which is invaluable for root-cause analysis and compliance with certification requirements that demand traceability (ARM Ltd., 2009; 2011). By combining these traces with assertion failure data, operators gain context-rich evidence to disambiguate transient faults from design defects.

Selective assertions balance detection coverage against execution overhead: Full instruction duplication or triple modular redundancy (TMR) is expensive in cost and power (Arifeen et al., 2020). In contrast, the S-seta methodology demonstrates that carefully selected assertions can detect a high fraction of failure modes relevant to system safety with significantly lower runtime penalty (Chielle et al., 2015). We infer that a hybrid approach — TMR or lockstep only for ultra-critical control loops while applying selective assertions across less-critical but still safety-relevant software — gives an effective cost-performance compromise.

Hardware mitigation is essential for memory and configuration errors: Empirical findings on neutron-induced latchup and other SEEs (Dodd et al., 2003) and standard practices for FPGA-based designs (Xilinx Inc., 2014) affirm that hardware-level correction and periodic scrubbing are non-negotiable for sustained reliability. The presence of an SEM controller reduces the probability of latent configuration errors and allows swift correction without full system reboot.

Containment policies informed by trace and assertion data reduce false positives and unnecessary failovers: Not all detected anomalies require system shutdown; many transient anomalies can be corrected or tolerated. The decision layer, therefore, must be evidence-based, combining immediate telemetry from CoreSight with assertion context and hardware error counters to choose the least disruptive corrective action that maintains safety.

Analytical Trade-off Evaluation

Detection latency: Assertion-based detection can have

sub-millisecond latency if assertions check lightweight invariants at critical control points; CoreSight telemetry provides near-real-time cues depending on the transport and processing pipeline (ARM Ltd., 2009; 2011). Hardware ECC and SEM correction are immediate for single-bit errors but require scheduled scrubbing for certain multi-bit or configuration errors (Xilinx Inc., 2014).

Performance overheads: Selective assertions impose a runtime cost proportional to assertion frequency and complexity. Chielle et al. report that carefully chosen assertions keep overhead modest while still detecting many faults (Chielle et al., 2015). Hardware mitigation typically has negligible run-time performance impact but uses area and power (Xilinx Inc., 2014). Lockstep redundancy doubles execution resources, which is often unacceptable for zonal controllers that must consolidate many functions (Abdul Karim, 2023).

Diagnosability: Trace plus assertions yield high diagnosability. Entrena et al.'s fault-injection studies emphasize that observability is critical to understanding microprocessor sensitivities (Entrena et al., 2012). The proposed architecture maximizes observable state without imposing prohibitive bandwidth requirements by using selective sampling and event-driven trace capture.

Certiability: Safety certification requires traceability and demonstrable fault-tolerance strategies. The architecture's layered approach aids certification by providing auditable telemetry, deterministic mitigation policies, and documented design-time hazard analyses that map detected events to corrective actions.

Discussion

This section provides deep interpretation of the proposed design, explores limitations, and outlines a research and validation roadmap. We expand on the nuanced relationships among observability, detection, mitigation, performance, and certification.

Interplay Between Observability and Detection

Observability (via CoreSight-style tracing) and detection (via software assertions and hardware error counters) occupy complementary roles. Observability without detection is passive — it gathers data but does not actively flag anomalies. Detection without observability creates blind corrective actions that may lack context, leading to overcorrection or misdiagnosis. When combined, traces provide context to assertion failures (e.g., the program counter and recent branch

history show where an assertion failed), enabling the decision layer to choose context-appropriate responses such as retry, reconfiguration, or graceful degradation.

Selecting assertion targets is a rigorous engineering activity that must be evidence-based. Chielle et al. (2015) emphasize that blanket insertion of assertions is inefficient; instead, assertions should be chosen through program-slicing, static analysis, fault-injection insights, and domain knowledge of safety-critical variables. For example, invariants around sensor fusion outputs or actuator command queues may offer high-value detection points with acceptable overhead. The selection process must also consider potential timing effects: adding assertions can change scheduling and may expose or hide timing-related faults (i.e., Heisenbugs). Therefore, validation must include temporal analyses.

Hardware Mitigation: Scope and Limitations

Hardware-based strategies such as SEM controllers and ECC address many physical fault modes but have bounded capability. They are highly effective against single-bit errors and transient configuration errors; however, multi-bit upsets in dense memories, latchup-induced destructive events, or faults in logic elements require additional mitigations such as spatial redundancy or scrubbing strategies complemented by error logging and workload migration (Xilinx Inc., 2014; Dodd et al., 2003). Designers should avoid over-reliance on a single hardware mechanism; instead, hardware mitigation must be integrated within a defense-in-depth approach.

Redundancy strategies (e.g., lockstep, TMR) provide high assurance but at a cost. Lockstep reduces the need for external voting but requires matching of timing and deterministic behavior; it can be effective for hard real-time control loops where determinism is essential (Abdul Karim, 2023). TMR provides resilience to stuck-at faults and certain Byzantine failures but multiplies hardware costs and power consumption, making it infeasible for many zonal or cost-sensitive platforms (Arifeen et al., 2020). A hybrid approach where redundancy is only applied to the most critical functions yields a practical compromise.

Adversarial and Environmental Disturbances

Beckers et al. (2022) caution that electromagnetic and other adversarial disturbances can produce fault patterns that differ from natural SEEs. These disturbances can target communication links, induce

bit-flips in unexpected places, or cause timing violations. The layered architecture mitigates such threats through multiple lenses: hardware error counters can detect increased error rates suggestive of environmental attack; assertions can detect semantic anomalies; trace data can reveal correlated anomalies across cores or peripherals. Nonetheless, distinguishing malicious disturbances from benign transient faults remains challenging. The architecture should include anomaly-detection heuristics and be designed to preserve forensic evidence for post-incident analysis.

Certification and Standards Considerations

Safety standards for automotive and aerospace domains increasingly demand traceability, demonstrable mitigation strategies, and systematic hazard analyses (Arthur et al., 2022). The proposed architecture maps directly onto these requirements by: (1) offering traceability through the observability layer, (2) providing documented mitigation mechanisms at hardware and software levels, and (3) enabling fault-injection testing to quantify system sensitivity and residual risk (Entrena et al., 2012). Engineers must document the rationale for assertion selection, the configuration of SEM controllers, and the decision policies in the containment layer to satisfy assessors.

Limitations of the Present Synthesis

While the synthesis builds a coherent architectural proposal from the literature, it lacks empirical validation within this work. Entrena et al.'s emulation-based fault injection shows the importance of experimental characterization (Entrena et al., 2012), and the proposed design must be validated through similar fault-injection campaigns across representative workloads. Additionally, the references focus more on specific technologies (e.g., Xilinx SEM) and less on other vendors' offerings; designers must adapt the principles to target platforms. Finally, software assertion strategies require careful mapping to the code base and may not be equally effective across all application domains.

Future Work and Research Agenda

A prioritized roadmap for validation and refinement includes:

1. Controlled Fault-Injection Experiments: Use multilevel emulation to inject faults at device, microarchitectural, and software levels to measure detection coverage and false-positive

rates for the combined assertion-plus-trace approach (Entrena et al., 2012).

2. Real-World Workload Evaluation: Deploy the architecture in automotive zonal controller prototypes (e.g., NXP S32G-based platforms) to evaluate performance and certification readiness (Abdul Karim, 2023).
3. Adversarial Disturbance Testing: Subject prototypes to electromagnetic disturbance and stress tests to quantify detection and containment efficacy as discussed by Beckers et al. (2022).
4. Toolchain and Automation: Develop tooling to automate assertion selection via static/dynamic analysis and to integrate trace ingestion into real-time decision logic, reducing developer burden (Chielle et al., 2015)
5. Certification Case Studies: Work with standards bodies and certification authorities to develop evidence packages that demonstrate compliance through combined trace and mitigation evidence (Arthur et al., 2022).

Conclusion

This article synthesizes established resilience mechanisms to propose a layered, co-designed architecture for safety-critical heterogeneous embedded systems. By combining CoreSight-style observability, selective software assertions, hardware-level soft-error mitigation (SEM, ECC), and a top-level containment and decision logic, the architecture offers a pragmatic path to reconcile high performance with rigorous reliability and certifiability demands. The literature indicates that no single technique is sufficient in isolation: observability without detection leaves systems blind, hardware correction without telemetry hinders diagnosis, and blanket redundancy is too costly for many deployments. The combined approach proposed here leverages the strengths of each mechanism and provides meaningful trade-offs amenable to certification and deployment in domains such as automotive zonal controllers and avionics. Future empirical work—especially emulation-based fault injection, adversarial disturbance testing, and real-world deployment case studies—will be essential to quantify the architecture's benefits and to refine automation for assertion selection and trace analysis. The path ahead integrates rigorous engineering with

evidence-based validation, advancing the dependability of systems on which human safety increasingly depends.

CoreSight Components. Technical Reference Manual, ARM Ltd., DDI0314H, 2009.

References

1. CoreSight Program Flow Trace. Architecture Specification, ARM Ltd., IHI 0035B, 2011.
2. Soft Error Mitigation Controller v4.1. Product Guide, Xilinx Inc., PG036, Nov. 2014.
3. Chielle, E., et al. S-seta: Selective software-only error-detection technique using assertions. *IEEE Transactions on Nuclear Science*, vol. 62, no. 6, pp. 3088–3095, Dec. 2015.
4. Dodd, P. E., et al. Neutron-induced latchup in SRAMs at ground level. In *2003 IEEE International Reliability Physics Symposium Proceedings*, 41st Annual, 2003, pp. 51–55.
5. Domeika, M. *Software Development for Embedded Multi-core Systems. A Practical Guide Using Embedded Intel Architecture*. Elsevier Inc., 2008. ISBN 978-0-7506-8539-9.
6. Dubrova, E. *Fault Tolerant Design: An Introduction*. 2008. Available from: <http://www.pld.ttu.ee/IAF0530/draft.pdf>. [Accessed September 2017].
7. Entrena, L., et al. Soft error sensitivity evaluation of microprocessors by multilevel emulation-based fault injection. *IEEE Transactions on Computers*, vol. 61, no. 3, pp. 313–322, March 2012.
8. ESA. *ESA/SCC Basic specification n. 25100: Single Event Effects Test Method and Guidelines*. Noordwijk, Netherlands, 2005.
9. Alcaide Portet, S., 2023. Hardware/Software solutions to enable the use of high-performance processors in the most stringent safety-critical systems.
10. Abdul Salam Abdul Karim. 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, 2023. Retrieved from <https://ijisae.org/index.php/IJISAE/article/view/7749>

11. Arifeen, T., Hassan, A. S., & Lee, J. A. Approximate triple modular redundancy: A survey. *IEEE Access*, vol. 8, pp. 139851–139867, 2020.
12. Arthur, D., Becker, C., Epstein, A., Uhl, B., & Ranville, S. Foundations of automotive software (No. DOT HS 813 226). United States Department of Transportation, National Highway Traffic Safety Administration, 2022.
13. Beckers, A., Guilley, S., Maurine, P., O'Flynn, C., & Picek, S. (Adversarial) electromagnetic disturbance in the industry. *IEEE Transactions on Computers*, vol. 72, no. 2, pp. 414–422, 2022.
14. Chamorro, W., Sola, J., & Andrade-Cetto, J. Event-based line SLAM in real-time. *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 8146–8153, 2022.