The American Journal of Applied Sciences

ISSN 2689-0992 | Open Access

TYPE Original Research
PAGE NO. 33-38

Check for updates

Toward Resilient and Future-Proof Automotive E/E Architectures: Integrating TSN Ethernet, Cross-Domain Control, and Fault-Tolerant Zonal Computing

OPEN ACCESS

SUBMITTED 01 April 2024
ACCEPTED 13 April 2024
PUBLISHED 26 April 2024
VOLUME Vol.06 Issue 04 2024

Ankit R. Mehra
Department of Electrical and Computer Engineering, Global Institute of Automotive Research, Pune, India

Abstract: Modern automotive electrical/electronic (E/E) architectures are evolving from distributed bus-centric topologies toward zonal, domain-controlled and centralized paradigms that must simultaneously address bandwidth growth, real-time determinism, cybersecurity, and functional safety (Stolz et al., 2010; Brunner et al., 2017; Kugele et al., 2017). Time-Sensitive Networking (TSN) over Ethernet has been proposed to meet deterministic communication needs, while zonal controllers and cross-domain master units are proposed to reduce wiring and increase computational consolidation (Ashjaei et al., 2021; Tavella et al., 2022). At the same time, hardware and software fault modes — including soft errors induced by radiation and single-event upsets — demand architectural approaches such as lockstep dual-core designs and fault injection testing to provide automotive-grade resilience (Abdul Salam Abdul Karim, 2023; Oliveira et al., 2017; Nazar, 2012).

Methods: This article synthesizes the technical themes and empirical findings from the supplied literature to construct a coherent, publication-ready exposition. It develops an integrative conceptual architecture that couples TSN-enabled zonal networks with hypervisor-based consolidation, lockstep dual-core safety islands, intrusion-aware middleware tailored for SOME/IP, and fault-injection informed validation. The methodological narrative contrasts tradeoffs across timing, safety, and security dimensions while deriving design heuristics and verification flows from the referenced research corpus (Gehrmann & Duplys, 2020; Baic et al., 2018; Kugele et al., 2017).

Results: The analysis shows that a structured combination of zonal ECUs, TSN-backboned domain control units, and lockstep redundancy significantly reduces end-to-end latency variance, wiring complexity, and overall system cost while maintaining ASIL-compliant safety targets when paired with selective hardware redundancy and software partitioning (Brunner et al., 2017; Haas & Langjahr, 2016; Abdul Salam Abdul Karim, 2023). However, integrating intrusion detection for SOME/IP and addressing multi-vector attacks remain essential to guard aggregated, high-value centralized controllers (Gehrmann & Duplys, 2020). Fault injection and radiation studies underscore the non-negligible incidence of soft errors in consolidated controllers, demanding a combination of hardware redundancy, error detection/correction, and runtime monitoring (Normand, 2001; Oliveira et al., 2017).

Conclusions: A future-proof E/E architecture is hybrid: zonal consolidation for harness reduction, TSN for timing determinism, cross-domain master controllers for coordinated vehicle behavior, and selective lockstep hardware for safety-critical functions. Architectures must be designed holistically: safety, security, real-time, and validation strategies are interdependent and cannot be treated in isolation (Kugele et al., 2017; Saidi et al., 2018). The article concludes with prescriptive design patterns, a prioritized verification agenda, and research opportunities to close gaps identified in the literature.

Keywords: Zonal E/E, Time-Sensitive Networking, Lockstep, Cross-Domain Control, SOME/IP Intrusion Detection, Fault Injection

Introduction

The last decade has witnessed unprecedented change in automotive electrical/electronic (E/E) architecture design. Vehicle functionality that was once isolated in discrete, low-power Electronic Control Units (ECUs) has converged into high-bandwidth, software-defined features spanning advanced driver assistance systems (ADAS), infotainment, and vehicle dynamics control. This convergence has exposed three interrelated pressures on architecture design: a demand for deterministic high-bandwidth communication, the need for consolidated computing to host complex software stacks securely and safely, and the need for validation techniques that ensure resilience against hardware and software faults (Stolz et al., 2010; Brunner et al., 2017; Kugele et al., 2017).

Time-Sensitive Networking (TSN) has emerged as a pivotal technology to bring deterministic capabilities to Ethernet networks, enabling real-time, scheduled traffic alongside best-effort communication (Ashjaei et al., 2021). TSN is not merely a faster bus; it is a set of standards and mechanisms that provide time-aware shaping, frame preemption, and precise time synchronization — attributes that are critical to automotive control loops. Meanwhile, the zonal architecture — grouping sensors and actuators by physical location and connecting them via localized zonal controllers to domain or central compute units — promises a drastic reduction in wiring harness mass and complexity while enabling partial decentralization of compute workloads (Haas & Langjahr, 2016; Stolz et al., 2010).

However, these transitions introduce complex safety and security interactions. Centralized controllers and cross-domain ECUs aggregate more functionality and increase blast radius for faults and cyber intrusions; therefore, detection, mitigation and architectural segregation strategies must be incorporated from the ground up (Gehrmann & Duplys, 2020; Saidi et al., 2018). Hardware reliability concerns, particularly radiation-induced soft errors and single-event upsets, gain prominence when high-performance processors host mixed-criticality workloads, which has driven research into lockstep dual-core designs and fault-tolerant processor configurations for zonal controllers (Abdul Salam Abdul Karim, 2023; Oliveira et al., 2017; Normand, 2001).

This article undertakes a rigorous, theory-driven synthesis of these elements. Rather than offering another taxonomy, the objective is to construct a practically actionable, academically grounded architecture blueprint that reconciles competing demands: deterministic networking (TSN), zonal consolidation, cross-domain control, intrusion detection for SOME/IP application traffic, and fault tolerance via lockstep and validation by fault injection. The literature provides discrete contributions in each area, but a gap persists in integrated design guidance that maps technology choices to safety and security outcomes — the lacuna this article addresses (Kugele et al., 2017; Ashjaei et al., 2021).

Methodology

The methodological approach of this article is purely synthetic and integrative: it analyses the supplied corpus of primary and secondary references, extracts recurring motifs, quantifies tradeoffs qualitatively, and constructs a cohesive architecture and verification strategy. The method follows three sequential steps.

First, thematic extraction: each reference was read for its factual contribution to one of five pillars — network determinism (TSN), zonal vs. domain/central topologies, cross-domain control strategies, fault tolerance hardware/software (including lockstep), and cybersecurity for vehicle middleware (notably SOME/IP). Observations were attributed via (Author, Year) citations for traceability (Brunner et al., 2017; Ashjaei et al., 2021; Gehrmann & Duplys, 2020; Abdul Salam Abdul Karim, 2023).

Second, cross-mapping and architectural synthesis: the themes were cross-mapped to identify synergies and conflicts. For example, TSN's need for precise time synchronization informs where to place time sources and how to partition traffic classes; zonal architectures create multiple aggregation points which influence intrusion detection placement and fault containment strategies (Brunner et al., 2017; Tavella et al., 2022). This mapping produced architectural variants parameterized by levels of centralization, redundancy approaches, and security posture.

Third, verification and validation flow construction: leveraging fault injection literature and radiation studies, along with lockstep implementation case studies, a verification flow was assembled. This flow prescribes fault injection campaigns targeted at critical timing paths, software-in-the-loop (SiL) and hardware-in-the-loop (HiL) setups for TSN scheduling, and intrusion detection testbeds for SOME/IP traffic flows (Nazar, 2012; Oliveira et al., 2017; Gehrmann & Duplys, 2020).

Throughout, the article avoids empirical fabrication and grounds every major claim in the cited literature: when making prescriptive recommendations, the supporting references are cited to ensure claims remain within the evidence provided by the corpus (Kugele et al., 2017; Saidi et al., 2018).

Results

The results section describes the architecture blueprint, details of component selection, and derived tradeoffs. The narrative translates the synthesis into actionable design patterns and verification sequences.

Architecture Blueprint: A Hybrid Zonal–Domain TSN Backbone

The proposed architecture is a hybrid that blends zonal controllers with domain control units (DCUs) linked by an Ethernet TSN backbone. Sensors and actuators are connected to local zonal ECUs that perform low-latency pre-processing and safety-critical actuation commands. Zonal ECUs aggregate traffic and forward traffic classes over TSN to domain controllers, which host consolidated application software and coordinate cross-domain functions such as propulsion-chassis harmonization (Tavella et al., 2022; Haas & Langjahr, 2016). For the most critical functions (e.g., primary braking or steering control), a lockstep capable safety island — either on a central ECU or on zonal hardware depending on latency constraints — provides hardware redundancy and fault detection (Abdul Salam Abdul Karim, 2023; Baic et al., 2018).

TSN Selection and Traffic Partitioning

TSN is leveraged for deterministic flows: time-sensitive control messages, scheduled sensor fusion data, and synchronization frames. Best-effort Ethernet is reserved for infotainment and non-critical telemetry. The literature shows that TSN's time-aware scheduler, frame preemption, and credit-based shaping can guarantee bounded latency for scheduled flows when network dimensioning aligns with the scheduling horizon (Ashjaei et al., 2021; Brunner et al., 2017). Practically, this requires precise time synchronization (IEEE 802.1AS), disciplined stream policing at ingress points (zonal ECUs), and allocation of guard bands that are calculated with headroom for retransmissions and redundancy.

Lockstep and Safety Islands

For safety-critical tasks requiring ASIL D compliance or equivalent, lockstep dual-core architectures provide a deterministic means of error detection by executing identical instruction streams and comparing results cycle-by-cycle or at defined checkpoints (Abdul Salam Abdul Karim, 2023; Oliveira et al., 2017). The synthesis suggests using lockstep in a selective manner: deploy lockstep cores where short detection latency is essential to avoid hazardous control divergence (e.g., primary braking control), while using software redundancy and monitoring for functions where detection windows can be larger.

Cross-Domain Control Strategies

Cross-domain units, or single master controllers for multiple domains, reduce redundant computation and facilitate coordinated control strategies (Tavella et al., 2022; Haas & Langjahr, 2016). The blueprint recommends hybrid cross-domain mapping: controllers that perform real-time coordination for tightly coupled domains (e.g., propulsion and chassis) but delegate ultra-low-latency I/O to zonal ECUs. This preserves the benefits of reduced software duplication and simplified orchestration without violating latency constraints.

Security — SOME/IP Intrusion Detection and Attack Surface Reduction

SOME/IP is a prevalent in-vehicle middleware for service discovery and remote procedure calls; however, it presents attack vectors when transported over aggregated networks. Intrusion detection tailored to SOME/IP payloads — including semantic inspection and anomaly detection on RPC patterns — is necessary at aggregation points (Gehrmann & Duplys, 2020). The architecture places intrusion detection systems (IDS) at zonal aggregation points and at the TSN edge toward domain controllers. IDS policies must balance false positives against the elevated risk of missing lateral movement in an aggregated system.

Validation and Fault Injection Findings

Fault injection literature and radiation testing indicate that consolidated controllers are vulnerable to single-event upsets and that software optimizations can affect lockstep efficacy (Nazar, 2012; Normand, 2001; Oliveira et al., 2017). The article recommends an integrated validation approach: use fast single-FPGA fault injection platforms to model bit flips and timing perturbations, complement with radiation test data when available, and deploy HiL tests that emphasize both timing (TSN schedule adherence under fault) and functional integrity (lockstep mismatch scenarios) (Nazar, 2012; Oliveira et al., 2017).

Tradeoffs and Quantitative Expectations

While not presenting new empirical data, the synthesis derives qualitative tradeoffs grounded in cited studies. Wiring harness mass and complexity are expected to reduce meaningfully under zonal consolidation, consistent with industry experiences reported by Stolz et al. and Haas et al. (Stolz et al., 2010; Haas & Langjahr, 2016). However, aggregated compute increases the requirement for robust security and verification, raising validation effort and the need for hardware redundancy. TSN provides timing guarantees, but achieving strict latency bounds imposes deterministic scheduling overhead and design discipline in traffic shaping (Ashjaei et al., 2021; Brunner et al., 2017).

Discussion

The discussion unpacks implications, counters potential objections, and outlines a prioritized research agenda.

Holistic Co-Design Imperative

A core conclusion is that E/E architecture design must be co-designed across networking, compute, safety and security. Siloed decisions — e.g., selecting TSN without rethinking intrusion detection placement, or consolidating compute without assessing radiation resilience — will likely produce brittle systems (Kugele et al., 2017; Saidi et al., 2018). The literature reflects a consensus that these domains are coupled; for instance, TSN scheduling choices influence how quickly a safety island can receive a diagnostic heartbeat, and consolidation choices influence where to locate IDS sensors (Ashjaei et al., 2021; Gehrmann & Duplys, 2020).

Selective Redundancy as a Cost-Effective Strategy

A common counter-argument to lockstep or full hardware redundancy is cost and energy overhead. The literature supports a selective redundancy strategy: allocate hardware redundancy where failure modes present intolerable risk or where timing constraints make software recovery infeasible (Abdul Salam Abdul Karim, 2023; Baic et al., 2018). For other functions, such as non-primary actuators or comfort features, software fault-tolerance strategies and watchdog monitoring can suffice. This mixed approach manages cost while aligning safety to function criticality.

SOME/IP IDS — Practical Challenges and Opportunities

Deploying intrusion detection for SOME/IP presents practical constraints: high throughput, encrypted payloads, and evolving RPC semantics complicate deep packet inspection (Gehrmann & Duplys, 2020). The recommended strategy is a layered IDS: lightweight, stateless filters at the network ingress; behavior-based anomaly detection at zonal aggregation; and signature or semantic analyzers at higher compute levels where latency permits. Additionally, IDS systems must be co-designed with TSN scheduling to ensure their monitoring traffic is not preempted or otherwise loses visibility into scheduled flows.

Validation Complexity and the Role of Fault Injection

Fault injection and radiation testing reveal that consolidated controllers face complex failure modes that are sometimes non-intuitive: microarchitectural optimizations can change the manifestation of soft errors, and interactions between software optimizations and lockstep comparisons can create undetected divergences (Oliveira et al., 2017; Normand, 2001). Thus, a verification flow must include: (1) focused fault injection on safety islands and timing paths, (2) system-level HiL tests under scheduled TSN traffic, and (3) iterative tuning of both hardware and software checksums, watchdogs, and mismatch resolution policies. Fast single-FPGA fault injection platforms are effective for early stage validation but must be complemented by radiation data for high-confidence assessments (Nazar, 2012).

Limitations of the Current Synthesis

This article synthesizes and interprets existing literature rather than presenting new experimental data. As such, the recommendations are contingent on the completeness and accuracy of the underlying studies. Some referenced papers are conference-aged or focused on specific implementations; generalizing across manufacturers and diverse vehicle classes introduces uncertainty (Kugele et al., 2017; Saidi et al., 2018). Furthermore, the field is rapidly evolving — for example, TSN standards and implementations continue to mature, and new processor families change the tradeoffs for lockstep and redundancy — so the blueprint should be treated as a design framework rather than an immutable prescription (Ashjaei et al., 2021).

Future Research Opportunities

Several research directions emerge as high priority:

1. **Runtime Adaptive Scheduling for TSN:** Investigate hybrid scheduling that adapts to system state (e.g., degraded mode) while preserving deterministic bounds. This could reduce overprovisioning of guard bands while maintaining safety. The literature suggests significant opportunity but limited practical implementations to date (Ashjaei et al., 2021)
2. **Lightweight SOME/IP Semantic Models for IDS:** Develop compact, learnable semantic models of RPC patterns and service dependencies that enable high-fidelity intrusion detection with low resource cost at zonal ECUs (Gehrmann & Duplys, 2020).
3. **Holistic Verification Platforms:** Create integrated HiL testbeds that combine TSN schedule verification, fault injection, and IDS testing in a single, repeatable workflow to capture cross-domain emergent behaviors (Nazar, 2012; Oliveira et al., 2017).
4. **Security and Safety Co-Validation Methods:** Methodologies that jointly validate security properties and safety compliance when a system is under attack or in degraded mode remain underdeveloped (Saidi et al., 2018).

Practical Implementation Guidance

Manufacturers and system integrators should adopt an incremental rollout strategy: start by deploying zonal controllers with basic TSN scheduling and conservative traffic partitioning; gradually consolidate functions to domain controllers while instrumenting IDS and increasing safety island coverage for the most critical functions. Simultaneously, build a verification pipeline that includes both traditionally required ASIL tests and security testbeds for lateral movement and SOME/IP semantics (Haas & Langjahr, 2016; Gehrmann & Duplys, 2020).

Conclusion

This article synthesizes an evidence-based blueprint for future automotive E/E architectures that integrates TSN, zonal consolidation, cross-domain control, lockstep fault tolerance, and intrusion detection tailored to SOME/IP. The literature supports a hybrid approach: use zonal ECUs for locality and harness reduction, TSN for deterministic networking, and selective lockstep redundancy for the highest safety integrity functions (Stolz et al., 2010; Brunner et al., 2017; Abdul Salam Abdul Karim, 2023). Security — especially for middleware like SOME/IP — and validation via fault injection and radiation-informed testing are non-negotiable elements of a robust architecture (Gehrmann & Duplys, 2020; Oliveira et al., 2017; Nazar, 2012).

The principal message for designers is clear: architectural choices are interdependent, and achieving a future-proof E/E architecture requires co-design across networking, compute, safety, security, and validation workflows. The proposed blueprint and verification agenda outline a pragmatic path forward, but the dynamic nature of automotive hardware, standards (e.g., TSN), and threat landscapes necessitates continuous reassessment and empirical validation in the field.

References

1. Brunner, S., Roder, J., Kucera, M., et al.: Automotive E/E-architecture enhancements by usage of

ethernet TSN. Paper presented at 13th Workshop on Intelligent Solutions in Embedded Systems (WISES), Hamburg, Germany, 12–13 June 2017.

2. Gehrmann, T., Duplys, P.: Intrusion detection for SOME/IP: Challenges and opportunities. Paper presented at 2020 23rd Euromicro Conference on Digital System Design (DSD), Kranj, Slovenia, 26–28 August 2020.
3. Kugele, S., Cebotari, V., Gleirscher, M., et al.: Research challenges for a future-proof E/E architecture. Paper presented at 15th Workshop Automotive Software Engineering, Lecture Notes in Informatics (LNI), Gesellschaft für Informatik (GI), Bonn, 1 September 2017.
4. Ashjaei, M., Bello, L.L., Daneshtalab, M., et al.: Time-sensitive networking in automotive embedded systems: state of the art and research opportunities. J. Syst. Architect. 117, 102137 (2021). https://doi.org/10.1016/j.sysarc.2021.102137
5. Stolz, W., Kornhaas, R., Krause, R., et al.: Domain control units — the solution for future E/E architectures? SAE Int. (2010). https://doi.org/10.4271/2010-01-0686
6. Haas, W., Langjahr, P.: Cross-domain vehicle control units in modern E/E architectures. Paper presented at 16 Internationales Stuttgarter Symposium: Automobil-und Motorentechnik, Springer Fachmedien Wiesbaden, 2016.
7. Baic, D., Langjahr, P., Haas, W., et al.: Safe computing with central ECUs. Paper presented at 18 Internationales Stuttgarter Symposium: Automobil-und Motorentechnik, Springer Fachmedien Wiesbaden, 2018.
8. Saidi, S., Steinhorst, S., Hamann, A., et al.: Special session: Future automotive systems design: Research challenges and opportunities. Paper presented at 2018 International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS), Turin, Italy, 30 September 2018 — 05 October 2018.
9. 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
10. Tavella, D., Tolkacz, J., Wasacz, B., et al.: Cross-domain control architecture — Single master controller for propulsion and chassis automotive domains. SAE Tech. Paper. 1–9 (2022). https://doi.org/10.4271/2022-01-0746
11. Yu, D.: The digital foundation platform — A multi-layered SOA architecture for intelligent connected vehicle operating system. arXiv preprint arXiv. 1–11 (2022). https://doi.org/10.4271/2022-01-0107
12. NAZAR, G. L.; CARRO, L. Fast single-FPGA fault injection platform. In: 2012 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT). [S.l.: s.n.], 2012. p. 152–157. ISSN 1550-5774.
13. NORMAND, E. Correlation of inflight neutron dosimeter and SEU measurements with atmospheric neutron model. IEEE Transactions on Nuclear Science, v. 48, n. 6, p. 1996–2003, Dec 2001. ISSN 0018-9499.
14. OLIVEIRA, Á. B. de et al. Analyzing the impact of software optimizations in lockstep dual-core ARM A9 under heavy ion induced soft errors. In: European Conference on Radiation and Its Effects on Components and Systems (RADECS). [S.l.: s.n.], 2017. p. 1–4.
15. OLIVEIRA, Á. B. de; TAMBARA, L. A.; KASTENSMIDT, F. L. Applying lockstep in dual-core ARM Cortex-A9 to mitigate radiation-induced soft errors. In: 2017 IEEE 8th Latin American Symposium on Circuits Systems (LASCAS). [s.n.], 2017. p. 1–4. Available from Internet: http://dx.doi.org/10.1109/LASCAS.2017.7948063