The OSCAR PUBLISHING Services
American Journal of Applied Science and Technology

VOLUME Vol.05 Issue 10 2025
PAGE NO. 262-269

# Integrated 10G Automotive Ethernet Architectures for Secure, EMC-Resilient ADAS and V2X Systems: A Multi-Protocol, Post-Quantum-Aware Framework

John A. Mercer
Department of Electrical and Computer Engineering, Northbridge Technical University

**Received:** 07 October 2025 **Accepted:** 19 October 2025; **Published:** 31 October 2025

**Abstract:** Background: Modern vehicles increasingly rely on high-bandwidth, low-latency communication fabrics to support Advanced Driver-Assistance Systems (ADAS), camera and sensor subsystems, and vehicle-to-everything (V2X) services. The migration toward 10 Gigabit Automotive Ethernet (10G AE) raises technical challenges across electromagnetic compatibility (EMC), multi-protocol integration, cyber resilience, and the software/hardware coordination of electronic control units (ECUs) and inertial measurement units (IMUs). Simultaneously, optical and high-speed digital transmission research offers perspectives on managing error rates and link stability. This paper synthesizes evidence from Ethernet-for-automotive analyses, V2X gateway implementations, ADAS ECU/IMU roles, EMC mitigation case studies, and post-quantum cryptography considerations to propose an integrated architecture and evaluation narrative. (Ioana et al., 2022; Kern, 2013; Katari et al., 2024; KARIM, 2025; Prest et al., 2017).

Methods: We conduct a methodical cross-domain synthesis of published engineering analyses and experimental work, mapping known technical constraints onto a cohesive design space for 10G Automotive Ethernet deployments within ADAS and V2X ecosystems. We organize the synthesis by physical layer concerns (EMC, shielding, cabling), link and protocol conversion (multi-protocol gateways), ECU/IMU data flows and timing, error control strategies, and cryptographic readiness including post-quantum algorithms. (Ioana et al., 2022; KARIM, 2025; A. Naughton et al., 2012; Prest et al., 2017).

Results: The integrated framework emphasizes: (1) hybrid shielding and PCB layout validated EMC approaches for camera and lighting control to minimize 10G interference risks; (2) multi-protocol gateway designs that preserve timing and determinism necessary for ADAS sensor fusion; (3) a layered error-management strategy incorporating forward error correction learning from optical comms studies; and (4) a migration path for incorporating post-quantum cryptographic primitives such as FALCON into vehicular keying infrastructures. (KARIM, 2025; Ioana et al., 2022; Brandonisio et al., 2017; Prest et al., 2017).

Conclusion: Transitioning to 10G AE in safety-critical vehicles requires a systems-level approach that integrates EMC-aware hardware design, protocol bridging that preserves deterministic behaviors, robust error mitigation drawing from optical network research, and anticipatory cryptographic upgrades. We identify research priorities and practical engineering guidelines to operationalize this transition while maintaining ADAS safety and V2X security requirements. (Kern, 2013; Katari et al., 2024).

**Keywords:** Automotive Ethernet, 10G, ADAS, V2X, EMC mitigation, Multi-protocol gateway, Post-quantum cryptography

## INTRODUCTION:

The contemporary vehicle is evolving from a collection of mechanically coupled subsystems toward a distributed cyber-physical system where network bandwidth, latency, timing determinism, and security collectively determine functional safety and user experience. The push toward high-resolution cameras, LiDAR, radar, and the dense sensor arrays demanded by Advanced Driver-

Assistance Systems (ADAS) — alongside the expanding domain of vehicle-to-everything (V2X) communications — makes the adoption of higher-rate in-vehicle networks both attractive and necessary. Historically, automotive network topologies relied on fieldbuses and lower-rate Ethernet variants; however, the performance requirements of modern ADAS and V2X use cases increasingly point to 10 Gigabit Automotive Ethernet (10G AE) as a viable backbone (Kern, 2013; Ioana et al., 2022).

The migration to 10G AE is not merely a matter of replacing cables and switches with higher-speed versions. It invokes a cascade of interdependent engineering issues: electromagnetic interference (EMI) and electromagnetic compatibility (EMC) become more challenging as data rates increase and signal edges become sharper; printed circuit board (PCB) layout and shield design for camera modules and associated lighting controllers require re-evaluation (KARIM, 2025). Simultaneously, the in-vehicle ecosystem is heterogeneous: CAN, LIN, FlexRay, MOST, and existing automotive Ethernet variants must interoperate with new 10G links, often realized via multi-protocol gateways that convert and preserve semantics and timing for ADAS-relevant messages (Ioana et al., 2022). On the software and algorithmic side, ECUs and IMUs play pivotal roles in sensor fusion, decision timelines, and control loops that are ultimately constrained by network characteristics (Katari et al., 2024). Moreover, the security landscape is evolving: the cryptographic primitives that protect vehicle communications face prospective threats from advances in quantum computing; post-quantum algorithms such as FALCON are being standardized for long-term security (Prest et al., 2017), and their integration into vehicular systems is a nascent challenge.

This article presents a comprehensive, theoretically grounded synthesis that draws primarily on the provided literature to delineate an integrative architecture for deployment of 10G AE in ADAS and V2X contexts. We frame the problem by exposing the critical technical tensions: bandwidth versus determinism, EMC-safety tradeoffs in PCB and cable design, multi-protocol semantic preservation, and cryptographic resilience in a post-quantum horizon. Building on empirical and analytical insights from automotive Ethernet deployments and optical/high-speed communications research, we propose a layered framework for hardware, firmware, and cryptographic adaptation that targets maintainable and certifiable ADAS performance.

The literature reveals several distinct but related gaps. First, while Ethernet migration studies and doctoral analyses provide strong conceptual frameworks for integrating IP/Ethernet into automotive electrical/electronic (E/E) architectures (Kern, 2013), they predate the practical engineering realities of 10G AE and its unique EMC challenges, particularly when camera PCBs and lighting controllers coexist on the same vehicle bus. Second, applied research into V2X multi-protocol gateways demonstrates proof-of-concept conversions and prototype deployments (Ioana et al., 2022), yet there is limited synthesis linking gateway design to ADAS timing constraints and to the hardware-level measures required for 10G operation. Third, while optical communications research and practical networking studies provide techniques for error control at 10 Gbps and above (A. Naughton et al., 2012; Li et al., 2015; Brandonisio et al., 2017), their translations to electrically wired automotive domains require careful adaptation. Fourth, the nascent literature on mitigating EMI in 10G Automotive Ethernet offers promising case studies (KARIM, 2025) but lacks a unified, cross-layer strategy that combines PCB shielding, connector design, link coding, and higher-level protocol strategies. Finally, the convergence with cryptography — specifically, how to prepare vehicular keying infrastructures for post-quantum algorithms without compromising real-time performance — remains underexplored in an automotive-grade engineering context (Prest et al., 2017).

This paper addresses these gaps by synthesizing these diverse threads into a coherent set of design principles, demonstrating how automotive engineers can navigate the multi-dimensional trade-space that 10G AE introduces for ADAS and V2X. We do so by explicating a methods framework for integrated systems analysis and then presenting detailed descriptive results and interpretative discussion that link low-level EMC and link reliability concerns to high-level safety and cryptographic imperatives.

**Methodology**

This work adopts a systems synthesis methodology, combining structured literature synthesis, comparative analysis, and conceptual frameworks to translate discrete technical findings into an integrated design narrative. The methodology is purely text-analytic and theoretical — consistent with the constraints of relying strictly on the supplied literature — and proceeds in four stages: corpus consolidation, cross-domain mapping, framework construction, and evaluative reasoning.

Corpus Consolidation: We begin by consolidating the provided references into domain clusters: (A)

Ethernet and IP migration studies relevant to automotive E/E architectures (e.g., Kern, 2013); (B) multi-protocol V2X gateway research that informs cross-domain message conversion (Ioana et al., 2022); (C) ADAS component roles and timing — especially relating to ECUs and IMUs (Katari et al., 2024); (D) EMC mitigation and PCB shielding studies for 10G Automotive Ethernet (KARIM, 2025); (E) high-speed optical and electrical link research bearing on error control and link stability (A. Naughton et al., 2012; Li et al., 2015; Brandonisio et al., 2017; Biswas et al., 2016); and (F) cryptographic readiness for post-quantum threats, with emphasis on FALCON and NIST activities (Prest et al., 2017; PQC-Forum). Consolidating the corpus provides discipline-specific insights that can be cross-referenced.

Cross-Domain Mapping: Each cluster was examined to identify the principal constraints, proposed mitigations, and open questions. For instance, the EMC mitigations recommended by shielding validation studies were mapped to the camera PCB and lighting controller contexts described in KARIM (2025). Gateway designs described by Ioana et al. (2022) were analyzed concerning timing preservation and message determinism required by ADAS ECUs and IMUs (Katari et al., 2024). Optical communication insights regarding forward error correction and burst-mode transmission were mapped to the 10G AE domain to derive analogies and candidate strategies. The cross-domain mapping emphasizes identifying where approaches in one domain (e.g., optical FEC techniques) can be meaningfully adapted to automotive electrical link problems.

Framework Construction: We construct a layered architecture that specifies requirements and mitigation strategies at the physical, link, gateway, and security layers. For each layer the framework enumerates: (1) core objectives (e.g., minimize jitter for sensor fusion); (2) primary hazards or constraints (e.g., EMC-induced bit errors); (3) proposed mitigations (e.g., shielded differential pair routing, duplex link configurations, FEC schemes); and (4) evaluation criteria (e.g., bit error rate thresholds consistent with ADAS control timelines). The framework draws its justifications from the corpus and aligns with the migration concepts proposed by Kern (2013) and the empirical EMC studies of KARIM (2025).

Evaluative Reasoning and Deliberative Synthesis: The final methodological step is a reasoned evaluation that brings together hardware, protocol, and cryptographic considerations. This involves layered trade-space analyses — for instance, quantifying qualitatively how tighter shielding might raise cost and weight but reduce bit error incidents that otherwise necessitate more aggressive link-level FEC, which in turn imposes latency. While no original experimental measurements are produced, the evaluative reasoning aims to be rigorous and conservative: it cites known empirical results (e.g., optical experiments and enterprise networking migration studies) and extrapolates those results into the automotive scenario with clearly stated inferential steps and uncertainties.

Throughout, every major claim and recommendation is anchored to at least one citation from the provided corpus. Where inferences extend beyond the literal text of a citation, we explicitly mark the reasoning as an adaptation or hypothesis driven by the synthesized literature.

**Results**

The results section presents the synthesized outcomes of the methodological process: a detailed layered architecture for integrating 10G Automotive Ethernet into ADAS and V2X systems, concrete hardware and protocol design recommendations, security adaptation strategies with post-quantum readiness, and an enumerated set of research priorities and evaluation criteria.

Layered Architecture Overview: The architecture is organized into four primary layers: Physical (PHY) and EMC, Link and Encoding, Gateway and Middleware, and Security and Key Management. Each layer is described below along with the specific recommendations and supporting citations.

Physical (PHY) and EMC Layer — Overview and Recommendations: As deployment moves toward 10G AE, the physical layer becomes the primary battleground for ensuring reliable operation. KARIM (2025) demonstrates that EMI issues at 10 Gbps are nontrivial for camera PCBs and lighting control circuits, and that shielded layouts validated via HyperLynx or equivalent electromagnetic simulation tools can significantly mitigate emissions and susceptibility risks. The following recommendations synthesize those findings:

• PCB Shielding and Grounding: Adopt multilayer PCB stacks for camera modules with dedicated ground planes, careful via stitching around high-speed traces, and local decoupling strategies. KARIM (2025) emphasizes shielding the camera PCB and camera interface connectors, which reduces common-mode radiation and coupling into adjacent systems.

• Differential Pair Routing and Impedance Control: Maintain strict differential pair length matching and impedance control for 10G traces; employ controlled dielectrics and constraining connectors to preserve

signal integrity. These are established practices in high-speed design and are implicitly supported by migration analyses that stress physical infrastructure (Kern, 2013).

• Shielded Cabling and Connectorization: Use shielded twisted pairs or, where weight and EMC considerations warrant, shielded twinax or coaxial assemblies with automotive-grade connectors. Field experience suggests that robust connector shielding reduces discontinuities that can radiate or reflect high-frequency components (KARIM, 2025).

• Thermal and Mechanical Coupling: Design mechanical housings to avoid ground loops and unintended resonances that can degrade EMC performance. EMC mitigation is holistic: mechanical design choices directly influence electromagnetic behavior (KARIM, 2025).

These measures collectively reduce the probability of transient or persistent bit errors induced by EMI and create a more stable baseline for higher-layer error management. The recommendations echo the thesis of Kern (2013) that physical infrastructure decisions fundamentally shape the migration potential of Ethernet within automotive E/E architectures.

Link and Encoding Layer — Error Management and Timing: 10G AE introduces link coding and error management choices that must reconcile throughput with latency and jitter constraints vital for ADAS:

• Error Control Strategy: Adopt a layered error mitigation approach that combines robust physical shielding (as above) with link-level techniques. Although optical network studies operate in a different medium, their work on forward error correction (FEC) and burst-mode transmission provides instructive analogies. Brandonisio et al. (2017) analyze FEC for 10Gb/s burst-mode transmission and indicate how strategic FEC placement can reduce effective bit error rates. While automotive Ethernet will not necessarily use the exact FEC schemes from optical networks, the principle—deploy FEC where residual bit error rates threaten application-level correctness—applies.

• Latency-Aware FEC: Any FEC or error correction must be tuned to ADAS timing constraints. Katari et al. (2024) underline the sensitivity of fusion and control loops to latency. Therefore, lightweight, low-latency FEC (or hybrid ARQ with strict latency bounds) is preferable to heavyweight error correction that would induce unacceptable delay.

• Determinism and Time-Sensitive Networking (TSN): Integrate TSN profiles that ensure bounded latency and prioritized scheduling for ADAS traffic. Kern (2013) and subsequent Ethernet adoption studies document the necessity of deterministic Ethernet extensions for automotive control traffic. Gateways and switches must implement TSN features to ensure that high-priority sensor and actuator messages traverse the 10G backbone with guaranteed bounds.

• Link Monitoring and Adaptive Rate Control: Implement continuous link health monitoring that can detect rising error rates or EMI events and trigger adaptive responses (e.g., increasing redundancy, switching to fail-safe modes, or rerouting via alternate paths). Lessons from live migration studies over 10Gb/s interfaces (Biswas et al., 2016) show that network performance fluctuations necessitate adaptive workflows; for vehicles, this must occur within safety constraints.

Gateway and Middleware Layer — Multi-Protocol Integration and Timing Preservation: V2X scenarios and in-vehicle heterogeneity require gateways that preserve semantics, timing, and safety properties when translating across protocols:

• Multi-Protocol Gateway Design: Ioana et al. (2022) present Ethernet-based communication technologies applied in a V2X context via a multi-protocol gateway. Building on their findings, gateways should be architected to maintain strict timing budgets for ADAS messages, perform protocol buffering that minimizes jitter, and support mapping rules that preserve critical metadata (timestamps, sequence numbers, QoS markers).

• Semantic Preservation and Message Prioritization: Gateways must not discard or re-order messages in ways that break sensor fusion or control assumptions. For instance, timestamp preservation and monotonic sequence delivery for IMU and camera data are essential for accurate sensor fusion (Katari et al., 2024). Gateways should maintain transparent timestamp translation or, when translation is necessary, provide explicit metadata and bounded correction windows.

• Gateway Hardware Acceleration: To support minimal latency, particularly in protocol translation, gateways should utilize hardware acceleration (FPGA or ASIC) for common translation paths. Ioana et al. (2022) show practical success in applying Ethernet-centric gateways in V2X; accelerating critical paths ensures the gateway does not become the dominant latency source.

• Security and Isolation: Gateways serve as security choke points; their design must include robust partitioning, secure boot sequences, and authenticated update mechanisms. This requirement dovetails with the security layer recommendations

below (Prest et al., 2017).

Security and Key Management Layer — Post-Quantum Readiness and Practical Integration: Vehicular networks require secure boot, authenticated messaging, and secure key management. The emerging post-quantum cryptographic (PQC) landscape influences how automotive systems should prepare:

• Anticipatory Cryptography: Prest et al. (2017) document FALCON within the NIST PQC project; adopting post-quantum algorithms preemptively for non-real-time or long-lived artifacts (e.g., firmware signing, long-term key exchange for OTA updates) reduces future rework. For latency-sensitive control paths, hybrid schemes that combine classical primitives with PQC variants can balance performance and forward secrecy.

• Hybrid Keying Strategies: Implement hybrid key exchange and signature schemes that pair well-understood classical algorithms with PQC candidates like FALCON for initial key establishment and for update signing where latency is acceptable. This staged approach aligns with broader NIST recommendations and PQC forum discussions (Prest et al., 2017; PQC-Forum).

• Performance Budgeting: PQC algorithms can have larger key sizes and higher computational costs; therefore, cryptographic operations must be budgeted into ECU cycles and gateway acceleration hardware. Karim (2025) does not directly address cryptography, but the hardware mitigation techniques for EMC invariably influence computational resource availability on constrained boards; cryptographic budgets must be accounted for early in hardware partitioning.

• Lifecycle and Update Management: Secure, authenticated over-the-air (OTA) update pipelines are fundamental to moving toward PQC. Migration strategies should emphasize secure boot chains and the ability to install larger PQC keys and algorithms without compromising safety during transition.

Synthesis of Cross-Layer Tradeoffs: The architecture above is deliberately layered to reveal interdependencies. For instance, more aggressive physical shielding reduces residual bit errors, which reduces the need for latency-inducing FEC at the link layer. However, improved shielding can raise cost, weight, and thermal constraints. Conversely, reliance on heavy link-level correction simplifies physical constraints but can compromise latency budgets and thus ADAS control performance. These tradeoffs must be evaluated on a use-case basis; for safety-critical subsystems like camera feeds used in emergency braking, the architecture favors physical investments plus low-latency FEC and TSN guarantees rather than deferred correction that risks transient control faults (Katari et al., 2024; KARIM, 2025).

Adaptation of Optical Network Lessons: The corpus contains several optical and high-speed network studies whose findings inform automotive strategies. Naughton et al. (2012) and Li et al. (2015) show that advanced modulation and error correction techniques can extend link reach and improve resilience in high-speed contexts. While automotive wiring differs materially from optical fiber, the conceptual lesson—that properly chosen modulation, coding, and error correction can reduce raw bit error rates and improve effective throughput—translates. Brandonisio et al. (2017) and Biswas et al. (2016) reinforce that forward error correction and link monitoring are practical and beneficial at 10Gb/s scales in other domains, suggesting their selective adoption in automotive 10G deployments with latency constraints in mind.

Quantitative Evaluation Criteria (Descriptive): Based on the literature, we propose the following evaluation criteria for design decisions (note that these are presented descriptively, not as measured values):

• Effective Bit Error Rate (BER): Target BER levels should be compatible with ADAS application tolerances; physical design and link corrections should aim for BER levels that keep retransmission or correction latency within control loop budgets (Katari et al., 2024).

• End-to-End Latency and Jitter: For sensor fusion tasks, end-to-end latencies must satisfy the temporal requirements of control algorithms. TSN and gateway translation must be evaluated against these budgets (Kern, 2013).

• EMC Susceptibility Thresholds: System designs must document susceptibility margins to identified vehicle-level EMI sources (ignition systems, power electronics, external RF sources), with HyperLynx or similar simulations validating margin adequacy (KARIM, 2025).

• Cryptographic Lifecycle Readiness: Systems should be able to accept and use PQC keys and signatures and have hardware or software pathways to upgrade cryptographic suites without compromising safety or boot integrity (Prest et al., 2017).

**Discussion**

The synthesis presented in the Results section underscores that successful deployment of 10G AE within ADAS and V2X contexts is a multi-dimensional

engineering challenge. This discussion expands upon the theoretical implications, practical tradeoffs, limitations of the current literature, and directions for future research.

Theoretical Implications and Systems Thinking: One central theoretical implication is the necessity of co-design: physical hardware design choices (e.g., shielding thickness, PCB stackup) have direct consequences on higher-layer protocol choices (e.g., whether to implement aggressive FEC). This interdependence exemplifies classical systems engineering principles but has particular urgency in automotive contexts because safety is non-negotiable. The layered architecture emphasizes that no single mitigation is sufficient; rather, engineering rigor must be distributed across layers with a measured allocation of redundancy and defense-in-depth.

From an information-theoretic perspective, increasing raw channel capacity (10G) reduces the headroom per bit for noise and interference; margins tighten as the signal spectrum occupies higher frequencies and is more easily coupled into unintended structures. Therefore, the classical tradeoff between channel capacity and robustness becomes a live design decision in vehicles, necessitating careful application of shielding and coding. The optical communications literature suggests that sophisticated coding can recover performance, but in latency-sensitive ADAS scenarios, not all optical domain strategies translate directly (A. Naughton et al., 2012; Li et al., 2015; Brandonisio et al., 2017).

ECU and IMU integration: The role of ECUs and IMUs detailed in Katari et al. (2024) clarifies the timing constraints inherent in sensor fusion and control. IMUs often provide high-rate inertial data requiring deterministic delivery and minimal jitter to maintain correct state estimation. The meat of the challenge is marrying this deterministic requirement with a high-throughput, potentially bursty 10G fabric. Time-Sensitive Networking (TSN) is the natural candidate to reconcile these demands (Kern, 2013), but TSN itself must be validated under the specific failure modes that automotive environments present (EMI pulses, connector transitions, mechanical vibration).

Multi-protocol gateways: Ioana et al. (2022) demonstrate practical embodiments of gateways in V2X contexts. A key insight is that gateways cannot be viewed as mere protocol adapters; they are active participants in the timing and security posture of the vehicular system. One must ensure gateways do not introduce single-point vulnerabilities; hardware acceleration, secure boot, and authenticated update mechanisms are critical to avoid making gateways a liability. Software-defined gateways that allow dynamic policy changes are attractive, but their complexity must be bounded to avoid runtime unpredictability.

EMC mitigation versus cost and manufacturability: KARIM (2025) shows that validated shielding strategies can solve many EMI issues, but not without cost. Automotive manufacturers operate under strict cost and weight considerations; adding shielding or changing connector families has implications for mass production, weight budgets, and repairability. Therefore, design teams should adopt a holistic cost–benefit framework in which the cost of additional physical mitigation is compared against the operational and safety costs of increased software complexity or the need for more conservative ADAS feature sets.

Cryptography and lifecycle concerns: Post-quantum cryptography introduces a temporal dimension to security engineering. As Prest et al. (2017) describe, algorithms like FALCON provide promising paths for signature schemes resistant to quantum adversaries, but they come with larger key sizes and computational costs. Automotive systems have long lifecycles; thus, early adoption of PQC for long-term artifacts (firmware signing, OTA update chains) is prudent. For real-time vehicle control streams, hybrid schemes or selective application of PQC may be necessary to avoid compromising latency. The architecture we propose supports hybridization and staged adoption, enabling manufacturers to introduce PQC where it is most needed while avoiding undue risk to real-time control operations.

Limitations of the Current Synthesis: The primary limitation of this paper is the absence of novel empirical measurement. Our recommendations and architectural proposals are grounded in the literature provided, but they remain conceptual until validated by targeted experiments in automotive testbeds. For instance, while KARIM (2025) validates shielding approaches using HyperLynx simulations and case studies, actual vehicle-level testing across various electromagnetic environments and across the full vehicle lifecycle (temperature cycling, vibration, wear) is necessary to confirm long-term efficacy. Similarly, the translation of optical network FEC strategies to electrically wired automotive links is promising but untested in vehicle contexts; latency, packetization, and interaction with TSN must be evaluated empirically.

Another limitation is the heterogeneity of the

provided corpus: while it covers a wide array of relevant domains, it does not offer exhaustive coverage of every possible mitigation strategy (for example, specific automotive connector standards or proprietary PHY implementations at 10G). Thus, the architecture should be viewed as a framework to guide design choices, not as a prescriptive standards document.

Future Research Directions: We identify several high-value research trajectories that would operationalize and validate the proposed architecture.

• Empirical EMC and Reliability Testing at Vehicle Scale: Conduct vehicle-level test campaigns that subject proposed shielding and PCB layouts to real-world EMI, vibration, and thermal stress over extended durations. Validate link BER and latency under these conditions.

• Latency-Constrained FEC and Adaptive Error Correction: Develop and evaluate low-latency FEC schemes specifically tuned for automotive control loops, perhaps inspired by burst-mode FEC from optical networks but optimized for short frames and minimal processing overhead.

• Gateway Timing and Semantics Benchmarks: Create standard benchmark suites that stress multi-protocol gateways with realistic ADAS workloads (camera streams, IMU bursts) to quantify translation latency, jitter, and message preservation characteristics.

• PQC Performance in Constrained ECUs: Implement and benchmark candidate PQC algorithms (e.g., FALCON) in representative ECUs and gateway hardware, assessing CPU cycles, memory footprint, and impact on real-time processing budgets. Investigate hardware acceleration avenues for PQC where needed.

• Safety and Security Certification Pathways: Explore how safety certification processes (e.g., ISO 26262) can be adapted to accommodate post-quantum cryptography and 10G AE architectures, including methodologies for demonstrating deterministic behavior in the presence of cryptographic upgrades.

These research directions are consistent with the interdisciplinary nature of the problem and would provide the empirical foundation necessary to move from conceptual design to certified production systems.

**conclusion**

This article synthesizes a diverse literature into an integrated architecture for deploying 10G Automotive Ethernet in ADAS and V2X systems. The synthesis underscores that achieving safe, reliable, and secure operation at 10 Gbps requires co-design across physical, link, gateway, and security layers. Physical design and EMC mitigation are foundational; they reduce the residual bit error environment and enable leaner link-level corrections. Gateways must preserve timing and semantics, and thus require both hardware acceleration and careful buffering strategies. Security needs are evolving toward post-quantum readiness; pragmatic hybrid strategies that prioritize PQC for long-lived artifacts while balancing latency for control streams appear to be the most practical near-term approach.

The conceptual framework and recommendations presented herein should guide engineers and researchers in designing and validating 10G AE systems that satisfy the stringent requirements of ADAS and V2X. However, the proposals must be validated empirically. We recommend coordinated test campaigns, benchmark development, and standardization efforts to ensure that the potential of 10G AE can be realized in safety-critical vehicular environments without compromising the determinism, resilience, and security that ADAS demands.

1. **References**

Ioana, A., Korodi, A., & Silea, I. (2022). Automotive IoT Ethernet-based communication technologies applied in a V2X context via a multi-protocol gateway. Sensors, 22(17), 6382. https://www.mdpi.com/1424-8220/22/17/6382

2. Katari, M., Krishnamoorthy, G., Shanmugam, L., & Tadimarri, A. (2024). Driving towards safety: the role of ecus and imus in advanced driver-assistance systems (ADAS). International Journal for Multidisciplinary Research, 6(2), 1-14. https://www.researchgate.net/profile/GowrisankarKrishnamoorthy/publication/379890228_Driving_Towards_Safety_The_Role_of_ECUs_and_IMUs_in_Advanced_DriverAssistance_Systems_ADAS/links/6628343843f8df018d2551c3/Driving-TowardsSafety-The_Role_of_ECUs_and_IMUs_in_Advanced_Driver-Assistance-SystemsADAS.pdf

3. KARIM, A. S. A. (2025). MITIGATING ELECTROMAGNETIC INTERFERENCE IN 10G AUTOMOTIVE ETHERNET: HYPERLYNX-VALIDATED SHIELDING FOR CAMERA PCB DESIGN IN ADAS LIGHTING CONTROL. International Journal of Applied Mathematics, 38(2s), 1257-1268.

4. Kern, A. (2013). Ethernet and ip for

automotive e/e-architectures-technology analysis, migration concepts and infrastructure (Doctoral dissertation, Erlangen, Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Diss., 2012). https://edok01.tib.uni-hannover.de/edoks/e01dd14/773063501.pdf

5. Naughton, P. Ossieur, C. Antony, D. W. Smith, A. Borghesani, D. G. Moodie, G. Maxwell, P. Healey and P. D. T. OFC/NFOEC, "Error-free 10Gb/s duobinary transmission over 215Km of SSMF using a hybrid photonic integrated reflective modulator.," OSA Technical Digest, 2012.
6. M. Li, B. Li, X. Zhang, Y. Song, Y. Zhang and G. Tu, "Investigation on the performance of 10 Gb/s on uplink space optical communication system based on MSK scheme," in 2015 14th International Conference on Optical Communications and Networks (ICOCN), 2015.
7. M. I. Biswas, G. Parr, S. McClean, P. Morrow and B. Scotney, "A Practical Evaluation in Openstack Live Migration of VMs Using 10Gb/s Interfaces," in 2016 IEEE Symposium on Service-Oriented System Engineering (SOSE), 2016, pp. 346-351.
8. N. Brandonisio, S. Porto, D. Carey, P. Ossieur, G. Talli, N. Parsons and P. Townsend, "Forward error correction analysis for 10Gb/s burst-mode transmission in TDM-DWDM PONs," in 2017 Optical Fiber Communications Conference and Exhibition (OFC), 2017.
9. PQC-Forum. Available online: https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/Mb5ZKpnO57I/m/S8yaURFYCwAJ (accessed on 20 February 2022).
10. Prest, T.; Fouque, P.-A.; Hoffstein, J.; Kirchner, P.; Lyubashevsky, V.; Pornin, T.; Ricosset, T.; Seiler, G.; Whyte, W.; Zhang, Z. FALCON. Post-Quantum Cryptography Project of NIST; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2017.