

Reliability Analysis and Improvement Strategies for IoT-Based Intelligent Connected Vehicle Computer Hardware Architecture

Shang Ma

*Northcote High School, Melbourne, Australia
info@nhs.vic.edu.au*

Abstract. As a pivotal IoT application, the hardware reliability of Intelligent Connected Vehicles (ICVs) is crucial for functional safety and autonomous driving. This paper systematically analyzes reliability challenges in ICV computing systems through theoretical and mechanistic approaches. Key issues identified include CPU/GPU overheating, single points of failure, and network latency, which lead to performance degradation, systemic risks, and decision-making interference. To address these, a multi-dimensional enhancement framework is proposed, integrating hardware optimization, software regulation, architectural redundancy, and edge computing. The study establishes a comprehensive reliability system covering thermal management, redundant backup, and latency compensation to ensure stable ICV operations.

Keywords: Connected Vehicles, hardware architecture, reliability, overheating, network latency

1. Introduction

Intelligent Connected Vehicles (ICVs) are one of the most representative applications of IoT in transportation. Simply put, they are like "supercomputers on wheels." Through high-speed networks, sensors, and their "brains" (control units), the car can perceive the environment in real-time and share information to make driving safer and more efficient. It doesn't just work alone; it talks to the cloud to become part of a giant smart traffic network. This system is powerful, but reliability is key. If the on-board computer overheats, a component fails, or the network lags, the vehicle might malfunction or fail, which poses a significant safety hazard. So, this report will focus on three typical "troubles": CPU/GPU overheating, single points of failure, and network lag. This study will analyze the causes of these issues and propose strategies to enhance the robustness and reliability of the system.

2. Hardware architecture reliability analysis

2.1. CPU/GPU overheating issues

In Intelligent Connected Vehicles (ICVs), high-performance CPUs and GPUs are essential for real-time sensor data processing. However, according to Joule's Law, these high-transistor-density chips generate significant thermal energy during continuous high-load operations.

Several factors exacerbate this heat accumulation. Structurally, computing units are often sealed in protective enclosures for durability and embedded in confined spaces like dashboards, which severely limits air circulation. Environmentally, extreme external conditions, such as direct summer sunlight, further reduce the thermal margin for hardware cooling.

This overheating significantly impairs system reliability by triggering "thermal throttling," a mechanism in which the system automatically reduces its clock speed to prevent damage. Such a reduction in computing power can lead to perception or decision-making lags. In high-speed autonomous driving scenarios, these delays hinder the vehicle's ability to react to its surroundings in real-time, posing severe safety risks.

In the long run, persistent high temperatures accelerate the aging of electronic components [1]. For example, semiconductor materials are more susceptible to "thermal aging" and electromigration in high-temperature environments, which significantly increases their failure rate as the temperature rises. This not only reduces system stability but also shortens the overall lifespan of the vehicle's computing system.

In extreme cases, overheating may even lead to sudden system crashes or automatic reboots. If this occurs while the vehicle is in motion, it could lose computational control within a short period of time. In complex traffic environments, such unforeseen failures would pose severe safety risks.

2.2. Single Point of Failure (SPOF) issues

In the hardware architecture of Intelligent Connected Vehicles (ICVs), the "Single Point of Failure" (SPOF) is another critical hazard threatening system reliability.

Modern smart vehicles are equipped with an array of sensors, including cameras, millimeter-wave radars, and LiDAR, as well as smaller sensors for measuring speed, temperature, and tire pressure. These devices generate massive amounts of data every second. To process this information, engineers have designed a "Centralized Computing Platform" [2]. This design concentrates all data into a single processor for analysis and decision-making, which improves efficiency and simplifies management. However, it also introduces a significant vulnerability: if the entire system over-relies on one core computing unit, the failure of this critical hardware could cause all intelligent functions of the vehicle to fail simultaneously.

As the central hub of the entire system, the reliability of the computing platform is directly linked to the normal operation of all vehicle functions [3]. Any malfunction may interrupt the various functions that rely on it for data processing. In extreme scenarios, the vehicle might lose control over critical functions such as steering, braking, or autonomous driving, leading to severe safety risks.

The occurrence of Single Point of Failure (SPOF) often stems from a design philosophy prioritizing cost-efficiency and simplification over safety redundancy. Without backup systems, a single chip failure—caused by voltage fluctuations or physical damage—can lead to total vehicle loss of control.

As autonomous functions expand, centralized platforms must manage complex tasks like perception and path planning simultaneously. This continuous high-load state, exacerbated by harsh environments (vibration and heat), increases hardware failure probabilities. Crucially, a single component malfunction can trigger a chain reaction, compromising the entire system. The impact of SPOF is catastrophic; for instance, a processor crash during high-speed driving eliminates "buffer time" for driver intervention, posing a severe threat to passenger safety.

2.3. Network latency issues

In the ICV IoT ecosystem, end-to-end latency is a decisive factor for real-time safety. As vehicles transition toward centralized architectures, the internal Ethernet backbone must handle massive data from high-resolution sensors. If bandwidth is insufficient or communication protocols lack efficient Quality of Service (QoS) prioritization, channel congestion will occur, resulting in "stale data". For a vehicle at high speeds, even a 100-millisecond delay in Forward Collision Warning (FCW) or Autonomous Emergency Braking (AEB) can result in a several-meter difference in braking distance, directly threatening passenger safety.

The causes of such latency involve complex internal and external factors. Externally, urban obstacles like skyscrapers and tunnels cause multi-path fading and signal shadowing, degrading 5G or V2X link reliability. Furthermore, transmitting sensor data to remote cloud servers exacerbates latency, highlighting the necessity of Edge-Cloud Synergy to deploy computing power closer to the vehicle terminal.

Ultimately, these latency fluctuations create a "communication bottleneck" that compromises closed-loop control stability. In V2V and V2I coordination, delayed updates on surrounding vehicles or traffic lights can cause path-planning modules to miscalculate safe windows or miss critical braking cues. As autonomous systems become increasingly dependent on high-frequency synchronization, addressing network latency remains a core challenge for enhancing overall hardware and system reliability.

3. Reliability improvement strategies

3.1. Strategies for overheating mitigation

To mitigate the risk of processor overheating, effective cooling strategies for the vehicle's "central brain" must be implemented through physical hardware optimization, software regulation, and task distribution.

At the hardware level, traditional cooling methods must be upgraded to meet the demands of high-performance computing. Relying solely on forced-air cooling (fans) is no longer sufficient for modern ICVs; a superior solution is the integration of a liquid cooling circulation system. By installing sealed coolant channels around the chips, the system can efficiently remove heat via the circulating liquid. This method offers significantly higher thermal dissipation efficiency than air cooling, ensuring that chips remain within safe operating temperatures even under extreme external heat. Furthermore, the application of advanced Thermal Interface Materials (TIMs), such as high-performance thermal grease or graphite sheets, can substantially enhance thermal conductivity between the chip and the heat sink. These materials function as a "thermal highway," preventing localized hotspots and ensuring uniform heat distribution across the computing unit.

In addition to hardware improvements, software-level regulation is essential to balance computational performance with thermal stability. A dynamic resource allocation mechanism can be

introduced to adjust processing power based on real-time driving conditions, utilizing techniques such as multi-domain DVFS to minimize power consumption and heat generation [4]. For instance, on straightforward road segments where data complexity is low, the system can automatically reduce the chip's operating frequency and voltage to minimize heat generation at the source. Conversely, full processing power is only deployed in complex scenarios like navigating busy intersections. This approach is further enhanced by integrating multiple thermal sensors within the chip, coupled with AI algorithms to predict temperature fluctuations. By preemptively increasing the cooling system's intensity before reaching critical temperatures, the system shifts from reactive cooling to proactive prevention.

Finally, Internet of Things (IoT) technology can be utilized to distribute the computational load. Not all data processing must occur on the vehicle's local computer; non-critical tasks, such as high-definition map downloads or media synchronization, can be offloaded to roadside units (RSUs) or cloud servers [5]. This "task offloading" strategy effectively reduces the burden on the vehicle's primary processors. By combining internal power optimization, advanced physical cooling, and cloud-based support, a comprehensive thermal management framework is established. This ensures the long-term stability and reliability of the hardware architecture, preventing system failures caused by thermal stress.

3.2. Strategies for single point of failure prevention

To address the risk of total system collapse caused by a Single Point of Failure (SPOF), the core approach to enhancing hardware reliability is to establish a "redundancy safeguard" mechanism. This ensures that the overall system remains operational even if localized damage occurs.

The most direct and effective solution is the implementation of hardware redundancy. Critical components of Intelligent Connected Vehicles, such as the Central Computing Unit (CCU) and the power supply system, should be equipped with dual, independent hardware sets—commonly referred to as "Dual-CPU" or "Dual-Power" configurations. Under this architecture, the primary and backup processors operate simultaneously, with the backup system maintained in a "Hot Standby" state. Should sensors detect voltage anomalies, data freezes, or physical damage to the primary processor, an automated failover mechanism can transfer control to the backup system within milliseconds. This seamless transition ensures that even if the primary "brain" fails during high-speed travel, the vehicle maintains fundamental braking and steering capabilities, providing the driver with critical intervention time.

Furthermore, to fundamentally mitigate risks, the highly centralized architecture can be optimized toward distributed redundancy. By strategically distributing computational tasks among various Zonal Controllers, the system can decentralize its "authority." For instance, dedicated chips can manage chassis control while others focus exclusively on environmental perception. This distributed design effectively lowers the workload of individual processors, thereby reducing failure rates associated with long-term high-load operation. More importantly, these dispersed modules can perform mutual monitoring via fault-detection algorithms. If the system identifies illogical data from a specific module, it can rapidly "isolate" the anomaly, preventing erroneous signals from spreading across the network uncontrollably.

Finally, establishing self-healing and functional degradation mechanisms is a cornerstone of hardware reliability. In the event of an irreversible failure that affects even the backup systems, the architecture should not simply crash but instead initiate a "Safe Degradation" (or Fail-Soft) mode. In this state, the hardware prioritizes core safety modules by deactivating non-essential functions, such as infotainment or ambient lighting, to conserve remaining power and computing resources. This

"essential-first" strategy allows the vehicle to execute a controlled deceleration and pull over safely using minimal hardware capacity. Through this multi-layered protection—ranging from hardware backups and algorithmic monitoring to emergency degradation—the hardware architecture of ICVs can maintain maximum stability even under extreme conditions.

3.3. Strategies for network latency optimization

To mitigate information lag, the core strategy for enhancing hardware reliability focuses on "minimizing data travel" and "prioritizing critical communication." First, the implementation of Edge Computing is essential. By deploying computational power to intelligent Roadside Units (RSUs), time-critical tasks—such as intersection collision warnings—are processed locally. This reduces transmission distances from kilometers to meters, reducing end-to-end latency from hundreds of milliseconds to just a few milliseconds. This "localized processing" acts as an on-site dispatcher, significantly accelerating emergency response times.

Second, 5G Network Slicing technology can designate "priority lanes" for autonomous driving, ensuring that safety-related perception data maintains highest bandwidth priority regardless of routine infotainment traffic [6]. This stabilizes latency fluctuations even during severe congestion. Complementing this, hardware should support multi-modal communication fusion, integrating 5G, DSRC, and satellite links. This allows the system to switch to interference-resistant backup links in signal-obscured areas such as tunnels, ensuring continuous data flow.

Finally, latency-aware and compensation algorithms must be established at the hardware base layer. When excessive latency is detected, the onboard computer initiates a "Prediction Mode." By synthesizing historical sensor trajectories with inertial navigation data, the system predicts the "shadow positions" of surrounding vehicles [7]. This algorithmic "filling" of hardware delays maintains closed-loop control stability during temporary network jitter. Through this three-pronged strategy—edge offloading, network slicing prioritization, and on-board latency compensation—ICVs achieve the real-time synchronization necessary for safe operation in complex IoT environments. Four

4. Conclusion

In summary, the reliability of Intelligent Connected Vehicles (ICVs) serves as the cornerstone for achieving high-level autonomous driving. As the physical carrier of computational power, the stability of the hardware architecture constitutes the vehicle's safety lifeline. Through in-depth analysis of processor overheating, Single Points of Failure (SPOF), and network latency, this paper concludes that hardware failures are rarely isolated; rather, they result from the synergy of extreme environments, high-load computation, and complex network fluctuations.

To address these challenges, the strategies proposed herein establish a "multi-dimensional defense system": ranging from physical-layer liquid cooling and software-driven dynamic frequency scaling to hardware redundancy via "hot standby" and distributed self-healing architectures. Furthermore, communication optimization through edge computing and network slicing not only alleviates pressure on individual nodes but also leverages IoT technology to achieve risk distribution and collaborative protection across vehicles, roads, and the cloud.

Looking ahead, with the evolution of semiconductor processes and the advent of 6G communication, ICV hardware architectures will become increasingly intelligent and flexible. Reliability analysis will shift from "reactive remediation" to "proactive prediction." Supported by precise AI forecasting algorithms and higher-bandwidth IoT frameworks, future intelligent vehicles

will possess superior fault tolerance and self-healing capabilities, ultimately providing the public with a safer, smoother, and more efficient mobility experience.

This study faces several limitations. First, due to restricted conditions, the findings rely on theoretical analysis rather than empirical simulation or laboratory testing, requiring further engineering validation. Second, the cost-benefit trade-offs of implementing high-performance cooling and redundancy systems have not been fully analyzed. Finally, the model simplifies the dynamic complexity of real-world IoT environments. Future research will utilize digital modeling and simulation tools to explore more robust hardware optimization strategies under diverse electromagnetic and high-concurrency communication scenarios.

References

- [1] Kuznetsov, G. V., Kravchenko, E. V., & Pribaturin, N. A. (2016). Reliability analysis of electrical engineering power semiconductor devices. *Russian Electrical Engineering*, 87(4), 235-239.
- [2] Bandur, V., Selim, G., Pantelic, V., & Lawford, M. (2021). Making the case for centralized automotive E/E architectures. *IEEE Transactions on Vehicular Technology*, 70(2), 1230-1245.
- [3] Li, Y., Liu, W., Liu, Q., Zheng, X., Sun, K., & Huang, C. (2024). Complying with iso 26262 and iso/sae 21434: A safety and security co-analysis method for intelligent connected vehicle. *Sensors*, 24(6), 1848.
- [4] Haj-Yahya, J., Alser, M., Kim, J., Yağlıkçı, A. G., Vijaykumar, N., Rotem, E., & Mutlu, O. (2020, May). SysScale: Exploiting multi-domain dynamic voltage and frequency scaling for energy efficient mobile processors. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)* (pp. 227-240). IEEE.
- [5] Bute, M. S., Fan, P., Liu, G., Abbas, F., & Ding, Z. (2022). A cluster-based cooperative computation offloading scheme for C-V2X networks. *Ad Hoc Networks*, 132, 102862.
- [6] Khan, H., Luoto, P., Bennis, M., & Latva-aho, M. (2018, May). On the application of network slicing for 5G-V2X. In *European Wireless 2018; 24th European Wireless Conference* (pp. 1-6). VDE.
- [7] Zhao, Y., Zhang, X., Mihalj, T., Schabauer, M., Putzer, L., Reichmann-Blaga, E., ... & Eichberger, A. (2025). A Communication-Latency-Aware Co-Simulation Platform for Safety and Comfort Evaluation of Cloud-Controlled ICVs. *IEEE Internet of Things Journal*.