PERSPECTIVE

From Fog-Enabled IoT to Cognitive Internet of Vehicles: A Perspective on Mutable–Immutable Blockchain Architectures

Saeed Javanmardi[1] **and Marco Scarpa**[1,*]

[1] Department of Engineering, University of Messina, 98166 Messina, Italy

Abstract

Fog-enabled Internet of Things (IoT-Fog) architectures and the Cognitive Internet of Vehicles (Cognitive IOV) are becoming key enablers for next-generation intelligent transport systems. At the same time, Blockchain is used to support integrity, transparency, and trust in distributed vehicular services. In this perspective article, we link these domains by discussing a hybrid Mutable-Immutable Blockchain Architecture in which mutable ledgers at the fog layer are combined with an immutable ledger in the cloud. We first outline an analytical model for Cognitive IOV over IoT-Fog, and then extend it to include Blockchain-related delays, throughput, and tamper resistance in the hybrid architecture. The main contribution is an analytical view of how key parameters, such as the fraction of transactions sent to the Blockchain or anchored in the immutable layer, influence the trade-off between latency, flexibility, and security in Cognitive IOV services.

Keywords: blockchain, IoT-Fog networks, cognitive internet of vehicles (IoV), mutability, latency-aware architecture.

Submitted: 22 December 2025
Accepted: 17 March 2026
Published: 25 March 2026

Vol. 1, **No.** 1, 2026.
10.62762/JSS.2025.861548

***Corresponding author:**
✉ Marco Scarpa
mscarpa@unime.it

1 Introduction

Fog-enabled IoT–Fog architectures use fog nodes between devices and the cloud to give local computing, storage, and control, which helps to reduce delay and network traffic. In transport systems, this change helps us to move from the normal Internet of Vehicles to the Cognitive Internet of Vehicles (CIoV). In CIoV, vehicles and roadside units use smart software to read data and improve the safety and quality of services. Blockchain helps these systems by keeping data honest and clear. However, a fully immutable Blockchain is too hard to change, and a fully mutable Blockchain can reduce trust. A hybrid Mutable–Immutable Blockchain Architecture [1] uses mutable ledgers at the fog layer and an immutable ledger in the cloud to keep a good balance between flexibility, security, and performance in Cognitive IOV over IoT–Fog.

The integration of Cognitive IoV with hybrid Mutable–Immutable Blockchain Architectures is still not well understood, especially because CIoV has stricter requirements than typical IoT systems, such as high vehicle speed, strong delay sensitivity, variable

Citation
Javanmardi, S., & Scarpa, M. (2026). From Fog-Enabled IoT to Cognitive Internet of Vehicles: A Perspective on Mutable–Immutable Blockchain Architectures. *Journal of Systems Scalability*, 1(1), 29–34.

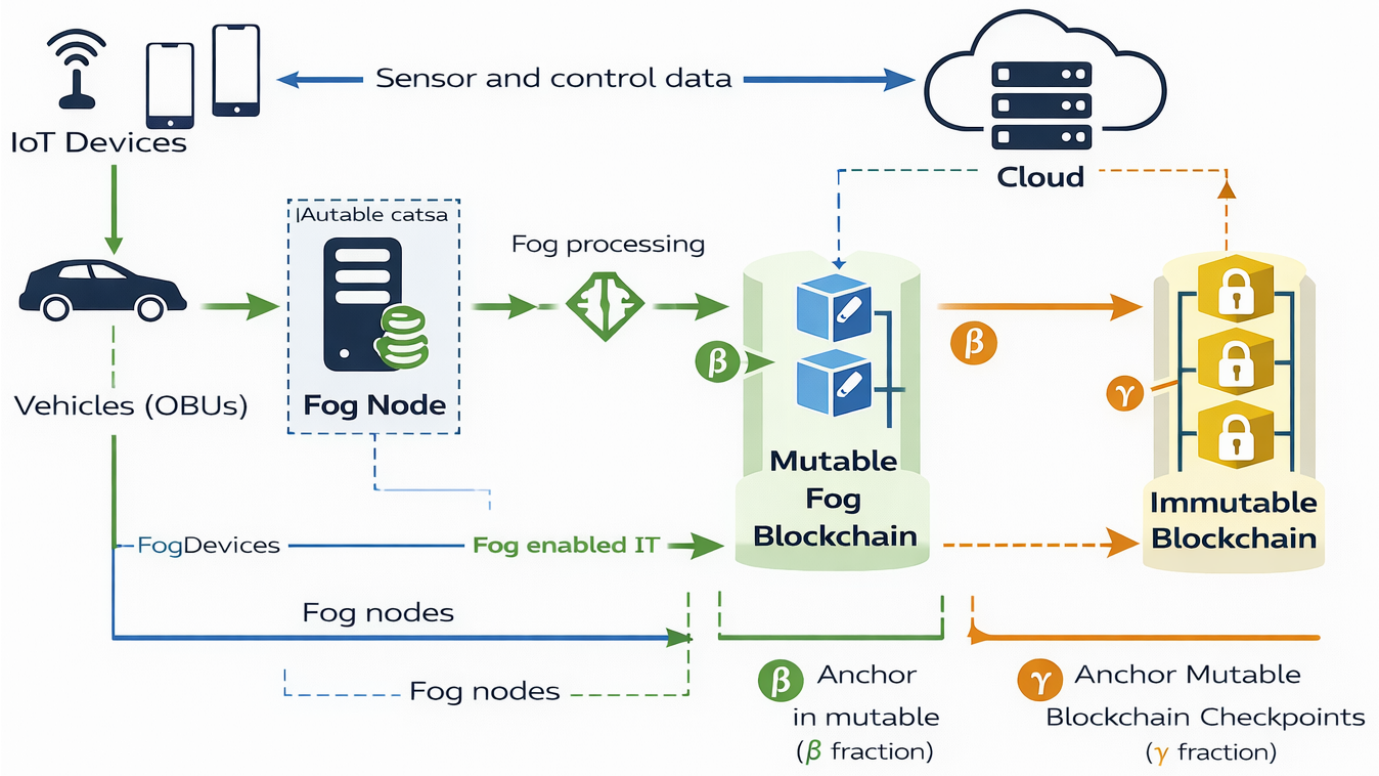

Figure 1. The reference architecture.

data rates, and tight privacy and safety constraints. Using only immutable Blockchains can cause high delay and very rigid data management, and using only mutable Blockchains can make the system easier to change and less trusted. This perspective article tries to fill this gap by giving one simple analytical view that connects Fog-enabled IoT, Cognitive IoV, and hybrid Mutable–Immutable Blockchain Architectures. We first present a Cognitive IoV model that describes the roles of devices, fog nodes, cloud, and cognitive engines, and then extend analytical models of mutable and immutable Blockchains to the Cognitive IoV context. The main contribution is an analytical discussion of how such hybrid architectures affect latency, throughput, and security, and how these insights can guide future research on secure and efficient vehicular services [2].

Figure 1 shows the architecture of the proposed fog-enabled hybrid Blockchain paradigm. IoT devices and vehicles generate data and send it to nearby fog nodes. The fog node performs local processing and records transactions in the mutable fog Blockchain. This mutable Blockchain provides low latency and allows fast updates close to the edge. A fraction of the data or checkpoints is then anchored to the immutable Blockchain. The immutable Blockchain provides strong security and tamper resistance. The cloud supports global coordination and storage but is not involved in time-critical processing. This design combines low delay at the fog layer with high security at the global layer.

2 Cognitive IOV: Analytical Modelling

In this section, we give an analytical view of a Cognitive Internet of Vehicles (Cognitive IOV) system built on top of Fog-enabled IoT [3]. A typical Cognitive IOV scenario has vehicles with on-board units (OBUs), roadside units (RSUs), fog nodes at the edge, and a cloud back-end. Cognitive engines can run at the vehicle, fog, or cloud level and use data from sensors, communication links, and past records to make decisions in real time.

To describe how Cognitive IOV behaves over time, we look at how messages move between vehicles and the fog layer. Let λ_v be the average rate (messages per second) at which one vehicle sends data (for example, position, speed, or alarms) to a nearby fog node. In a coverage area with N_v active vehicles, the total message arrival rate at the fog node can be written as

$$\Lambda_f = N_v \lambda_v. \tag{1}$$

The fog node handles these messages using local decision rules and cognitive models, for example to predict collisions or traffic jams. If we see the fog node

as a service system with service rate μ_f (messages per second), a simple M/M/1 model gives the average processing delay at the fog as

$$D_f = \frac{1}{\mu_f - \Lambda_f}, \tag{2}$$

with the stability condition $\Lambda_f < \mu_f$. This delay includes both the waiting time in the queue and the processing time at the fog node.

The Cognitive IOV system also needs to talk to the cloud, for example to update global models or to store selected events for later analysis. Let λ_c be the rate at which processed fog events are sent to the cloud (for instance, only a fraction α of messages that meet some relevance rule). In that case, $\lambda_c = \alpha\Lambda_f$, with $0 \le \alpha \le 1$. The cloud has a higher service capacity μ_c, but the end-to-end delay for these messages also includes the network delay between the fog and the cloud, denoted as D_{net}. If we again model the cloud as an M/M/1 queue, the average delay for cloud operations is

$$D_c = D_{net} + \frac{1}{\mu_c - \lambda_c}. \tag{3}$$

Even if D_c can be quite large, it is often acceptable, because most cloud tasks in Cognitive IOV are not strict real-time tasks but long-term learning and optimisation.

The total response time of a Cognitive IOV service, as seen by a vehicle, depends on how many parts are in the control loop. For local decisions that stay in the fog layer, the average decision time can be written as

$$T_{\text{local}} \approx D_{\text{tx}} + D_f + D_{\text{rx}}, \tag{4}$$

where D_{tx} and D_{rx} are the uplink and downlink communication delays between the vehicle and the fog node (including wireless access and RSU forwarding). If the decision also needs the cloud (for example, when a new model is requested), we can write

$$T_{\text{hybrid}} \approx D_{\text{tx}} + D_f + D_c + D_{\text{rx}}. \tag{5}$$

These simple formulas show that Cognitive IOV performance is very sensitive to the balance between local fog processing and remote cloud operations. In practice, Cognitive IOV systems try to keep safety-critical and low-latency tasks inside T_{local}, and use T_{hybrid} only for less urgent functions.

Finally, Cognitive IOV is not only about delay and throughput, but also about how information is chosen, combined, and used by cognitive engines. We can describe the cognitive workload at the fog node as a function of both the message rate and the complexity of the decision models. Let C_f be the average computational cost (for example, CPU cycles or processing time) per message for the local cognitive engine. The effective service rate μ_f then depends on the available resources and on C_f. This means that changes in model complexity (for example, moving from a simple rule-based algorithm to a more complex machine learning model) directly affect D_f and T_{local}. These analytical relations are important later when we add Blockchain to Cognitive IOV, because Blockchain operations will bring extra computational and communication costs on top of the existing cognitive and networking delays.

3 Cognitive IOV in Mutable–Immutable Blockchain Architectures: Analytical Modelling

In this section, we extend the previous Cognitive IOV model by adding a hybrid Mutable–Immutable Blockchain Architecture. We assume that each fog domain has a local mutable ledger, and that a global immutable ledger runs in the cloud. Vehicles send messages to the fog layer as described before, and some of these messages must be stored on the Blockchain (for example, safety events, service payments, or trust updates).

Let Λ_f be the total arrival rate of messages at a fog node, as defined in the previous section. We assume that a fraction β of these messages need a Blockchain transaction. The arrival rate of Blockchain transactions at the fog mutable ledger is then

$$\lambda_b = \beta\Lambda_f, \quad 0 \le \beta \le 1. \tag{6}$$

The fog node (or a small group of nodes) runs a consensus mechanism [4] to group transactions into blocks and add them to the mutable chain. For simplicity, we model the fog Blockchain as a service system with effective service rate μ_m (transactions per second). This rate depends on block size, block time, and the cost of verification and consensus. Using an M/M/1 model, the average waiting and processing delay on the mutable ledger is

$$D_m = \frac{1}{\mu_m - \lambda_b}, \tag{7}$$

with stability condition $\lambda_b < \mu_m$. We can also include a local propagation delay among fog nodes, D^m_{prop}, so the total ledger delay at the fog side is $D^{tot}_m \approx D_m + D^m_{prop}$.

The mutable ledger also supports controlled redaction, which is important to fix errors or to follow legal rules. Let λ_r be the arrival rate of redaction requests, and assume that $\lambda_r = \gamma\lambda_b$ with $0 \le \gamma \le 1$, meaning that only a fraction of transactions will later need a redaction. Each redaction needs extra checks and an update of the local ledger state, which creates extra load. If we call μ_r the service rate for redaction requests, we can write the redaction delay as

$$D_r = \frac{1}{\mu_r - \lambda_r}, \tag{8}$$

and the total processing capacity of the fog ledger now depends on both normal transactions and redactions. In practice, if γ is high, the available capacity for new Cognitive IOV transactions goes down and the average delay D_m^{tot} goes up, so system designers should keep γ small with good data and policy management.

The immutable cloud ledger is used to anchor or aggregate information from many fog domains. We assume that only a fraction δ of fog Blockchain transactions are chosen for anchoring in the cloud, based on their importance or legal needs. The arrival rate of anchoring transactions at the immutable ledger is

$$\lambda_i = \delta\lambda_b, \quad 0 \le \delta \le 1. \tag{9}$$

These transactions can be sent one by one or in batches of size B, which are committed to the cloud chain every T_a seconds. If we again see the immutable ledger as a service system with rate μ_i (anchoring operations per second), the average ledger delay at the cloud side can be written as

$$D_i = D_{net} + \frac{1}{\mu_i - \lambda_i}, \tag{10}$$

where D_{net} is the network delay between fog and cloud. When batching is used, the effective anchoring delay also includes a waiting term of about $T_a/2$ before a batch is full, so a more complete expression is

$$D_i^{tot} \approx D_{net} + \frac{T_a}{2} + \frac{1}{\mu_i - \lambda_i}. \tag{11}$$

For many Cognitive IOV services, this extra delay is acceptable, because anchoring is not part of the most urgent safety loop, but it is important for long-term trust and auditing.

We can now combine the previous formulas to estimate the end-to-end response time of a Cognitive IOV service that uses the mutable ledger. For a vehicle request that must be written to the fog Blockchain but does not need immediate cloud anchoring, the average decision time is

$$T_{\text{BC-local}} \approx D_{\text{tx}} + D_f + D_m^{tot} + D_{\text{rx}}, \tag{12}$$

where D_{tx} and D_{rx} are the wireless uplink and downlink delays, and D_f is the basic fog processing delay from the previous section (without Blockchain). If the service also needs confirmation that the transaction is anchored in the immutable ledger, the response time becomes

$$T_{\text{BC-hybrid}} \approx D_{\text{tx}} + D_f + D_m^{tot} + D_i^{tot} + D_{\text{rx}}. \tag{13}$$

These formulas show that Cognitive IOV designers must carefully choose which operations stay only in the mutable fog ledger and which ones must also be confirmed by the immutable cloud ledger, in order to keep delay inside the limits of each application.

Finally, we give a simple view of tamper resistance in this hybrid architecture. Let P_m be the probability that an attacker can successfully make an undetected change on the mutable ledger in a given time window (for example, by controlling some fog nodes or keys). Let P_i be the probability of a successful attack on the immutable ledger in the same window, which is usually much lower because of stronger consensus and higher decentralisation. If a fraction $(1-\delta)$ of Cognitive IOV decisions rely only on the mutable ledger, while a fraction δ are also anchored in the immutable ledger, we can write an effective tamper resistance metric as

$$P_{\text{eff}} \approx (1-\delta)P_m + \delta P_i. \tag{14}$$

When we increase δ, the system becomes closer to the security level of the immutable ledger, but the delay and the use of resources also increase. When we decrease δ and keep more operations only on the mutable layer, the system becomes more flexible and faster, but the chance of a successful undetected attack also becomes higher. This simple analytical view helps to understand the trade-offs that Cognitive IOV architects face when they design Mutable-Immutable Blockchain Architectures and motivates further work on adaptive strategies that tune parameters such as β, γ, and δ according to service type, risk level, and current network conditions.

4 Research Summary

This article gives an analytical view of Cognitive Internet of Vehicles (Cognitive IoV) running on Fog-enabled IoT networks and adds a hybrid Mutable–Immutable Blockchain Architecture to it.

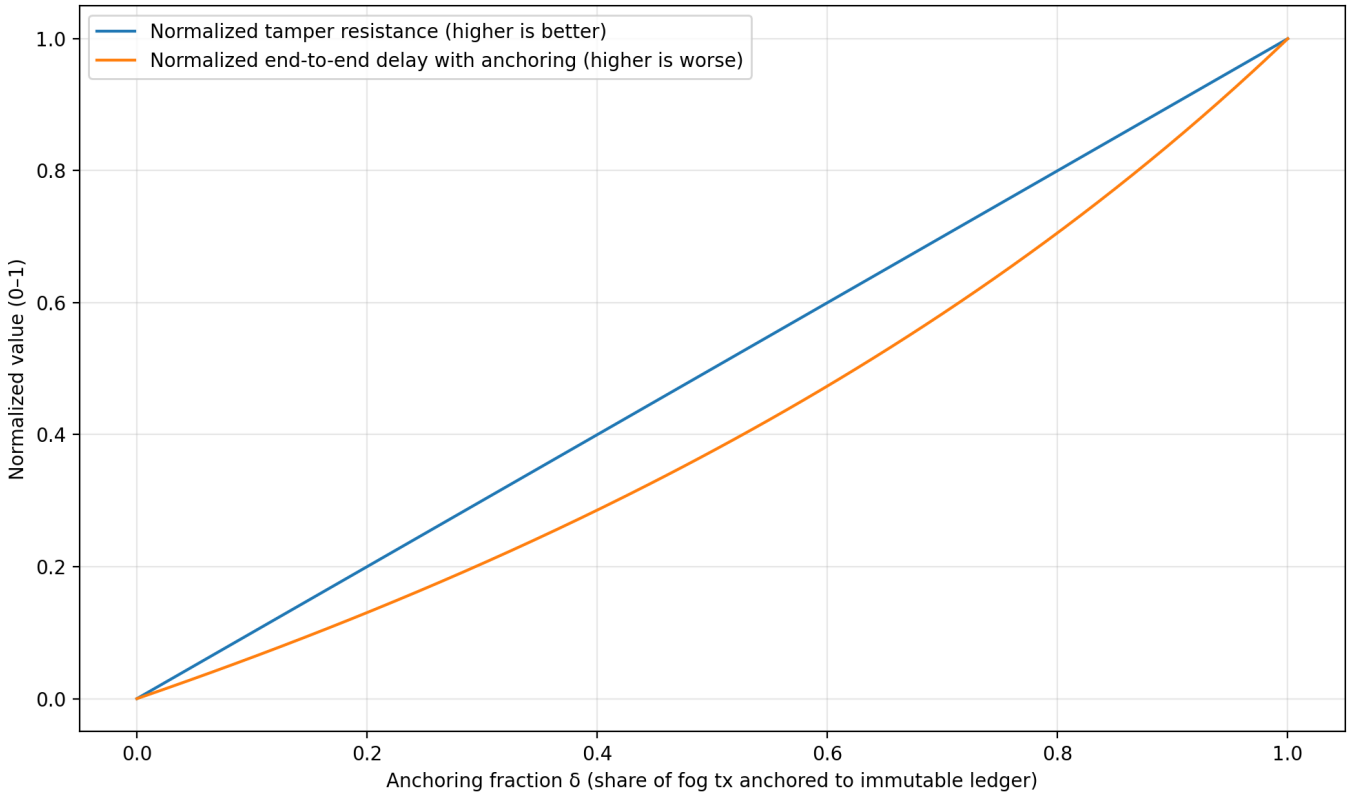

Figure 2. The trade-off between security and delay.

First, the paper models Cognitive IoV by defining message arrival rates from vehicles to fog nodes, fog service rate, and cloud communication delay. The authors use simple queueing models (M/M/1) to estimate average processing delay at fog nodes and cloud servers. They show how factors like the number of active vehicles and the cognitive workload change the delay and throughput of IoV services.

Then, the article extends this model to include Blockchain operations. It calculates transaction arrival rates for the mutable ledger at the fog, redaction request rates, and anchoring rates for the immutable ledger in the cloud. It also shows how block size, consensus, and batch anchoring affect service rates and total Blockchain delay. The paper provides formulas for local Blockchain delay and hybrid Blockchain delay, and it explains how parameters such as β (transaction fraction), γ (redaction fraction), and δ (anchoring fraction) influence performance.

The findings show that a higher value of δ increases security and makes the system closer to immutable reliability, but also increases delay. A lower value of δ keeps services fast but reduces tamper resistance. The paper also introduces a simple metric for effective tamper resistance and shows how mutable and immutable layers balance flexibility, latency, and protection. Overall, the article concludes that hybrid Blockchain can give better flexibility than fully immutable systems, and better trust than fully mutable ones, while still supporting delay-sensitive Cognitive IoV services.

Figure 2 shows the trade-off between **security** and **delay** when we change the **anchoring fraction** δ. When $\delta = 0$, fog transactions are not anchored to the immutable ledger. In this case, the system is faster, but the security is closer to the mutable ledger. So, the **tamper resistance** is lower. When δ increases, more fog transactions are anchored to the immutable ledger. This improves security. So, the **tamper resistance** increases as δ increases. However, anchoring adds extra steps, such as extra communication and processing on the cloud side. It can also create more waiting time in the queue. Because of this, the **end-to-end delay** increases when δ increases.

In summary, δ is a design choice:

- A small δ gives **lower delay**, but **weaker security**.
- A large δ gives **stronger security**, but **higher delay**.

5 Conclusion

In this perspective article, we connected Fog-enabled IoT, Cognitive Internet of Vehicles, and a hybrid Mutable-Immutable Blockchain Architecture through an analytical view. We first described how Cognitive IOV services can be modelled in terms of message rates, processing delays, and cognitive workloads at the vehicle, fog, and cloud layers. Then, we extended this model to include a mutable ledger at the fog side and an immutable ledger in the cloud, introducing simple expressions for Blockchain-related delay, throughput, and tamper resistance.Our analysis shows that design choices, like how many messages go to the Blockchain, how often data is changed (redacted), and how many transactions are saved in the immutable ledger, directly affect the balance between delay, flexibility, and security.

Data Availability Statement

Not applicable.

Funding

This work was partially funded by European Union - Next generation EU - PNRR - Missione 4, Componente 2, Investimento 1.1 - Bando PRIN 2022 PNRR - Decreto Direttoriale n. 1409 del 14-09-2022 - Progetto PaB-PIF, CUP J53D2301493, Project ID. P20227W8ZC.

Conflicts of Interest

The authors declare no conflicts of interest.

AI Use Statement

The authors declare that no generative AI was used in the preparation of this manuscript.

Ethical Approval and Consent to Participate

Not applicable.

References

[1] Javanmardi, S., Scarpa, M., Shojafar, M., Distefano, S., & Merlino, G. (2025). Mutable Blockchains in IoT-Driven Sustainable Urban Planning: Challenges, and Analytical Modeling. In *2025 IEEE International Conference on Smart Computing* (*SMARTCOMP*) (pp. 408–413). IEEE. [CrossRef]

[2] Chen, M., Tian, Y., Fortino, G., Zhang, J., & Humar, I. (2018). Cognitive internet of vehicles. *Computer Communications, 120*, 58–70. [CrossRef]

[3] Reshi, I. A., & Sholla, S. (2024). Securing IoT data: Fog computing, blockchain, and tailored privacy-enhancing technologies in action. *Peer-to-Peer Networking and Applications, 17*(6), 3905–3933. [CrossRef]

[4] Reshi, I. A., & Sholla, S. (2025). IBF network: enhancing network privacy with IoT, blockchain, and fog computing on different consensus mechanisms. *Cluster Computing, 28*(3), 208. [CrossRef]

Saeed Javanmardi is a postdoctoral researcher at the University of Messina in Italy. He received his PhD in Information Technology from the University of Naples Federico II. His research focuses on IoT-Fog networks, Blockchain security, and resource management. He has worked on intrusion detection systems and task scheduling, especially in drone and fog environments. During his academic career, he has published several journal and conference papers and collaborated with research groups and industry partners in Italy. His goal is to design secure and efficient network solutions for modern IoT applications. (Email: saeed.javanmardi@unime.it)

Marco Scarpa is a full professor at the Department of Engineering, University of Messina, Italy. He was born on 4 March 1969 in Italy and holds a PhD in Computer Science. His research interests include distributed computing, computer reliability, and wireless sensor networks. He teaches courses in computer systems and basic computer science, and he has contributed to many research projects and scientific publications in computing and engineering. (Email: marco.scarpa@unime.it)