

Optimizing Validator Energy in PureChain-Based Federated Learning for Software-Defined Vehicles

Miraculous Udurume¹, Love Allen Chijioke Ahakonye², Jae Min Lee¹, Dong-Seong Kim^{1*}

¹ IT-Convergence Engineering, Kumoh National Institute of Technology, Gumi, South Korea

² ICT Convergence Research Center, Kumoh National Institute of Technology, Gumi, South Korea

*NSLab Co. Ltd., Gumi, South Korea, Kumoh National Institute of Technology, Gumi, South Korea

udurumemiraculous@gmail.com, (loveahakonye, ljmpaul, dskim)@kumoh.ac.kr

Abstract—Safety-critical software-defined vehicle (SDV) functions require on-vehicle learning under tight power budgets. Blockchain-federated learning (FL) designs, however, run the proof-of-authority and association (PoA²) consensus per client update, which inflates validator energy. We present PureChain, a co-design that seals *one* PureChain block *per* FL round after aggregation, eliminating consensus per update. The mechanism reduces the validator energy and consistently improves the energy–delay products (EDP, ED²P) without degrading the detection accuracy.

Index Terms—Co-scheduling, Efficiency, Energy–Delay, Federated Learning, PoA², PureChain, Software-Defined Vehicles (SDV)

I. INTRODUCTION

Edge autonomy in Software-Defined Vehicles (SDVs) relies on local intelligence for safety-critical V2V/V2I functions, such as collision prevention and cooperative maneuvers, with strict energy and latency constraints [1]. Centralized authentication or anomaly monitoring is fragile in dynamic environments and lacks scalability [2]. A more robust approach is edge learning with trust anchored at the edge, where federated learning (FL) keeps data local and blockchain ensures tamper-proof coordination [3], [4]. Homomorphic encryption further enables privacy-preserving misbehavior detection in IoV, reinforcing the need for edge FL [5]. Concurrently, FL research is focusing on energy efficiency, incorporating client selection and system-level co-design to meet edge power limitations [6], [7].

A key inefficiency remains: many blockchain-FL stacks run consensus per client update, which inflates validator energy and round latency, the scarcest resources in vehicular edge networks [2]. We propose an energy-aware FL-IDS for SDVs that co-schedules PureChain (PoA²) consensus, finalizing a single block per federated round after aggregating the data from the roadside unit (RSU), rather than n blocks for n clients. This aligns the consensus cost with the progress of round-level FL, preserving privacy while improving traceability and efficiency for timely intrusion detection [8], [9].

II. SYSTEM ARCHITECTURE

We propose a four-layer energy-aware SDV stack: (i) on-vehicle SDV layer (sensing, IDS, local training), (ii) RSU layer with an energy-aware client scheduler and FedAvg aggregation with round co-scheduling, (iii) PureChain layer with PoA²

validators, and (iv) audit/policy (access control, log), see Fig. 1.

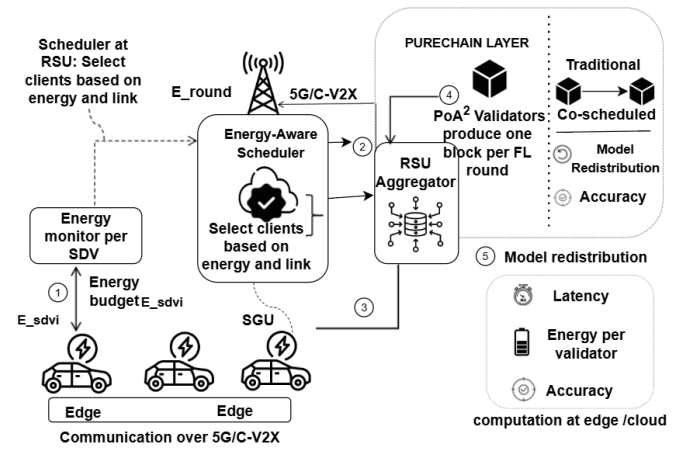


Fig. 1. Energy-aware PureChain workflow. (1) SDVs train locally and send updates; (2) RSU *energy-aware scheduler* selects clients (e.g., residual energy/link QoS); (3) RSU aggregates (FedAvg) and prepares a round payload; (4) PureChain (PoA²) finalizes *one* block *per* round after aggregation (co-scheduled consensus); (5) reported metrics: latency, validator energy, and accuracy.

Each vehicle V_i extracts features $\phi(D_i)$, runs a lightweight IDS, and performs a local secure gradient update (SGD) step to produce the client update $u_i^{(r)}$ as in Equation 1:

$$\theta_i^{(r)} = \theta^{(r)} - \eta \nabla_{\theta} \mathcal{L}(\theta^{(r)}; D_i), \quad u_i^{(r)} = \theta_i^{(r)} - \theta^{(r)}. \quad (1)$$

The RSU scheduler chooses a subset $\mathcal{S}_r \subseteq \{1, \dots, N\}$ for round r (energy/link awareness) and aggregates only $\{u_i^{(r)} : i \in \mathcal{S}_r\}$ through the FedAvg as in Equation 2; it forms a compact payload $p_r = (r, H(\theta^{(r+1)}), \text{stats})$:

$$\theta^{(r+1)} = \sum_{i=1}^N w_i (\theta^{(r)} + u_i^{(r)}), \quad w_i = \frac{|D_i|}{\sum_{j=1}^N |D_j|}. \quad (2)$$

Rather than validating a block per client update, PoA² validators finalize *one* block B_r per round when a ρ majority signs p_r , amortizing consensus from per update to per round as *quantified* in Equation 3:

$$t_{\text{cons}}^{\text{fixed}} = N t \quad \longrightarrow \quad t_{\text{cons}}^{\text{co}} = t. \quad (3)$$

Vehicles authenticate to RSUs; PoA² enforces RSU–validator bindings. Raw data never leaves V_i ; PureChain records only hashes and minimal metadata for accountability and forensics.

III. EXPERIMENTATION AND PERFORMANCE ANALYSIS

We evaluated on CICIoV2024 with client counts $n_c \in \{5, 10, 20\}$ on 10 federated rounds. Data are non-independent and identically distributed (non-IID) (Dirichlet $\alpha=0.7$). The learning stack (features, labels, model, optimizer) is identical across both modes; only consensus timing changes: per update PoA² (one block per client update) vs. round-batched co-scheduling (one block per round after RSU aggregation). Validators are set to $n_v=5$, $P_v=40$ W. Because the learning pipeline remains unchanged, accuracy/F1 score remains stable; we therefore focus on latency/throughput, total energy, and energy–delay (EDP/ED²P). Figure 2 shows the total energy in $R=10$ rounds. Co-scheduling reduces energy by -176 J, -376 J, and -776 J for $n_c=5, 10, 20$, respectively, reflecting the change from (n_c) to (1) consensus commits per round.




Fig. 2. Total energy over $R=10$: co-scheduled PureChain (1 block/round) vs. per-update PoA².

Figure 3 reports updates/s vs. clients at $R=10$. Co-scheduling improves rounds/s by $+0.079$, $+0.234$, $+0.959$ (updates/s by $+0.395$, $+2.339$, $+19.183$) for $n_c=5, 10, 20$, respectively, with greater gains at higher client density.




Fig. 3. Updates/s at $R=10$: co-scheduled PureChain increases throughput by removing per-update consensus.

In all settings, EDP and ED²P decrease as energy increases, while throughput increases, indicating simultaneous gains in efficiency and responsiveness without sacrificing detection quality. FL retains raw data on-vehicle, and the blockchain

ledger provides auditability and traceability, properties desirable for SDV deployments under tight power and latency budgets.

IV. CONCLUSION AND FUTURE WORK

We presented a PureChain-based energy-efficient FL co-design that locks a single PoA² block per FL round. Experiments on CICIoV2024 with 5, 10, and 20 clients over 10 rounds preserved detection accuracy while reducing validator energy and improving energy efficiency. Secondary benefits include lower round latency and higher updates per second, with gains strengthening as the fleet size increases. Future work will evaluate hardware-in-the-loop on vehicular edge platforms, co-optimizing adaptive round-level batching with compression/quantization, broadening robustness metrics, incorporating reputation-aware validator selection, and extending to additional datasets and road scenarios.

ACKNOWLEDGMENT

This work was partly supported by the Innovative Human Resource Development for Local Intellectualization program through the IITP grant funded by the Korea government (MSIT) (IITP-2025-RS-2020-II201612, 33%) and by the Priority Research Centers Program through the NRF funded by the MEST (2018R1A6A1A03024003, 33%) and by the MSIT, Korea, under the ITRC support program (IITP-2025-RS-2024-00438430, 34%).

REFERENCES

- [1] X. Zhu, Z. Li, Y. Jiang, J. Xu, J. Wang, and X. Bai, “Real-Time Vehicle-to-Vehicle Communication-Based Network Cooperative Control System Through Distributed Database and Multimodal Perception: Demonstrated in Crossroads,” in *International Congress on Information and Communication Technology*. Springer, 2024, pp. 133–146.
- [2] L. He, F. Li, H. Xu, W. Xia, X. Zhang, and X. Tao, “Blockchain-Based Vehicular Edge Computing Networks: The Communication Perspective,” *Science China Information Sciences*, vol. 66, no. 7, p. 172301, 2023.
- [3] P. Surapaneni, S. Bojjagani, V. C. Bharathi, M. Kumar Morampudi, A. Kumar Maurya, and M. Khurram Khan, “A Systematic Review on Blockchain-Enabled Internet of Vehicles (BIOV): Challenges, Defenses, and Future Research Directions,” *IEEE Access*, vol. 12, pp. 123 529–123 560, 2024.
- [4] L. A. C. Ahakonye, C. I. Nwakanma, and D.-S. Kim, “Tides of Blockchain in IoT Cybersecurity,” *Sensors*, vol. 24, no. 10, p. 3111, 2024.
- [5] H. L. Nakayiza, L. A. C. Ahakonye, D.-S. Kim, and J. M. Lee, “Homomorphic encryption for privacy-preserving misbehavior detection in the internet of vehicles,” in *2025 International Conference on Artificial Intelligence in Information and Communication (ICAIC)*, 2025, pp. 0320–0324.
- [6] F. Maciel, A. M. de Souza, L. F. Bittencourt, L. A. Villas, and T. Braun, “Federated learning energy saving through client selection,” *Pervasive and Mobile Computing*, p. 101948, 2024, early Access.
- [7] X.-T. Dang, B.-M. Vu, Q.-S. Nguyen, T.-T.-M. Tran, J.-S. Eom, and O.-S. Shin, “A survey on energy-efficient design for federated learning over wireless networks,” *Energies*, vol. 17, no. 24, p. 6485, 2024.
- [8] L. A. C. Ahakonye, C. I. Nwakanma, J. M. Lee, and D. S. Kim, “Purechain-Enhanced Federated Learning for Dynamic Fault Tolerance and Attack Detection in Distributed Systems,” *High-Confidence Computing*, p. 100354, 2025.
- [9] D.-S. Kim, I. S. Igboanusi, L. A. C. Ahakonye, and G. O. Anyanwu, “Proof-of-Authority-and-Association Consensus Algorithm for IoT Blockchain Networks,” in *The 43rd IEEE International Conference on Consumer Electronics (ICCE 2025)*, 2025.