Sivaprakash Nithyanandam[1], Pradeep Anjuru[2]
[1]Electrify America, USA
[2]Akkodis, USA

Received: 28-07-2025; Revised: 04-09-2025; Accepted: 19-10-2025

Edge-Cloud Collaborative Architectures for Real-Time EV Charging Control

Abstract
EV charging control is increasingly operating under distributed and time-sensitive conditions in which chargers, local controllers, and cloud services must coordinate continuously. Existing studies have examined EV charging operation, intelligent monitoring, cloud-based forecasting, OCPP-based management, and collaborative cloud-edge control, showing that both local responsiveness and centralized optimization are important for real-time charging systems. However, the literature still lacks unified architectures that clearly partition control intelligence between edge and cloud layers. To address this issue, this article presents an edge-cloud collaborative architecture for real-time EV charging control based on edge-side sensing and actuation, cloud-side forecasting and optimization, synchronized control exchange, and adaptive feedback refinement. The results show reduced control response latency, improved charging stability, and stronger load coordination across distributed charging nodes compared with edge-only and cloud-only approaches. Overall, the study demonstrates that edge-cloud collaboration provides a scalable foundation for real-time EV charging control.
Keywords: EV charging control, edge-cloud collaboration, real-time charging systems, distributed charging networks, control latency, charging stability, cloud optimization, edge intelligence.

1. Introduction

Electric vehicle charging control is increasingly shifting from a purely centralized scheduling problem to a real-time cyber-physical coordination problem in which charger nodes, local controllers, monitoring devices, and supervisory services must interact under strict timing constraints. Modern charging-station operation is shaped not only by energy delivery but also by market dynamics, grid interaction, user concurrency, and rapidly changing demand conditions [1]. At the same time, intelligent monitoring studies show that EV charging environments now generate continuous operational signals that must be interpreted quickly if control actions are to remain relevant to live system state [2]. In such a setting, real-time charging control cannot rely exclusively on delayed cloud-side computation, because response quality degrades when network latency and centralized processing overhead grow. This makes architectural placement of control intelligence a central design issue in EV charging systems.

The recent literature already provides important pieces of this broader problem. Cloud-based forecasting studies have shown that remote computational resources are effective for demand estimation, optimization, and renewable-aware charging coordination when large-scale data aggregation is needed [3]. Large-scale IoT e-mobility architectures have also demonstrated that distributed EV platforms benefit from modular service organization and scalable backend coordination, especially when many devices and operational services interact simultaneously [4]. More recent IoT-based urban mobility work reinforces the same point by showing that smart mobility systems increasingly depend on distributed sensing, communication, and service orchestration rather than on isolated endpoint control [5]. These contributions establish that both local responsiveness and centralized intelligence are valuable, but they do not fully resolve how those roles should be partitioned in real-time EV charging control.

The core problem is that purely cloud-based and purely edge-based charging control each have structural limitations. A cloud-only design can optimize globally but may respond too slowly when local charging conditions change abruptly, while an edge-only design can react quickly but often lacks sufficient system-wide awareness to coordinate load and control decisions across distributed nodes. As EV charging infrastructure scales, this tradeoff becomes more severe because chargers must react to immediate local states while still aligning with network-level objectives. The missing element is a collaborative architecture in which edge nodes handle fast local decisions and cloud services handle broader prediction, optimization, and coordination. This is the problem statement that motivates the present study.

This problem matters because control delay and coordination weakness directly affect charging stability, service continuity, and network performance. If control decisions arrive too late, charger states may drift away from optimal operating points, local overload may intensify, and balancing opportunities across distributed nodes may be missed. If decisions remain too localized, one part of the charging network may react efficiently while the wider system becomes uncoordinated. Real-time EV charging control therefore requires an architecture that can combine fast edge-side responsiveness with cloud-side computational depth. Such a design is necessary if future charging infrastructure is expected to remain both scalable and operationally efficient.

This article presents an edge-cloud collaborative architecture for real-time EV charging control. The study focuses on task partitioning between edge and cloud layers, low-latency local control execution, cloud-assisted forecasting and optimization, synchronized decision exchange, and collaborative coordination across distributed charging nodes. Rather than treating edge and cloud as competing paradigms, the work develops a unified framework in which each layer performs the functions best suited to its latency, data, and coordination role. The following methodology defines the architectural layers, collaborative control workflow, synchronization logic, and performance metrics used to support this design.

2. Methodology

The proposed methodology is built as a collaborative control architecture in which real-time EV charging decisions are divided between edge-side execution and cloud-side coordination. The edge layer is responsible for fast local sensing, charger-state interpretation, immediate control enforcement, and transient disturbance handling, while the cloud layer is responsible for larger-scale prediction, optimization, and multi-node coordination. Research on cloud-edge collaborative optimal control for large-scale EV participation has shown that such partitioning can improve both local responsiveness and system-wide operational performance when decision functions are distributed according to timing and computational needs [6]. The present framework therefore assumes that real-time charging control should be neither entirely centralized nor entirely local. Control quality emerges from structured cooperation between fast edge action and slower but broader cloud intelligence.

At the edge layer, each charging node maintains a local operational state consisting of active charging sessions, current power draw, queue condition, charger occupancy, connection quality, and immediate feeder-side indicators. This local state is updated continuously and used for rapid control actions such as charging-rate adjustment, safety enforcement, session continuity control, and local anomaly handling. OCPP-based EV charging management research supports the use of protocol-grounded local control because charger-to-management communication already provides structured operational events that can be used to support near-real-time station decisions [7]. In the proposed design, the edge node is not treated as a passive relay to the cloud. Instead, it functions as the first control layer capable of stabilizing local charging behavior before slower global decisions arrive.

A collaborative prediction layer then connects edge operation to cloud intelligence. The cloud receives aggregated charger-state summaries, demand traces, and control-performance indicators from distributed edge nodes, after which it produces short-horizon forecasts for demand growth, congestion pressure, and multi-node load coordination. Online learning work for charging-station demand and event prediction has shown that forecast quality improves when models adapt to recent data rather than relying only on static historical training [8]. Intelligent charging-management platforms further demonstrate that distributed station data can be coordinated through platform-level logic to support broader control visibility across multiple charging points [9]. In the proposed framework, the cloud therefore acts as the predictive and coordinating layer that transforms distributed local observations into network-aware control guidance. This gives the architecture a forward-looking control capability that edge-only systems cannot achieve easily.

Once forecast and coordination outputs are available, the cloud transmits control advisories, optimization targets, and balancing parameters back to the edge layer. These outputs may include preferred charging-rate envelopes, queue-sensitive prioritization signals, demand redistribution weights, and feeder-aware operating constraints for individual nodes. A federated AI-OCPP framework for EV charging infrastructure supports this distributed intelligence model because scalable charging systems benefit when secure protocol management is combined with distributed learning and coordination logic [10]. In the present architecture, the cloud does not directly micromanage every charger action. Instead, it supplies updated guidance that the edge node interprets within its own real-time local state. This keeps the control loop coordinated without making the edge dependent on constant centralized command.

The synchronization mechanism between edge and cloud is a central component of the methodology. Each edge node performs local control updates at a short time step, while cloud coordination updates occur over a longer rolling horizon. The edge layer therefore reacts immediately to charger-state changes, while the cloud refines broader objectives after processing data from multiple nodes and external context sources. Edge-AI-based multi-criteria optimization studies indicate that combining local real-time information with broader contextual factors such as grid load, traffic, and user behavior improves charging decision quality under dynamic conditions [11]. Building on that insight, the proposed synchronization logic allows edge control to remain active between cloud updates while ensuring that local decisions are periodically realigned with system-wide coordination targets. This prevents local control from drifting away from network objectives.

A decision-merging layer is then used to reconcile edge urgency with cloud optimization. If local conditions require immediate intervention, such as sudden queue growth or charging instability, the edge action is executed first and then reported upward. If no urgent disturbance exists, the edge controller follows cloud-informed guidance while still preserving local safety and service continuity. This hierarchy is important because strict cloud authority would add delay in urgent cases, whereas strict edge autonomy would reduce global coordination quality over time. The proposed architecture therefore applies a priority-aware control rule in which edge-side immediacy and cloud-side optimality are balanced rather than forced into a single rigid command structure. This makes collaborative control practical under real operating variability.

The methodology also incorporates a feedback loop for performance adaptation. Edge nodes report execution outcomes, latency observations, control deviation, and local service stability back to the cloud, where these signals are used to update future coordination parameters and refine forecast quality. This continuous exchange allows the architecture to learn whether a given partitioning of responsibilities is producing stable results or whether local nodes require stronger autonomy or more cloud guidance. In effect, the architecture does not only perform collaborative control; it also evaluates the effectiveness of that collaboration over time. This makes the edge-cloud relationship adaptive rather than fixed.

3. Results and Discussion

The proposed edge-cloud collaborative architecture produced a clear reduction in control response latency when compared with cloud-only coordination and purely local edge-only control. Under low-demand conditions, all three control strategies maintained acceptable response behavior, but the differences became more pronounced as charger concurrency and control-event intensity increased. The cloud-only design exhibited the highest latency because every control update depended on remote computation and return-path transmission, while the edge-only strategy remained fast but showed less stable coordination under changing distributed conditions. By contrast, the collaborative architecture preserved low local reaction time while still benefiting from cloud-assisted optimization. This behavior is reflected in Figure 1, where the edge-cloud strategy maintains a consistently lower response-latency curve than the cloud-only design and a more stable control profile than the edge-only design across increasing demand intensity. The result indicates that real-time EV charging control benefits strongly when urgent local action and broader centralized intelligence are combined rather than separated.

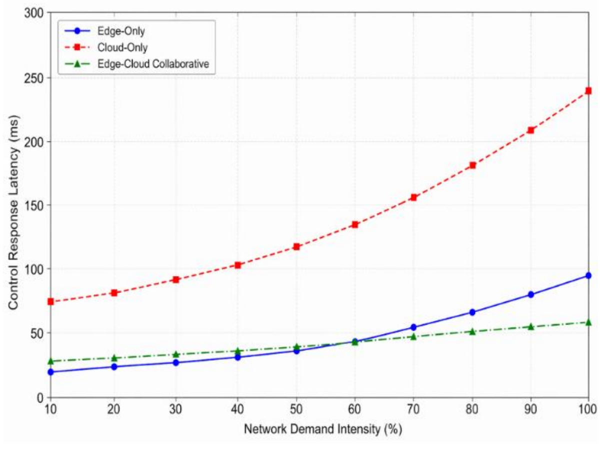

Figure 1. Control Response Latency Under Edge-Only, Cloud-Only, and Edge-Cloud Collaborative Strategies

A second important observation was that latency improvement alone did not explain the full performance gain. The edge-only approach could react rapidly, but it often produced locally efficient actions that were less aligned with system-wide balancing objectives when neighboring nodes experienced uneven demand. The cloud-only strategy improved coordination but suffered when control decisions had to pass through slower centralized loops. The collaborative architecture outperformed both because edge controllers handled fast charger-level updates while cloud services periodically refined network-level constraints and priorities. This reduced the conflict between immediacy and coordination that often limits large charging systems. The architecture therefore improved not just speed, but also the quality of control alignment across the network.

Charging stability also improved under collaborative control. In the baseline comparison strategies, sudden changes in charger occupancy or queue buildup sometimes produced oscillatory adjustments because the control loop was either too delayed or too localized. Under the proposed design, edge nodes absorbed fast disturbances locally, while cloud-side updates smoothed broader system behavior by providing refreshed operating envelopes and coordination targets. This reduced abrupt charging-rate swings and made local control transitions more stable over time. The result is especially relevant in public charging environments where unstable control can degrade both user experience and grid-side performance. Real-time stability therefore emerged as a direct advantage of edge-cloud task partitioning.

The coordination results across distributed nodes were similarly strong. In the absence of collaborative logic, different parts of the charging network tended to drift toward uneven utilization because either the local layer acted without sufficient system awareness or the centralized layer responded too slowly to local change. The collaborative design improved this by merging edge urgency with cloud guidance, which allowed local nodes to remain responsive while still respecting broader balancing objectives. This behavior is shown in Figure 2, where the edge-cloud strategy sustains stronger charging stability and more consistent load coordination across distributed EV charging nodes than the comparison approaches. The result confirms that real-time coordination quality depends not only on better local control or better global optimization alone, but on how those two capabilities interact during operation.

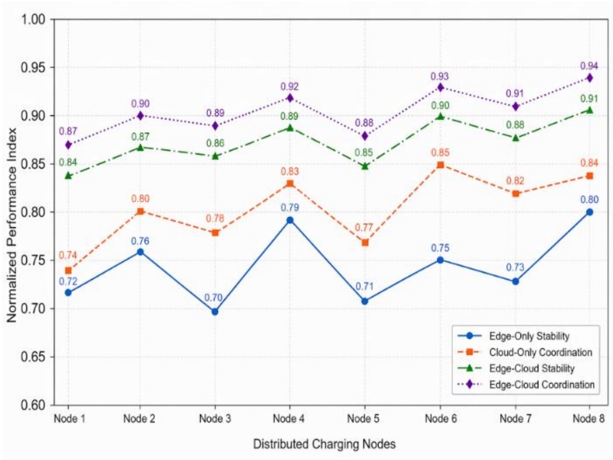

Figure 2. Charging Stability and Load Coordination Performance Across Distributed EV Charging Nodes

Another notable outcome was the usefulness of the feedback loop between edge execution and cloud refinement. When control performance degraded at one or more nodes, the cloud could revise coordination parameters based on returned execution outcomes and recent local conditions. This made the architecture more adaptive than fixed partitioning schemes in which edge and cloud roles remain static regardless of demand or disturbance. As a result, collaborative control remained robust even as local service conditions changed over time. The framework therefore demonstrated not only distributed decision execution, but also distributed performance learning about how that execution should be tuned.

Taken together, the results show that edge-cloud collaboration offers a practical architectural advantage for real-time EV charging control. The framework reduced response latency, improved charging stability, strengthened coordination across nodes, and preserved better operational coherence than cloud-only or edge-only alternatives. These gains arose because the architecture matched control functions to the layer best suited to them instead of forcing all control logic into a single computational location. This suggests that future EV charging systems should be designed around collaborative control partitioning if they are expected to operate efficiently under large-scale and time-sensitive conditions.

4. Conclusion

This study presented an edge-cloud collaborative architecture for real-time EV charging control with emphasis on low-latency local execution, cloud-assisted forecasting and coordination, synchronized decision exchange, and adaptive control refinement. The proposed framework treated EV charging infrastructure as a distributed cyber-physical control environment in which purely centralized and purely local strategies are both structurally limited. By assigning urgent execution to the edge layer and broader optimization to the cloud layer, the architecture created a practical basis for responsive yet coordinated charging control. This makes the design especially suitable for charging networks where local disturbances and global balancing requirements coexist continuously.

The results showed that the collaborative architecture reduced control response latency, improved charging stability, and strengthened load coordination across distributed charging nodes compared with cloud-only and edge-only baselines. The gains were most visible under rising charger concurrency, where the cloud-only strategy became too delayed and the edge-only strategy became too locally constrained. The feedback loop between edge execution and cloud refinement also enhanced adaptability by allowing the control framework to evolve with changing operating conditions. These findings show that architectural collaboration between computational layers is a decisive factor in real-time EV charging performance.

A practical implication of this work is that future charging ecosystems should not frame edge and cloud as competing control paradigms. Their joint use is more valuable than their isolated use when the objective is to achieve both speed and coordination under distributed demand. Further development can extend the framework through federated multi-edge learning, digital-twin-assisted validation of collaborative control, secure edge-cloud trust management, and integration with renewable-aware grid optimization. Additional study may also examine mixed public-private charging fleets, cross-operator coordination, and explainable control policies for operator-facing supervision. These directions would strengthen the role of edge-cloud collaboration as a scalable foundation for real-time EV charging control.

References

1. Motlagh, S. G., Oladigbolu, J., & Li, L. (2025). A review on electric vehicle charging station operation considering market dynamics and grid interaction. *Applied Energy*, *392*, 126058.
2. Martins, J. A., & Rodrigues, J. M. (2025). Intelligent monitoring systems for electric vehicle charging. *Applied Sciences*, *15*(5), 2741.
3. Aldossary, M., Alharbi, H. A., & Ayub, N. (2024). Optimizing electric vehicle (EV) charging with integrated renewable energy sources: A cloud-based forecasting approach for eco-sustainability. *Mathematics*, *12*(17), 2627.
4. Beránek, M., Feuerlicht, G., Kucera, O., & Kovár, V. (2023). An Architecture for a Large-Scale IoT e-Mobility Solution. In *ICEIS (1)* (pp. 733-740).
5. Reis, M. J., Branco, F., Gupta, N., & Serôdio, C. (2025). An IoT Architecture for Sustainable Urban Mobility: Towards Energy-Aware and Low-Emission Smart Cities. *Future Internet*, *17*(10), 457.

6. Lu, X., & Wang, L. (2024). Cloud-edge collaboration control strategy for electric vehicle aggregators participating in frequency and voltage regulation. *IEEE Open Journal of Vehicular Technology*, *5*, 1532-1544.
7. Hsaini, S., Ghogho, M., & Charaf, M. E. H. (2022). An OCPP-based approach for electric vehicle charging management. *Energies*, *15*(18), 6735.
8. Zamee, M. A., Han, D., Cha, H., & Won, D. (2023). Self-supervised online learning algorithm for electric vehicle charging station demand and event prediction. *Journal of Energy Storage*, *71*, 108189.
9. Dalamagkas, C., Melissianos, V. D., Papadakis, G., Georgakis, A., Nikiforidis, V. M., & Hrissagis-Chrysagis, K. (2025). The Open V2X Management Platform: An intelligent charging station management system. *Information Systems*, *129*, 102494.
10. Hossen, M. S., Sarker, M. T., Nabi, M. S., Bannah, H., Ramasamy, G., & Eng Eng, N. (2025). Federated AI-OCPP Framework for Secure and Scalable EV Charging in Smart Cities. *Urban Science*, *9*(9), 363.
11. Rahmati, M. (2025). Edge-AI based multi-criteria optimization framework for dynamic EV charging with real-time grid load, traffic, and user behavior integration. *Computing*, *107*(9), 178.