

**Supplementary Information for  
“Unlocking electric vehicle flexibility in  
microgrids through lightweight  
on-charger coordination”**

Yehui Li<sup>1</sup>, Xueyuan Cui<sup>2</sup>, Chongqing Kang<sup>3</sup>, Thomas Morstyn<sup>4</sup>,  
Yi Wang<sup>1</sup>

<sup>1</sup>Department of Electrical and Computer Engineering, The University  
of Hong Kong, Hong Kong SAR, China.

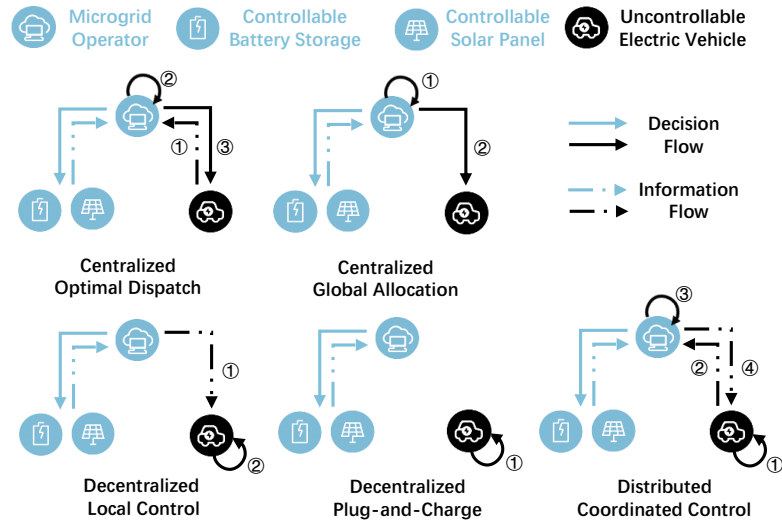
<sup>2</sup>Department of Engineering and Applied Sciences,, Harvard University,  
Cambridge, USA.

<sup>3</sup>Department of Electrical Engineering, Tsinghua University, Beijing,  
China.

<sup>4</sup>Department of Engineering Science, University of Oxford, Oxford, UK.

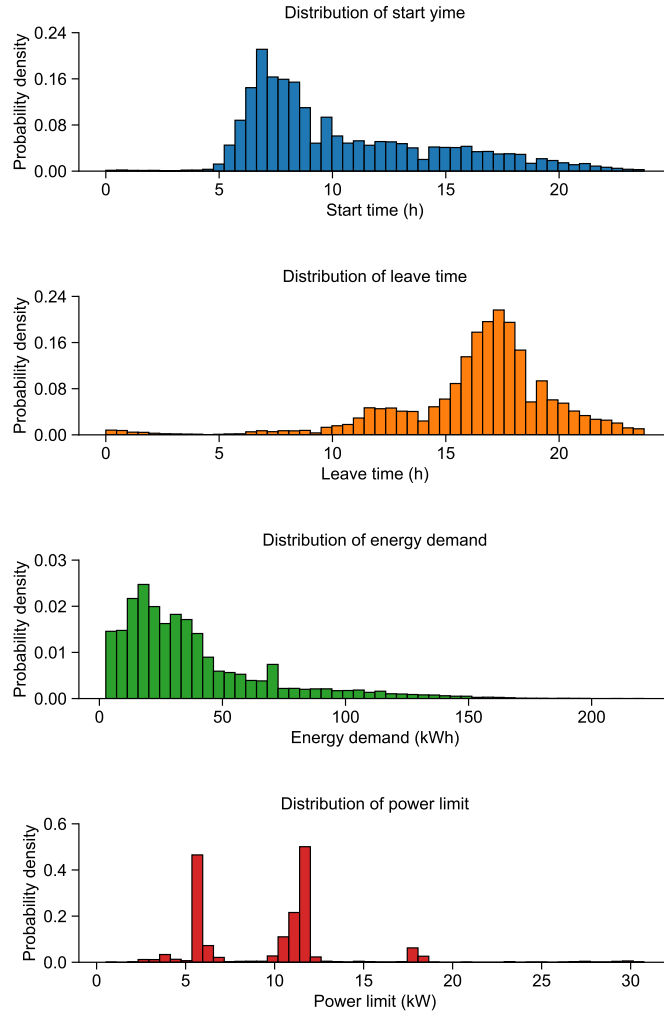
Contributing authors: [yhli@eee.hku.hk](mailto:yhli@eee.hku.hk); [xcui@seas.harvard.edu](mailto:xcui@seas.harvard.edu);  
[cqkang@tsinghua.edu.cn](mailto:cqkang@tsinghua.edu.cn); [thomas.morstyn@eng.ox.ac.uk](mailto:thomas.morstyn@eng.ox.ac.uk);  
[yiwang@eee.hku.hk](mailto:yiwang@eee.hku.hk);

## Supplementary Figures



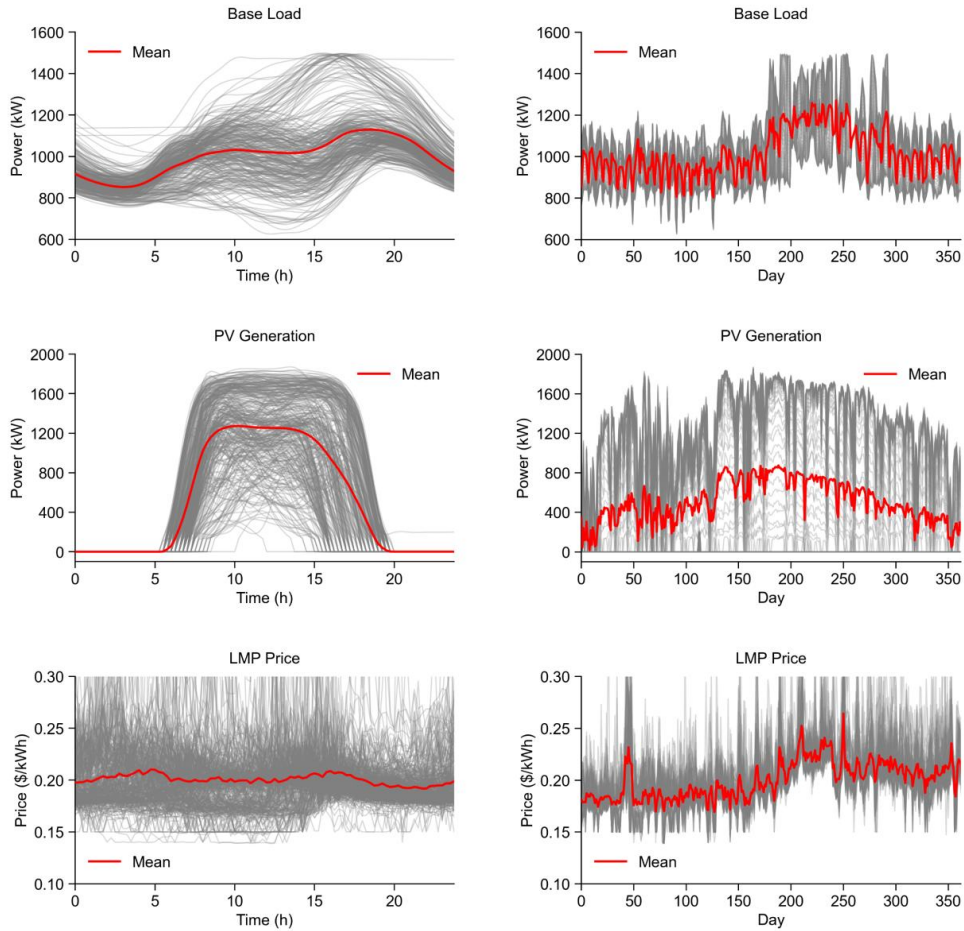
**Fig. 1** Schematic comparison of decision and information flows across distinct EV control architectures.

While the microgrid operator maintains direct control over stationary resources (batteries and solar panels) in all scenarios, the interaction with EVs varies significantly. In Centralized Optimal Dispatch, the operator collects detailed state information (including sensitive data) from EVs to issue direct regulation commands. In Centralized Global Allocation, the operator assigns a uniform average power limit without requiring upward data transmission. In Decentralized Local Control, EVs optimize schedules based on broadcast price signals. In Decentralized Plug-and-Charge, the unmanaged EVs draw maximum charging power upon connection. In Distributed Coordinated Control, our framework enables a privacy-preserving mechanism where the microgrid and EVs iteratively exchange decision variables to converge on an optimal solution without exposing raw user data.




**Fig. 2** Probability distribution of start time, leave time, energy demand, and power limit for 30000 charging sessions.

The ACN dataset is collected from 120 EV service equipment of three workplaces charging sites in California, including California Institute of Technology (Caltech), Jet Propulsion Laboratory (JPL), and an office building. The peak arrival times for EVs are concentrated during the morning commute, while the peak departure times are concentrated during the afternoon commute. The energy demand is concentrated in the 10 kWh to 40 kWh range, indicating that most EV users charge for daily commuting rather than waiting until the battery is completely depleted. The maximum charging power is concentrated in the 3 kW to 20 kW range, indicating that most chargers in this region are Level 1 or Level 2 commercial chargers.



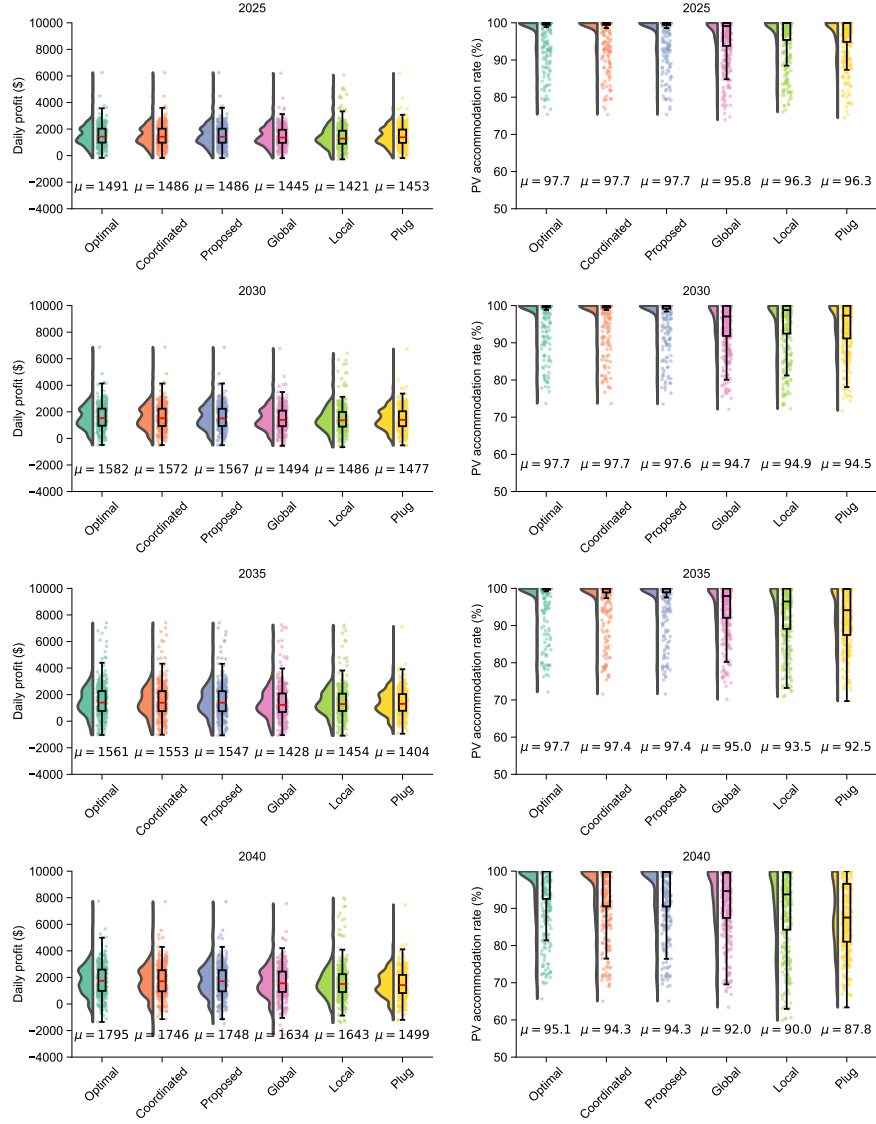
**Fig. 3** Stochastic profiles of electrical load, photovoltaic (PV) generation, and electricity prices.

The figure illustrates the temporal variability and uncertainty modeled for base load (top row), PV generation (middle row), and locational marginal price (LMP) (bottom row). The left column (0–24 h) captures intraday dynamics, identifying a characteristic evening peak in demand and the diurnal cycle of solar output. The right column (0–365 days) reveals long-term seasonal trends, notably the variance in solar intensity throughout the year. Crucially, while the mean trajectories (solid red lines) exhibit stable patterns, the ensemble of stochastic realizations (grey lines) demonstrates significant volatility and uncertainty, particularly in market prices and peak-load conditions. This variance highlights the necessity of robust planning against extreme scenarios in energy system modeling.



**Fig. 4 Heatmap visualization of PV accommodation efficiency and carbon footprint of different months throughout the year.**

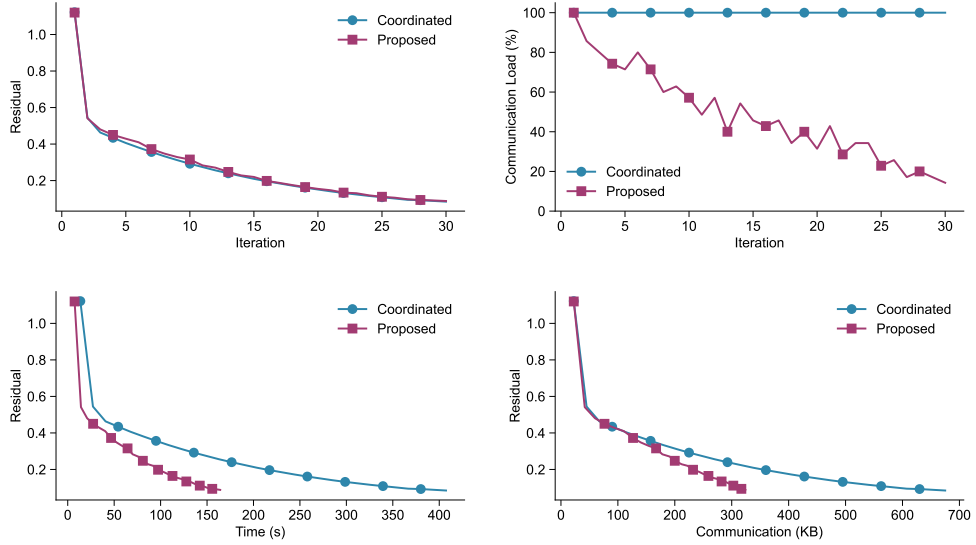
Rows represent distinct control frameworks, while columns correspond to monthly aggregates. Top: PV Accommodation Rate (higher is better). During summer and winter, higher heating and cooling demands enable nearly all renewable energy generation to be absorbed. There is a distinct seasonal dip in late spring (May–June), driven by surplus generation exceeding load capacity. In this challenging interval, the proposed method achieves an accommodation rate of 89.5%, effectively bridging the gap between the theoretical optimal (90.2%) and the unmanaged plug-and-charge baseline (83.5%). Bottom: Carbon Emission (lower is better). Since load is positively correlated with carbon emissions, high energy consumption during winter and summer leads to elevated carbon emissions. The proposed approach consistently yields the lowest emission tier alongside the optimal strategy. In summary, these results validate that the proposed distributed coordinated control can realize the benefits of centralized optimization in maximizing renewable uptake and minimizing emissions.



**Fig. 5** Projection of microgrid operational performance under different control frameworks from 2025 to 2040.

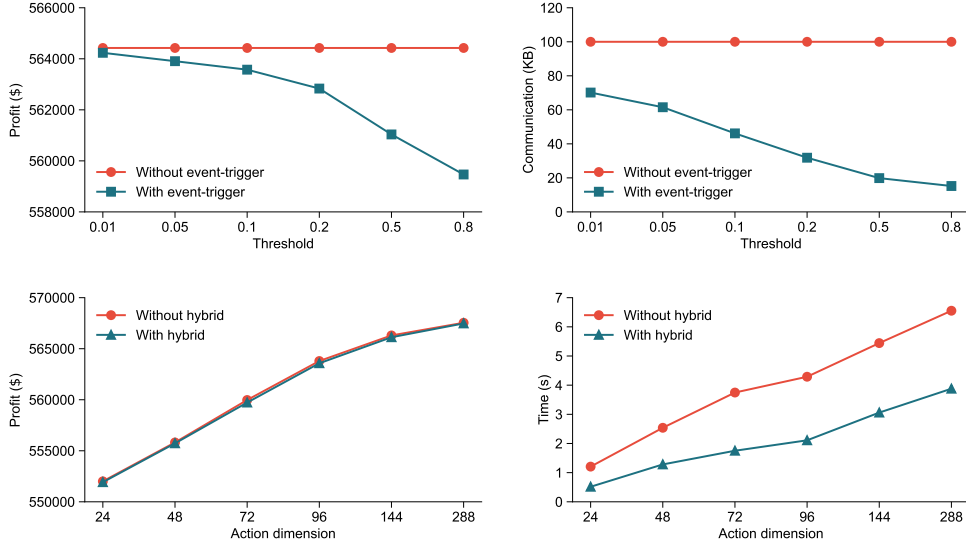
The figure presents the statistical distribution of daily profit and PV accommodation rate across four milestone years. The violin plots, overlaid with raw data points, visualize the variance, with  $\mu$  denoting the mean value. Left column: Distributions of daily economic profits. By 2040, the proposed distributed control exhibits a tighter distribution around the high-profit mean compared to the decentralized benchmarks, indicating more consistent daily returns and reduced financial volatility. Right column:

Distributions of PV accommodation rates. The long tail observed in the plug-and-charge and local plots in 2035 and 2040 indicates frequent occurrences of low-efficiency days with high curtailment. In contrast, the proposed method maintains a compact distribution near the 100% ceiling, close to the optimal strategy. Overall, the results confirm that the proposed method effectively approximates the global optimum in both economic and technical metrics, sustaining the advantage as grid complexity increases over the 15-year horizon.



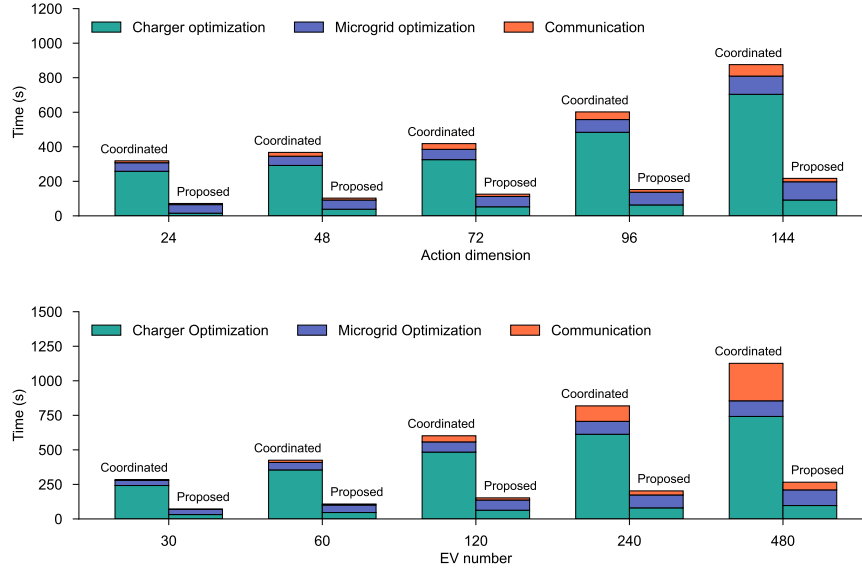
**Fig. 6** Convergence dynamics and communication efficiency analysis of the proposed distributed framework.

The figure compares the iterative performance of the proposed method with the traditional coordinated baseline in terms of mathematical convergence and resource consumption. Top left: Residual evolution over iterations. Both methods exhibit similar asymptotic convergence rates, indicating comparable mathematical stability. Top right: Communication load per iteration. Unlike the coordinated benchmark, which maintains a static 100% communication overhead, the proposed method implements an adaptive mechanism that progressively reduces data exchange as the system approaches equilibrium. Bottom: Computational efficiency in terms of computational time and communication volume. Consequently, the reduction in per-iteration communication load translates into significant practical benefits: the proposed framework achieves the same convergence precision with approximately three times faster execution time and substantially lower bandwidth requirements. This confirms the method’s superior scalability for bandwidth-constrained microgrid environments.



**Fig. 7 Ablation study quantifying the contributions of event-triggered communication and hybrid optimization modules.**

The top row evaluates the sensitivity to the event-trigger threshold parameter. As the threshold increases, the communication volume drops drastically, while the total profit decreases only marginally. This indicates a favorable trade-off where the proposed method effectively filters redundant data transmission with negligible impact on economic optimality (e.g., at threshold 0.2, communication is reduced by 70% with  $\leq 1\%$  profit loss). The bottom row validates the scalability of the hybrid mechanism across varying action dimensions. Ideally, the profit curves for the hybrid mode and traditional one perfectly align, confirming that the hybrid module enables acceleration without compromising solution accuracy. Conversely, the computation time analysis (d) demonstrates that the hybrid approach significantly mitigates computational burden, exhibiting superior scalability for high-dimensional decision-making tasks compared to the baseline.



**Fig. 8 Decomposition of computational latency and scalability analysis across varying problem complexities.**

The stacked bar charts detail the execution time breakdown into three critical phases: local charger optimization, microgrid dispatch, and data communication. The coordinated baseline suffers from a computational bottleneck dominated by the charger optimization phase, which grows linearly with dimension. In contrast, the proposed method significantly suppresses this local computation time, reducing total latency by 75%. As the network size expands, the coordinated approach exhibits severe scalability issues, where both computation and communication burdens increase drastically at 480 EVs. The proposed framework, however, maintains a low and stable time profile ( $\leq 250$ s even at 480 EVs), demonstrating its capability for real-time decision-making in large-scale infrastructure systems.

# Supplementary Notes

## Frequency Control Optimization

This section compares the classical droop control with the proposed strategy for regulating microgrid frequency. The system frequency dynamics are governed by the swing equation:

$$\frac{df_t}{dt} = \frac{P_t^{\text{net}}}{M}, \quad (1)$$

where  $f_t$  denotes the system frequency at time  $t$ , and  $M$  represents the system inertia constant.  $P_t^{\text{net}}$  reflects the net power imbalance, defined as:

$$P_t^{\text{net}} = P_t^{\text{pv}} + P_t^{\text{es,dis}} - P_t^{\text{base}} - P_t^{\text{es,ch}} - \sum_{n \in \mathcal{N}} P_{n,t}^{\text{ev}} \quad (2)$$

## Classical Droop Control Strategy

In the classical decentralized control, EVs adjust their charging power linearly in response to frequency deviations. The total regulation power requirement  $P_t^{\text{req}}$  is determined by the system droop characteristic:

$$P_t^{\text{req}} = -K^{\text{droop}} (f_t - f^{\text{nom}}), \quad (3)$$

where  $K^{\text{droop}}$  is the frequency droop coefficient and  $f^{\text{nom}}$  is the nominal frequency (e.g., 60 Hz).

Ideally, the remaining power gap  $P_t^{\text{req}}$ , after accounting for the ES response, is allocated equally among all connected EVs. The target power setpoint for each EV  $n$  is derived as:

$$P_{n,t}^{\text{ev}} = -\frac{P_t^{\text{req}}}{|\mathcal{N}|}, \quad (4)$$

where  $|\mathcal{N}|$  denotes the total number of EVs projected to participate.

The actual charging power  $P_{n,t}^{\text{ev}}$  is constrained by the physical capacity limits and ramp rate limits  $R_n^{\text{ramp}}$  of the charger:

$$\begin{aligned} -P_{n,\max}^{\text{ev}} &\leq P_{n,t+\Delta t}^{\text{ev}} \leq P_{n,\max}^{\text{ev}} \\ \left| \frac{P_{n,t+\Delta t}^{\text{ev}} - P_{n,t}^{\text{ev}}}{\Delta t} \right| &\leq R_{n,\text{ev}}^{\text{max}} \end{aligned} \quad (5)$$

where  $\Delta t$  is the control time step;  $P_{n,\max}^{\text{ev}}$  and  $R_{n,\text{ev}}^{\text{max}}$  is the maximum ramp rate of EV  $n$ . While simple, this method reacts passively to current deviations without anticipating future system dynamics.

## Proposed Control Strategy

Since individual EVs lack knowledge of the total participating vehicle population  $|\mathcal{N}|$  and system frequency response characteristics, an adaptive estimation mechanism is

employed. Each EV estimates the aggregate frequency response relationship:

$$P_t^{\text{req}} = -\hat{K}(f_t - f^{\text{nom}}) \quad (6)$$

where  $\hat{K}$  is the estimated system frequency response coefficient.

Assuming symmetric behavior among EVs, the total regulation power can be approximated as:

$$P_t^{\text{req}} = \hat{N}_t P_{n,t}^{\text{ev}} \quad (7)$$

where  $\hat{N}_t$  is the estimated number of participating EVs. By observing frequency deviations  $\Delta f_t$  and the local regulation effort  $P_{n,t}^{\text{ev}}$ , each EV iteratively updates these estimates at each time step.

At time  $t$ , each EV  $n$  solves a predictive optimization problem to determine its optimal regulation power:

$$\min_{P_{n,t+\Delta t}^{\text{ev}}} J = w^{\text{freq}} \left( \tilde{f}_{t+\Delta t} - f^{\text{nom}} \right)^2 + w^{\text{smooth}} \left( P_{n,t+\Delta t}^{\text{ev}} - P_{n,t}^{\text{ev}} \right)^2 \quad (8)$$

where the weights  $w^{\text{freq}}$  and  $w^{\text{smooth}}$  balance frequency tracking performance and control smoothness.

The frequency prediction incorporates the estimated collective EV response:

$$\tilde{f}_{t+\Delta t} = f_t + \frac{\tilde{P}_{t+\Delta t}^{\text{net}}}{M} \Delta t, \quad (9)$$

where the predicted net power imbalance is:

$$\tilde{P}_{t+\Delta t}^{\text{net}} = P_t^{\text{pv}} + P_t^{\text{es,dis}} - P_t^{\text{base}} - P_t^{\text{es,ch}} - \hat{N}_t P_{n,t+\Delta t}^{\text{ev}} \quad (10)$$

Here, the local decision variable  $P_{n,t+\Delta t}^{\text{ev}}$  is scaled by the estimated population size  $\hat{N}_t$  to predict the aggregate EV contribution to frequency regulation.

The actual charging power  $P_{n,t}^{\text{ev}}$  is also constrained by the physical capacity limits and ramp rate limits  $R_n^{\text{ramp}}$  of the charger:

$$\begin{aligned} -P_{n,\max}^{\text{ev}} &\leq P_{n,t+\Delta t}^{\text{ev}} \leq P_{n,\max}^{\text{ev}} \\ \left| \frac{P_{n,t+\Delta t}^{\text{ev}} - P_{n,t}^{\text{ev}}}{\Delta t} \right| &\leq R_{n,\text{ev}}^{\text{max}} \end{aligned} \quad (11)$$

## Method description

To ensure algorithmic clarity, it is essential to distinguish between the two nested iterative loops in our distributed optimization framework. The detailed process is described in pseudocode Algorithm 1.

- The outer loop, indexed by  $(k)$ , governs the global coordination between the cloud microgrid and the distributed edge chargers. In each outer iteration  $k$ , the cloud computers  $\mathcal{H}_t^{(k)}$ , and each charger computes the optimal response  $P_t^{\text{ev},(k)}$ .
- The inner loop, indexed by  $(j)$ , represents the internal solver-free computational process executed within on-charger microcontrollers for efficient computation. For each outer iteration  $k$ , given a fixed outer-loop signal  $\widehat{\mathcal{H}}_t^{(k)}$ , the charger iterates to solve the local constrained optimization sub-problem. The inner loop converges to a final value  $P_{n,t}^{\text{ev},(j)}$  as the update for the outer loop  $P_{n,t}^{\text{ev},(k)}$ . This value is communicated back to the cloud, completing the cycle for outer iteration  $k$ .

---

**Algorithm 1:** Lightweight On-charger Coordination

---

**Input:** Dispatch horizon  $\mathcal{T}$ , Fleet  $\mathcal{N}$ , Charging horizon  $\mathcal{T}_n$ , Thresholds

$\Delta_{\mathcal{H}}, \Delta_P$ , Tolerances  $\epsilon_{out}, \epsilon_{in}$ .

**Output:** Optimal charging variables  $\mathbf{P}_n^{\text{ev}}, \mathbf{E}_n^{\text{ev}}$ .

```
1 while outer residuals >  $\epsilon_{out}$  do // Outer Loop: Cloud-Edge Coordination
   // ----- Phase 1: Cloud Processing (Microgrid) -----
2    $\mathbf{P}^{\text{ev},(k)} \leftarrow \arg \min$  Upper-Level Problem
3   Compute signal:  $\mathcal{H}_t^{(k)} \leftarrow (\sum_n \hat{P}_{n,t}^{\text{ev},(k)} - P_t^{\text{ev},(k)}) - \rho^{-1} \lambda_t^{\text{ev},(k)}$ 
   /* Event-Triggered Downlink */
4   if  $\|\mathcal{H}^{(k)} - \hat{\mathcal{H}}^{(k-1)}\|_2 > \Delta_{\mathcal{H}}$  then
5     Update Cache:  $\hat{\mathcal{H}}^{(k)} \leftarrow \mathcal{H}^{(k)}$  and Broadcast  $\hat{\mathcal{H}}^{(k)}$ .
   // ----- Phase 2: Edge Processing (Parallel EVs) -----
6   for  $n \in \mathcal{N}$  in parallel do
7     while inner residuals >  $\epsilon_{in}$  and  $j \leq J$  do // Inner Loop: Tiny MPC
8       /* Riccati Recursion */
9       Update LQR parameter  $(\tilde{Q}, \tilde{R}, \tilde{q}, \tilde{r})$ 
10      if  $|\mathcal{T}_n| \gg T_{\text{thresh}}$  then
11        Load offline steady-state gains:  $K_{n,t} \leftarrow K_{\infty}, P_{n,t} \leftarrow P_{\infty}$ 
12        Backward recursion for linear terms only:  $d_{n,t}, p_{n,t}$ 
13      else
14        Backward Riccati Recursion:  $P_{n,t}, K_{n,t}, d_{n,t}, p_{n,t}$ 
15      /* Forward Pass */
16       $P_{n,t}^{\text{ev},(j+1)} \leftarrow -K_{n,t} E_{n,t}^{\text{ev},(j+1)} - d_{n,t}$ 
17       $E_{n,t+1}^{\text{ev},(j+1)} \leftarrow A E_{n,t}^{\text{ev},(j+1)} + B P_{n,t}^{\text{ev},(j+1)}$ 
18      /* Projection & Dual Update */
19      Update Slack via box constraints:
20       $Z_{n,t}^{(j+1)} \leftarrow \min(\max(E_{n,t}^{\text{ev},(j+1)} + \delta^{-1} \mu_{n,t}^{(j)}, 0), E_n^{\text{req}})$ 
21       $W_{n,t}^{(j+1)} \leftarrow \min(\max(P_{n,t}^{\text{ev},(j+1)} + \delta^{-1} \sigma_{n,t}^{(j)}, 0), P_n^{\text{max}})$ 
22      Update duals:
23       $\mu_n^{(j+1)} \leftarrow \mu_n^{(j)} + \delta(E_n^{(j+1)} - Z_n^{(j+1)})$ 
24       $\sigma_n^{(j+1)} \leftarrow \sigma_n^{(j)} + \delta(P_n^{(j+1)} - W_n^{(j+1)})$ 
25      Increment iteration:  $j \leftarrow j + 1$ 
26      Obtain optimal trajectories:  $\mathbf{P}_n^{\text{ev},(k+1)} \leftarrow \mathbf{P}_n^{\text{ev},(j)}$ 
27      /* Event-Triggered Uplink */
28      if  $\|\mathbf{P}_n^{\text{ev},(k+1)} - \hat{\mathbf{P}}_n^{\text{ev},(k)}\|_2 > \Delta_P$  then
29        Update cache:  $\hat{\mathbf{P}}_n^{\text{ev},(k+1)} \leftarrow \mathbf{P}_n^{\text{ev},(k+1)}$  and upload to cloud.
   // ----- Phase 3: Consensus Update -----
30   Update Price:  $\lambda_t^{\text{ev},(k+1)} \leftarrow \lambda_t^{\text{ev},(k)} + \rho \left( |\mathcal{N}|^{-1} \sum_n \hat{P}_{n,t}^{\text{ev}} - |\mathcal{N}|^{-1} P_t^{\text{ev}} \right)$ 
31   Increment iteration:  $k \leftarrow k + 1$ 
```

---