

# Scalable Uncertainty-Aware Decision Frameworks for Energy-Mobility Cyber-Physical Systems

Rishav Sen  
Vanderbilt University  
Nashville, USA  
rishav.sen@vanderbilt.edu

## ABSTRACT

Energy-mobility cyber-physical systems (EM-CPS) couple stochastic mobility demand with time-coupled energy tariffs, creating decision problems where local charging actions can trigger system-level peaks and long-horizon costs. This work progresses from classical optimization to online decision-making and incentives. First, we study mixed-fleet public transit scheduling under charger availability and state-of-charge (SoC) constraints, using hierarchical MILP decomposition to obtain feasible, low-cost schedules on real transit data. Second, we address workplace vehicle-to-building (V2B) charging under persistent uncertainty in arrivals and building loads, developing methods that combine rolling-horizon control with constraint-aware learning to improve peak-driven objectives while preserving feasibility. Third, we design CONSENT, a negotiation mechanism for strategic EV users guaranteeing strategy-proofness, individual rationality, and budget feasibility under uncertainty. We validate the full stack in calibrated simulation grounded in real-world survey, telemetry, and building-load data.

## KEYWORDS

Energy-mobility cyber-physical systems; Vehicle-to-Building; Sparse rewards; Uncertainty handling; Mechanism design; Constrained Decision making

### ACM Reference Format:

Rishav Sen. 2026. Scalable Uncertainty-Aware Decision Frameworks for Energy-Mobility Cyber-Physical Systems. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 3 pages. <https://doi.org/10.65109/PDLS8233>

## 1 INTRODUCTION

Energy-mobility cyber-physical systems require decisions that couple discrete assignments with time-coupled energy constraints. Feasibility depends on limited resources (e.g., chargers and depot power) and dynamic states (e.g., battery levels), while costs depend on operational objectives and tariffs. These coupled constraints and peak-driven objectives are now well documented across electrified fleets and managed charging: mixed-fleet transit scheduling becomes a large combinatorial assignment problem under energy and charging constraints [17], and unmanaged charging can materially increase site peaks and demand-charge exposure [10].

We begin with mixed-fleet bus scheduling, where diesel and electric vehicles (EVs) must be assigned to fixed service blocks under charger availability and SoC constraints, yielding a structured, largely offline combinatorial optimization problem [15]. We then turn to workplace Vehicle-to-Building (V2B) charging, where decisions are made online under uncertain arrivals and building loads, and costs are dominated by sparse billing-period peak events [8, 16, 21]. Unlike Vehicle-to-Grid (V2G) systems, which target grid-level services such as frequency regulation and energy arbitrage on behalf of a utility or aggregator, V2B positions the building operator as the primary decision-maker minimizing a *demand charge*: a fee determined by the single highest power draw over the billing period; empirical evidence confirms that EV charging demand can materially reshape monthly peaks under common rate structures [6]. V2B decision-making is shaped by two coupled challenges: (i) sparse, peak-dominated objectives induced by demand charges alongside uncertainty in arrivals, departures, energy requirements, and building loads; and (ii) strategic participation, as users retain private information about their flexibility and may reject operator-assigned schedules when incentives are misaligned. The remainder of this paper describes these challenges, covering hierarchical combinatorial optimization, learning under sparsity and uncertainty, incentive-compatible negotiation, and simulation-based validation.

## 2 CLASSICAL OPTIMIZATION: A HIERARCHICAL APPROACH

We begin with a setting where the main difficulty is feasibility under coupled physical constraints. The task is to assign a limited fleet to fixed service blocks while respecting charger availability, depot power limits, and battery state-of-charge (SoC) dynamics for electric buses. We formulate this as a mixed integer linear program (MILP) with binary trip to bus assignment variables, time-coupled SoC updates, and charging infrastructure constraints (e.g., charger occupancy) [15]. While a monolithic MILP captures the full problem, it can become slow as the number of blocks, buses, and charging decisions increases. To improve tractability, we use a hierarchical decomposition that partitions resources into tiers based on requirements. We schedule the first tier (electric buses) first, since their assignments are tightly coupled with charging feasibility. The second tier (hybrid/diesel buses) covers the remaining blocks without requiring explicit charging decisions. If the second-stage assignment is infeasible, we iteratively demote the least efficient electric bus and re-solve until feasibility is restored. On one month of CARTA data (from Chattanooga, TN, USA), this approach produces feasible mixed-fleet schedules and reduces operating cost relative to the observed deployment and a state-of-the-art baseline [17].



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems ([www.ifaamas.org](http://www.ifaamas.org)). <https://doi.org/10.65109/PDLS8233>

### 3 LEARNING UNDER SPARSITY & UNCERTAINTY

Month-long demand charges make the objective peak-driven and sparse, while uncertain EV arrivals, requirements, and building loads hinder predictability, complicating both optimization and data-driven control. In V2B, a single rare peak can dominate cost, making long-horizon credit assignment brittle, and uncertain arrivals/departures and volatile loads render deterministic schedules fragile. Prior EV-charging work typically addresses these challenges via (i) receding-horizon optimization for peak shaving and bill reduction under hard constraints [21], or (ii) reinforcement-learning (RL) controllers that adapt online without perfect forecasts but can be sample-inefficient under delayed, sparse rewards and may violate constraints without careful structure [2, 20]. More broadly, feasibility cannot be enforced post-hoc: safe and constrained RL formalizes it as essential, yet still depends on problem structure to remain stable in long-horizon, high-variance [1].

The online V2B formulation in [16] addresses uncertainty by combining receding-horizon decision making with Monte-Carlo sampling of plausible futures, estimating peak risk without scenario-tree blowup and exposing a compute-performance trade-off suitable for real-time operation. Overall, peak-dominated tariffs and hard operational limits require hybrid decision pipelines: feasibility and uncertainty handling are enforced by a domain-specific optimization layer, while learning captures long-horizon demand-charge effects that are difficult to represent with purely myopic objectives [8]. We validate these methods in OPTIMUS [19], a calibrated discrete-event testbed for workplace V2B that captures stochastic arrivals and departures, SoC dynamics, infrastructure limits, and peak-driven tariffs, and supports negotiation routines (Section 4). Realism is anchored by a survey of EV behavior and flexibility, and by telemetry data from real building-load traces that govern peak formation, enabling stress tests that probe robustness under rare peak events and correlated demand patterns [19].

### 4 AGENTS WITH PRIVATE INFORMATION

The learning-and-control pipeline above assumes that computed schedules are followed, but in practice, V2B operators can only shape charging based on user requests and opt-in constraints. EV users may refuse charging delays, misreport flexibility, or strategically manipulate stated requirements to extract incentives, especially when peak-driven objectives impose non-obvious trade-offs at the individual level. As a result, operator-side control is insufficient unless paired with a mechanism layer that elicits flexibility while preserving voluntary participation and operator viability.

As a result, operator-side control is insufficient unless paired with a mechanism layer that elicits flexibility while preserving voluntary participation and operator viability. There is a substantial line of work in multi-agent systems on incentive-compatible EV charging: Gerding et al. [5] design a strategy-proof online auction for EV charging under capacity constraints, and Stein et al. [18] extend this with a pre-commitment mechanism that guarantees truthful reporting while remaining flexible about charging timing—a setting structurally related to our per-vehicle departure constraints. Two-sided market formulations [4] further highlight that both operator and user incentives must be jointly managed.

CONSENT [14] (to be presented at AAMAS 2026) builds on this tradition but targets the distinct V2B setting, where demand-charge tariffs create non-obvious peak-driven trade-offs not present in network-capacity-constrained V2G markets. It proposes a negotiation framework that offers incentive-backed flexibility menus and couples acceptance outcomes with downstream uncertainty-aware control. The design is meant to form a “win-win” situation for the EV user and the building operator. Each EV user is treated as a strategic agent with private information about willingness to shift departure time or accept reduced SoC. The mechanism is constructed to satisfy three properties that are central for deployability: *strategy-proofness* (truthful reporting is the dominant strategy for each negotiation), *individual rationality* (participation benefits each user relative to opting out), and *budget feasibility* (total incentives are bounded by operator savings). These guarantees constrain menu construction and pricing, ensuring that economic benefits from flexibility are shared without creating incentives for manipulation. Empirical results reported indicate that incentive alignment can reduce schedule rejection and improve both user- and operator-side outcomes under uncertainty without sacrificing feasibility, thereby converting user autonomy into a controlled degree of freedom.

### 5 FUTURE WORK

A concrete next step is community-scale V2B coordination across multiple buildings, depots, and fleets, allocating aggregate power budgets over long horizons while each site enforces hard local feasibility [9]. This setting also amplifies EV user persistence: many users return repeatedly across days and weeks, enabling longitudinal tracking of recurring arrival patterns and energy needs, and allowing flexibility to be allocated proactively over the billing period to reduce peak risk rather than reacting day-by-day [11].

Additionally, the objective can be expanded beyond long-term demand charges to incorporate demand response (DR) signals that are operationally distinct and grid-critical. DR is triggered by external events (e.g., emergency curtailment requests, critical peak pricing, heat-wave stress) that are temporally sparse but can have a disproportionate impact on reliability and peak management [3, 12]. These signals add a sparse-event decision layer: the controller must curtail or reshape load on demand without violating EV service guarantees or destabilizing demand-charge performance. We plan to formulate joint demand-charge and DR objectives and adapt learning and negotiation for sparse DR events, motivating hierarchical architectures that separate community-level coordination from local execution [7, 13]. Another promising direction is the use of multi-agent RL (MARL), though robustness under rare events remains an open challenge [22].

### 6 ACKNOWLEDGEMENT

I would like to thank my advisor, Abhishek Dubey, my mentor, Ayan Mukhopadhyay, my labmates in the SCOPE Lab at Vanderbilt University, and our research partners at the Nissan Advanced Technology Center–Silicon Valley. This work was supported in part by the Chameleon Testbed, funded by the National Science Foundation.

## REFERENCES

- [1] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. 2017. Constrained policy optimization. In *International conference on machine learning*. PMLR, 22–31.
- [2] Abdullah Al Zishan, Moosa Moghimi Haji, and Omid Ardakanian. 2020. Adaptive control of plug-in electric vehicle charging with reinforcement learning. In *Proceedings of the eleventh ACM international conference on future energy systems*. 116–120.
- [3] Mohamed H Albadi and Ehab F El-Saadany. 2008. A summary of demand response in electricity markets. *Electric power systems research* 78, 11 (2008), 1989–1996.
- [4] Enrico Gerding, Sebastian Stein, Valentin Robu, Dengji Zhao, and Nicholas R Jennings. 2013. Two-sided online markets for electric vehicle charging. (2013).
- [5] Enrico H Gerding, Valentin Robu, Sebastian Stein, David C Parkes, Alex Rogers, and Nicholas R Jennings. 2011. Online mechanism design for electric vehicle charging. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*. 811–818.
- [6] Madeline Gilleran, Eric Bonnema, Jason Woods, Partha Mishra, Ian Doebber, Chad Hunter, Matt Mitchell, and Margaret Mann. 2021. Impact of electric vehicle charging on the power demand of retail buildings. *Advances in Applied Energy* 4 (2021), 100062.
- [7] Parisa Hajjaligol, Panayiotis Papadopoulos, Amin Moazami, and Mohammadreza Aghaei. 2025. A hierarchical energy management system for a cluster of buildings: Reinforcement learning and model predictive control (RL-MPC) approach. *Energy and Buildings* (2025), 116879.
- [8] Fangqi Liu, Rishav Sen, Jose Paolo Talusan, Ava Pettet, Aaron Kandel, Yoshinori Suzue, Ayan Mukhopadhyay, and Abhishek Dubey. 2025. Reinforcement Learning-based Approach for Vehicle-to-Building Charging with Heterogeneous Agents and Long Term Rewards. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems*. 1345–1353.
- [9] Milos Manic, Dumidu Wijayasekara, Kasun Amarasinghe, and Juan J Rodriguez-Andina. 2016. Building energy management systems: The age of intelligent and adaptive buildings. *IEEE Industrial Electronics Magazine* 10, 1 (2016), 25–39.
- [10] Matteo Muratori and Arthur Yip. 2024. *Projecting Electric Vehicle Electricity Demands and Charging Loads*. Technical Report. National Renewable Energy Laboratory (NREL), Golden, CO (United States).
- [11] Alfredo Nespoli, Emanuele Ogliari, and Sonia Leva. 2023. User behavior clustering based method for EV charging forecast. *IEEE Access* 11 (2023), 6273–6283.
- [12] Peter Palensky and Dietmar Dietrich. 2011. Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE transactions on industrial informatics* 7, 3 (2011), 381–388.
- [13] Giuseppe Pinto, Marco Savino Piscitelli, José Ramón Vázquez-Canteli, Zoltán Nagy, and Alfonso Capozzoli. 2021. Coordinated energy management for a cluster of buildings through deep reinforcement learning. *Energy* 229 (2021), 120725.
- [14] Rishav Sen, Fangqi Liu, Jose Paolo Talusan, Ava Pettet, Yoshinori Suzue, Mark Bailey, Ayan Mukhopadhyay, and Abhishek Dubey. 2026. CONSENT: A Negotiation Framework for Leveraging User Flexibility in Vehicle-to-Building Charging under Uncertainty. *arXiv preprint arXiv:2601.01581* (2026).
- [15] Rishav Sen, Amutheezan Sivagnanam, Aaron Laszka, Ayan Mukhopadhyay, and Abhishek Dubey. 2024. Grid-Aware Charging and Operational Optimization for Mixed-Fleet Public Transit. In *2024 IEEE 27th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 4172–4179.
- [16] Rishav Sen, Yunuo Zhang, Fangqi Liu, Jose Paolo Talusan, Ava Pettet, Yoshinori Suzue, Ayan Mukhopadhyay, and Abhishek Dubey. 2025. Online Decision-Making Under Uncertainty for Vehicle-to-Building Systems. In *Proceedings of the ACM/IEEE 16th International Conference on Cyber-Physical Systems (with CPS-IoT Week 2025)*. 1–12.
- [17] Amutheezan Sivagnanam, Afiya Ayman, Michael Wilbur, Philip Pugliese, Abhishek Dubey, and Aron Laszka. 2021. Minimizing energy use of mixed-fleet public transit for fixed-route service. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 14930–14938.
- [18] Sebastian Stein, Enrico Gerding, Valentin Robu, and Nick Jennings. 2012. A model-based online mechanism with pre-commitment and its application to electric vehicle charging. (2012).
- [19] Jose Paolo Talusan, Rishav Sen, Ava Pettet, Aaron Kandel, Yoshinori Suzue, Liam Pedersen, Ayan Mukhopadhyay, and Abhishek Dubey. 2024. OPTIMUS: Discrete Event Simulator for Vehicle-to-Building Charging Optimization. In *2024 IEEE International Conference on Smart Computing (SMARTCOMP)*. IEEE, 223–230.
- [20] Felix Tuchnitz, Niklas Ebell, Jonas Schlund, and Marco Pruckner. 2021. Development and evaluation of a smart charging strategy for an electric vehicle fleet based on reinforcement learning. *Applied Energy* 285 (2021), 116382.
- [21] Gilles Van Kriekinge, Cedric De Cauwer, Nikolaos Sapountzoglou, Thierry Coosemans, and Maarten Messagie. 2021. Peak shaving and cost minimization using model predictive control for uni- and bi-directional charging of electric vehicles. *Energy reports* 7 (2021), 8760–8771.
- [22] José R Vázquez-Canteli and Zoltán Nagy. 2019. Reinforcement learning for demand response: A review of algorithms and modeling techniques. *Applied energy* 235 (2019), 1072–1089.