Honda Research Institute EU
Honda Research Institute Europe GmbH
https://www.honda-ri.de/

Quantum-Classical Framework for the Joint Optimization of Bidirectional EV Charging and Trip Assignment

Mustafa Musab, Daniel Müssig, Sebastian Schmitt, Markus Wappler, Jörg Lässig

2026

Preprint:

This is an accepted article published in GI Quantum Computing Workshop, Informatik Festival 2026. The final authenticated version is available online at: https://doi.org/[DOI not available]

Quantum-Classical Framework for the Joint Optimization of Bidirectional EV Charging and Trip Assignment

Mustafa Musab [1], Sebastian Schmitt [2], Markus Wappler[3], Jörg Lässig [3, 4], and Daniel Müssig [3]

Abstract: The integration of bidirectional (V2G) charging allows electric vehicle (EV) fleets to operate as distributed energy assets. However, the joint optimization of charging schedules and trip assignments is NP-hard. Moreover, bidirectional power flows create non-monotonic state-of-charge (SoC) dynamics that weaken standard dual-based Benders cuts. We propose a hybrid quantum-classical Benders decomposition framework. To overcome the qubit overhead associated with encoding feasibility cuts, we introduce logic-based conflict cuts, which require near-zero auxiliary qubits. Furthermore, we develop a sparse recursive SoC formulation that reduces the constraint-matrix density by a factor of $\sim T/3$ in the horizon length T. We demonstrate that our framework recovers MILP-optimal solutions on benchmark instances with a low auxiliary-qubit budget, and, as a minimal proof of concept, we execute a small instance end-to-end on IBM quantum hardware.

Keywords: Hybrid Quantum-Classical, Logic-Based Benders Decomposition, QAOA, V2G, MILP

1 Introduction and Literature Review

The accelerating adoption of electric vehicles is reshaping both the transportation and energy sectors. Global EV sales exceeded 17 million in 2024, accounting for more than 20% of all car sales worldwide, and projections indicate continued growth [In25]. This transition presents a dual opportunity, as EVs can contribute to reducing CO_2 emissions while also serving as flexible energy storage assets for power systems through bidirectional charging technologies such as Vehicle-to-Grid (V2G) and Vehicle-to-Home (V2H) [Ku25].

Driven by this commercial potential, a growing body of literature has explored different formulations of the EV problem, ranging from trip scheduling to bidirectional energy dispatch. Due to the inherent complexity of these models, decomposition-based methods have become the standard tool for tackling large-scale EV optimization. Benders decomposition, in particular, has been applied to EV charging infrastructure planning [AK16; GH+25], optimizing car-sharing station location [ÇF19], and V2V routing problems [Re25].

[1] Fraunhofer IOSB-AST , Am Vogelherd 90, 98693 Ilmenau, Germany, mustafa.musab@iosb-ast.fraunhofer.de, https://orcid.org/0009-0002-2964-6529
[2] Honda Research Institute Europe GmbH, Carl-Legien-Straße 30, 63073 Offenbach/Main, Germany, sebastian.schmitt@honda-ri.de, https://orcid.org/0000-0001-7130-5483
[3] Fraunhofer IOSB-AST, Wilhelmsplatz 11, 02826 Görlitz, Germany, markus.wappler@iosb-ast.fraunhofer.de; jlaessig@hszg.de, https://orcid.org/0000-0002-3169-3802; daniel.muessig@iosb-ast.fraunhofer.de, https://orcid.org/0000-0002-2345-1475
[4] University of applied Sciences Zittau/Görlitz, Wilhelmsplatz 11, 02826 Görlitz, Germany, jlaessig@hszg.de, https://orcid.org/0000-0002-3169-3802

Tab. 1: Classification of relevant literature by problem characteristics and computing paradigm.

Reference	Type	Trip assign.	Bidir. V2G	Paradigm	Decomposition
Veshchezerova et al. [Ve23]	EVCA	✓	–	Hybrid	Graph
Deller et al. [De23]	EVC	–	✓	Quantum	–
Varga et al. [VRL22]	EVSA	✓	–	Classical	Benders
Müssig et al. [Mü23]	EVC	–	✓	Quantum	–
Du et al. [Du25]	EVOP	✓	✓	Classical	–
Ekstrøm et al. [EWS26]	EVCA	✓	✓	Quantum	–
Our work	**EVCA**	✓	✓	**Hybrid**	**Benders**

Type: EVC = EV charging; EVCA = EV charging & trip assignment; EVSA = EV fleet scheduling & assignment; EVOP = EV charging & orienteering problem. **Indicators:** ✓ yes; – not applicable.

Quantum computing offers a different paradigm for the combinatorial aspect of this problem. For instance, quantum algorithms such as the Quantum Approximate Optimization Algorithm (QAOA) [FGG14] and quantum annealing [RS23] can be utilized to explore discrete solution spaces. Previous work from our group studied the bidirectional EV charging problem and successfully benchmarked QAOA on small instances [Mü23]. Recognizing the large qubit cost required for encoding integer variables, our subsequent research investigated d-level qudit systems to encode the EV charging problem more efficiently [De23; EWS26]. In the current work, we extend this foundation to the joint optimization of charging planning and trip scheduling for realistic large scale systems.

Despite the wide range of EV use cases studied in the literature, bidirectional charging remains comparatively underexplored. While bidirectional V2G may appear to be a straightforward extension of the unidirectional case, it fundamentally alters the structure of the feasible region and introduces non-monotonic SoC behavior [5] that complicates the optimization process. To the best of our knowledge, this is the first hybrid quantum-classical framework applied to the joint optimization of bidirectional EV charging and fleet scheduling. Table 1 outlines the studies most closely related to our study.

This paper makes the following contributions:

1. **Logic-Based Hybrid Framework:** We propose a hybrid architecture for bidirectional EV charging and trip assignment using logic-based conflict cuts that admit a near-zero-qubit QUBO encoding.
2. **Recursive SoC Formulation:** We formulate the SoC constraints in a recursive way, resulting in sparse constraint-matrix density.

[5] In unidirectional charging, the SoC monotonically increases during idle periods, allowing solvers to simplify the model by evaluating battery limits solely at trip boundaries. Bidirectional discharging breaks this assumption, requiring dense constraints at every timestep.

The remainder of the paper is structured as follows. Section 2 describes the scenario considered and the mathematical formulation of the problem. Section 3 details the development of our hybrid approach. Finally, Section 4 presents the experimental results and concludes.

2 Problem Description

We consider a bidirectional EV charging and trip assignment scenario where an EV fleet operator allocates vehicles to customer requests (trips) while optimizing bidirectional energy dispatch during idle periods. Table 2 summarizes the key parameters and notations used in our mathematical formulation.

Tab. 2: Notation and Problem Parameters

Symbol	Description	Symbol	Description
N_{EV}	Number of EVs, $n \in \{1, \ldots, N_{\text{EV}}\}$	E_n^0	Initial SoC of EV n (at $t = 0$)
T	Number of time steps, $t \in \{1, \ldots, T\}$	E_n^{cap}	Battery capacity of EV n
R	Number of trip requests, $i \in \{1, \ldots, R\}$	$E_n^{\min}$	Required final SoC of EV n (at $t = T$)
Δt	Time step duration (hours)	E_i^{trip}	Energy required for serving trip i
c_t	Electricity price at time t	$\mathcal{T}_i = [t_i^{\text{start}}, t_i^{\text{end}}]$	Time window of trip i
$l_{\max}$	Maximum charging/discharging power	$P^{\min}, P^{\max}$	Grid fuse limits on aggregate power
λ_i	Reward for serving trip i	$O = \{(i, j) : \mathcal{T}_i \cap \mathcal{T}_j \neq \emptyset\}$	Overlapping trip pairs

To ensure valid operations, the system must adhere to strict physical and logical requirements:

- Each trip can be served by at most one EV, and overlapping trips cannot be assigned to the same vehicle.
- An EV cannot be charged or discharged while actively serving a trip.
- The total charging and discharging power of the fleet at any time t is bounded by the grid fuse constraints.

Given these settings, the objective is to maximize the number of trips served while minimizing charging costs, subject to battery capacities, grid limits, and assignment constraints. The first term in (1a) represents the charging cost, where $l_{n,t}$ is the charging power of EV n at time t. Negative levels ($l_{n,t} < 0$) indicate discharge (V2G) while positive levels ($l_{n,t} > 0$) indicate charging, and ($l_{n,t} = 0$) mean that no charging or discharge occurs at this time for that EV. The charging cost c_t generally depends on the charging power, in particular, with different prices for selling ($l < 0$) and buying ($l > 0$) energy, but we typically use the same price for both cases, i.e., c_t is independent of the power level. The second term of Eq. (1a) rewards the trips served by setting the reward parameter $\lambda_i \geq c_{\max} \cdot E_i^{\text{trip}}$. This ensures that any feasible trip is economically preferable to leaving it unassigned, thereby implementing a "serve-if-feasible" logic. Constraint (1b) ensures that each trip is served by at most one EV. Constraint (1c) prevents an EV from serving two trips with overlapping time windows. The constraint (1d) enforces the logic that an EV cannot charge while serving. The SoC of EV must always be within the allowable range $[0, E_n^{\text{cap}}]$. This is enforced by Eqs. (1e) and

$$\min \quad \Delta t \sum_{t=1}^{T} \sum_{n=1}^{N_{\mathrm{EV}}} c_t \cdot l_{n,t} - \sum_{i=1}^{R} \sum_{n=1}^{N_{\mathrm{EV}}} \lambda_i \cdot r_{n,i} \tag{1a}$$

$$\text{s.t.} \quad \sum_{n} r_{n,i} \leq 1 \qquad \forall i \tag{1b}$$

$$r_{n,i} + r_{n,j} \leq 1, \qquad \forall n,\ (i,j) \in \mathcal{O} \tag{1c}$$

$$-l_{\max}(1 - r_{n,i}) \leq l_{n,t} \leq l_{\max}(1 - r_{n,i}) \qquad \forall n,\ i,\ t \in \mathcal{T}_i \tag{1d}$$

$$E_n^0 + \Delta t \sum_{k \leq t} l_{n,k} - \sum_{i:\, t_i^{\mathrm{start}} \leq t} E_i^{\mathrm{trip}} \cdot r_{n,i} \geq 0, \qquad \forall n, t \tag{1e}$$

$$E_n^0 + \Delta t \sum_{k \leq t} l_{n,k} - \sum_{i:\, t_i^{\mathrm{start}} \leq t} E_i^{\mathrm{trip}} \cdot r_{n,i} \leq E_n^{\mathrm{cap}}, \qquad \forall n, t \tag{1f}$$

$$E_n^0 + \Delta t \sum_{t} l_{n,t} - \sum_{i} E_i^{\mathrm{trip}} \cdot r_{n,i} \geq E_n^{\min}, \qquad \forall n \tag{1g}$$

$$P^{\min} \leq \sum_{n} l_{n,t} \leq P^{\max}, \qquad \forall t \tag{1h}$$

$$r_{n,i} \in \{0,1\}, \quad l_{n,t} \in [-l_{\max}, l_{\max}].$$

(1f). The final SoC must be greater than $E_n^{\min}$ (1g). Constraint (1h) ensures that the fuse limits of the grid connection are not exceed.

The total count of SoC and grid inequalities in (1e), (1f), (1g), and (1h) scales as $2N_{\mathrm{EV}}T + N_{\mathrm{EV}} + 2T$. In a standard QUBO formulation, each inequality is typically converted to an equality by introducing a slack variable. Each slack must then be discretized using binary expansion, requiring a substantial number of qubits. This overhead alone can easily exceed the practical capacity of near-term quantum hardware. To illustrate, even for a minimal instance with ($N_{\mathrm{EV}} = 2,\ T = 24$), i.e., scheduling two EVs over a single day, the model contains 146 inequalities. This motivates a decomposition-based strategy, which we develop in the following sections.

3 Quantum-Classical Benders Decomposition

In this section, we detail the proposed hybrid framework. Our approach is built on two foundational pillars. First, we develop a recursive SoC formulation in Sect. 3.2. Then, we derive logic-based conflict cuts in Sect. 3.4.

3.1 Classical Benders Decomposition

Benders decomposition is a classical technique for solving large-scale mixed-integer linear programs (MILPs). The key idea is to identify a set of complicating variables such that,

once fixed, the remaining problem reduces to a linear program that is considerably easier to solve. The optimization then proceeds iteratively: a *master problem* (MP) explores the discrete assignment space, while a continuous *subproblem* (SP) evaluates the feasibility and cost of those assignments, generating cutting planes that are fed back to refine the MP [Be05; Ge72]. In the context of our bidirectional EV scheduling model, we allocate the binary trip-assignment variables $r_{n,i}$ to the MP, while charging schedules $l_{n,t}$ and SoC dynamics are handled in the linear SP.

3.2 Recursive SoC Formulation

The SoC bounds are the densest part of the model. Each bound in (1e)–(1f) is an inequality row that references the entire charging history up to t and every trip already departed by time t. We quantify the constraints density by counting non-zero elements (nnz).

For EV n, the SoC bounds at time t carries t nnz from the charging variables. Summing this over the horizon $t = 1 \dots T$ yields $T(T+1)/2$ nnz. Similarly, a trip i is present from its departure onward, i.e. in $(T - t_i^{\text{start}} + 1)$ rows per bound. Assuming trip start times are uniformly distributed across the horizon, a trip appears in $\mathbb{E}\left[T - t_i^{\text{start}} + 1\right] = (T+1)/2$ equations on average. Across all R trips, this contributes $R(T+1)/2$ nnz.

Finally, the end-of-horizon SoC requirement (1g) adds one row with T charging nnz and R assignment nnz. Summing these contributions across all N_{EV} vehicles, and accounting for both SoC bounds (upper and lower), the total nnz count is:

$$\mathbb{E}\left[\text{nnz}_{\text{cum}}^{\text{SoC+fin}}\right] = N_{\text{EV}}\Big[\underbrace{T(T+1) + R(T+1)}_{\text{both SoC bounds}} + \underbrace{(T+R)}_{\text{final SoC}}\Big] = N_{\text{EV}}\,(T+R)(T+2). \quad (2)$$

We introduce an explicit SoC variable $E_{n,t}$ and rewrite the battery dynamics recursively:

$$E_{n,t} = E_{n,t-1} + \Delta t \cdot l_{n,t} - \sum_{i:\, t_i^{\text{start}} = t} E_i^{\text{trip}} \cdot r_{n,i}, \qquad \forall n, t, \quad (3)$$

with the initial condition $E_{n,0} = E_n^0$. Observe that the dense upper and lower SoC inequalities reduce to simple variable bounds: $0 \le E_{n,t} \le E_n^{\text{cap}}$. Additionally, the end-of-horizon requirement (1g) is enforced as $E_{n,T} \ge E_n^{\min}$. This results on reduced nnz count:

$$\text{nnz}_{\text{rec}}^{\text{SoC+fin}} = N_{\text{EV}}(3T + R + 1), \quad (4)$$

which is exact for every instance. The resulting density reduction is

$$\frac{\text{nnz}_{\text{cum}}^{\text{SoC+fin}}}{\text{nnz}_{\text{rec}}^{\text{SoC+fin}}} = \frac{(T+R)(T+2)}{3T+R+1} \xrightarrow{T \gg R} \frac{T}{3}, \quad (5)$$

i.e., the density of the SoC block drops from $O(T^2)$ to $O(T)$ per vehicle.

Tab. 3: Comparison of the cumulative and recursive SoC formulations. *nnz*: constraint-matrix nonzero count; *Speedup*: cumulative-to-recursive ratio; *OOM*: Gurobi ran out of memory during model build.

	Cumulative		**Recursive**		**Speedup**	
Instance	**nnz**	**time**	**nnz**	**time**	**nnz**	**time**
(32, 1, 4)	1K	0.01s	236	0.01s	5.8×	1×
(192, 5, 20)	220K	2.4m	15K	15.1s	14.4×	9.6×
(768, 20, 80)	14M	*OOM*	725K	20.0m	19.1×	*n/a*

This simple reformulation has significant impact on the optimizer performance. Table 3 compares performance on the dataset of Ref. [VRL22], reporting the total model nnz (across all constraints) to highlight the overall sparsity gain. Observe that for the large instance (T=768, N_{EV}=20, R=80), Gurobi exhausted the memory (OOM) building the cumulative model. By contrast, the recursive formulation remains highly tractable. We adopt it hereafter for both the Benders subproblem and direct-MILP baseline.

3.3 Hybrid Decomposition Architecture

Our goal is to decouple the NP-hard combinatorial trip assignment from the continuous energy dispatch. To achieve this, we formulate MP tailored for quantum execution. By rewriting constraints (1b) and (1c) in quadratic form, we eliminate the auxiliary qubits otherwise required to convert inequalities into equalities. All continuous variables and dispatch-side constraints are handled within the classical SP. The MP at iteration k therefore retains only the binary variables $r_{n,i}$ and a continuous cost surrogate θ:

$$
\begin{aligned}
\min_{r,\theta} \quad & \theta - \sum_{i=1}^{R} \sum_{n=1}^{N_{\mathrm{EV}}} \lambda_i \, r_{n,i} && && (\mathrm{MP}^{(k)}) \\
\text{s.t.} \quad & \sum_{n \neq m} r_{n,i} \cdot r_{m,i} = 0 && \forall i && (6a) \\
& r_{n,i} \cdot r_{n,j} = 0 && \forall n,\ (i,j) \in O && (6b) \\
& \theta \geq \pi_0^{(k')} + \sum_{n,i} \pi_{n,i}^{(k')} \, r_{n,i} && k' \in \mathcal{K}_{\mathrm{opt}} && (6c) \\
& \sum_{i \in C_n^{(k')}} r_{n,i} \ \leq \ |C_n^{(k')}| - 1 && k' \in \mathcal{K}_{\mathrm{cnf}} && (6d) \\
& r_{n,i} \in \{0,1\}, \quad \theta \in \mathbb{R}, &&
\end{aligned}
$$

where $\mathcal{K}_{\mathrm{opt}}$ and $\mathcal{K}_{\mathrm{cnf}}$ are the sets containing the optimality and logic-based cuts, respectively, accumulated up to iteration k.

For a candidate assignment r^* returned by the MP, the SP is the linear dispatch problem:

$$\min_{l} \quad \Delta t \sum_{t=1}^{T} \sum_{n=1}^{N_{\mathrm{EV}}} c_t \, l_{n,t} \qquad (\mathrm{SP}(r^*))$$

$$\text{s.t.} \quad E_{n,t} - E_{n,t-1} - \Delta t \, l_{n,t} = - \sum_{i:\, t_i^{\mathrm{start}} = t} E_i^{\mathrm{trip}} \, r^*_{n,i} \qquad \forall n \; t \qquad (7a)$$

$$0 \le E_{n,t} \le E_n^{\mathrm{cap}} \qquad \forall n, \; t \qquad (7b)$$

$$E_{n,T} \ge E_n^{\mathrm{min}} \qquad \forall n \qquad (7c)$$

$$-\, l_{\max}(1 - s^*_{n,t}) \le l_{n,t} \le l_{\max}(1 - s^*_{n,t}) \qquad \forall n, \; t \qquad (7d)$$

$$P_{\min} \le \sum_{n} l_{n,t} \le P_{\max}, \qquad \forall t \qquad (7e)$$

where $s^*_{n,t} = \min(1, \sum_{i\,:\,t \in \mathcal{T}_i} r^*_{n,i})$ indicates whether vehicle n is serving a trip at time t. The interaction between these two problems is governed by Benders cuts. When $\mathrm{SP}(r^*)$ is feasible, its dual variables yield an optimality cut (6c). When $\mathrm{SP}(r^*)$ is infeasible, it returns feasibility cut to prevent the MP from selecting r^* again. The derivation of these cuts and the associated qubit-encoding cost are the subject of the next subsection.

3.4 Logic-based feasibility cuts

For some candidates r^*, there is no feasible l, and the above SP is infeasible. In standard Benders, this is handled by extracting a Farkas certificate from the dual ray of the infeasible SP and deriving a feasibility cut. However, in our V2G settings, this dual-based cut is structurally weak and exhausts many iterations. We therefore adopt logic-based cuts [HO03].

First, we assume that the grid power (7e) is sufficiently provisioned: $P_{\max} \ge N_{\mathrm{EV}} \cdot l_{\max}$ (e.g., this can be enforced by limiting the number of EVs entering the scheduling phase). Under this assumption, the SP is completely decoupled per EV. Therefore, an infeasible solution r^* implies that at least one EV cannot commit to the assigned subset of trips $\bar{A}_n = \{i : r^*_{n,i} = 1\}$. Our task is then to identify an irreducible inconsistent subsystem (IIS) $C_n \subseteq \bar{A}_n$ that excludes this subset (and any superset of C_n) from the solution space. We solve N_{EV} independent LPs, and for each infeasible LP we apply a deletion filter to extract the IIS C_n. Starting from $\bar{A}_n$, each trip is temporarily removed and per-EV LP is re-solved. If the remaining subset is still infeasible even without serving that trip, the trip is deemed non-essential to the conflict and is permanently dropped. Conversely, if removing the trip restores feasibility, the trip is essential and is retained. The resulting set C_n is minimal in the sense that no proper subset of it is infeasible, and yields the combinatorial cut:

$$\sum_{i \in C_n} r_{n,i} \le |C_n| - 1$$

Algorithm 1 Hybrid Quantum-Classical Benders Decomposition

Initialise MP with structural constraints (6a)-(6b) and $\bar{z} \leftarrow +\infty$, $\underline{z} \leftarrow -\infty$
Add pre-seed optimality cut from SP($\bar{r} = 0$)
repeat
    Solve MP on quantum hardware $\rightarrow$ candidate $\bar{r}$, lower bound $\underline{z}$
    Solve SP($\bar{r}$) classically
    if SP feasible **then**
        $z_{\text{cand}} \leftarrow \text{SP}(\bar{r}) - \sum_{n,i} \beta_i \bar{r}_{n,i}$; if $z_{\text{cand}} < \bar{z}$, set $\bar{z} \leftarrow z_{\text{cand}}$ and store incumbent
        Add optimality cut (6c) from SP duals
    else
        for each infeasible vehicle n **do**
            $C_n \leftarrow \text{IIS-EXTRACT}(\bar{A}_n)$ ▷ $\bar{A}_n = \{i : \bar{r}_{n,i} = 1\}$; deletion filter
            Add conflict cut (6d)
        end for
    end if
until $(\bar{z} - \underline{z})/|\bar{z}| < \varepsilon$ or iteration limit
return $(r^\star, l^\star, E^\star)$

Empirically, we have observed that 70–85% of extracted conflicts have a size $|C_n| \leq 2$. This means they reduce to either single-variable fixings ($r_{n,i} = 0$) or pairwise exclusions ($r_{n,i} + r_{n,j} \leq 1$). Note that we can express these pairwise cuts in a quadratic form ($r_{n,i} \cdot r_{n,j} = 0$). This has significant implications for QUBO encoding: the majority of the conflict cuts do not require auxiliary slack qubits or binary expansions to encode. We revisit this in Sec. 4 and report per-instance auxiliary-qubit budgets. The IIS extraction procedure is illustrated in Fig. 1, and the hybrid framework is given in Algorithm 1.

3.5 MP Quantum Formulation

We solve the MP using the QAOA algorithm [FGG14]. For this, we first transform the MP to its quantum formulation by selecting a computation basis and replacing the search variables with the corresponding quantum operators. Then we incorporate all constraints as penalty terms. The computational basis states are indexed by the qubit variables $r_{n,i} \in \{0, 1\}$. The corresponding z-angular momentum operators are defined as:

$$
\begin{aligned}
R^z_{n,i} \left| r_{n,i} \right\rangle &= r_{n,i} \left| r_{n,i} \right\rangle, \\
R^+_{n,i} \left| r_{n,i} \right\rangle &= \sqrt{(1 - r_{n,i})(1 + r_{n,i})} \left| r_{n,i} + 1 \right\rangle, \\
R^-_{n,i} \left| r_{n,i} \right\rangle &= \sqrt{r_{n,i}, (2 - r_{n,i})} \left| r_{n,i} - 1 \right\rangle.
\end{aligned}
$$

The x and y operators are defined as $R^x_{n,i} = \frac{1}{2}(R^+_{n,i} + R^-_{n,i})$ and $R^y_{n,i} = \frac{1}{2i}(R^+_{n,i} - R^-_{n,i})$. The assignment register spans a Hilbert space of dimension $2^{N_{\text{EV}} R}$. The encoding additionally

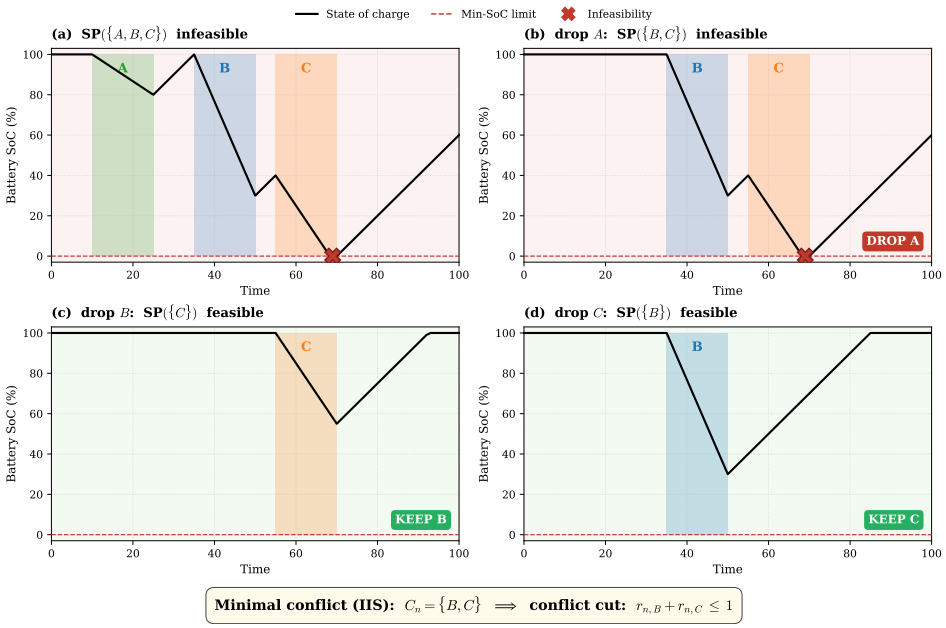

Fig. 1: IIS extraction illustrated on a toy instance with three trips A, B, and C assigned to a single EV. Panel (a) shows the candidate assignment $\bar{A}_n = \{A, B, C\}$ exhausting the battery during trip C. Panel (b) demonstrates that removing A leaves the subset $\{B, C\}$ infeasible, identifying A as non-essential to the conflict. Panels (c) and (d) verify that the remaining set $\{B, C\}$ is irreducible, as removing either element restores feasibility. The resulting minimal conflict $C_n = \{B, C\}$ generates the combinatorial cut $r_{n,B} + r_{n,C} \leq 1$, strictly excluding any master assignment that allocates vehicle n a superset of $\{B, C\}$. If $\bar{A}_n$ contains multiple independent conflicts, the deletion filter extracts one minimal IIS per Benders iteration, systematically eliminating the remaining conflicts in subsequent iterations.

allocates a binary-expanded register for the surrogate variable Θ. The MP cost Hamiltonian at iteration k is

$$
\begin{aligned}
H_{\mathrm{MP}}^{(k)} \;=\; \Theta \;&-\; \sum_{i=1}^{R} \sum_{n=1}^{N_{\mathrm{EV}}} \lambda_i \, R_{n,i}^{z} && \text{(trip rewards)} \\
&+\; P \sum_{i=1}^{R} \sum_{n<m} R_{n,i}^{z} \, R_{m,i}^{z} && \text{(at most one EV per trip)} \\
&+\; P \sum_{n=1}^{N_{\mathrm{EV}}} \sum_{(i,j)\in O} R_{n,i}^{z} \, R_{n,j}^{z} && \text{(no overlapping trips)} \\
&+\; P \sum_{k' \in \mathcal{K}_{\mathrm{opt}}} H_{\mathrm{opt}}^{(k')} \;+\; P \sum_{k' \in \mathcal{K}_{\mathrm{cnf}}} H_{\mathrm{cnf}}^{(k')}, && (8)
\end{aligned}
$$

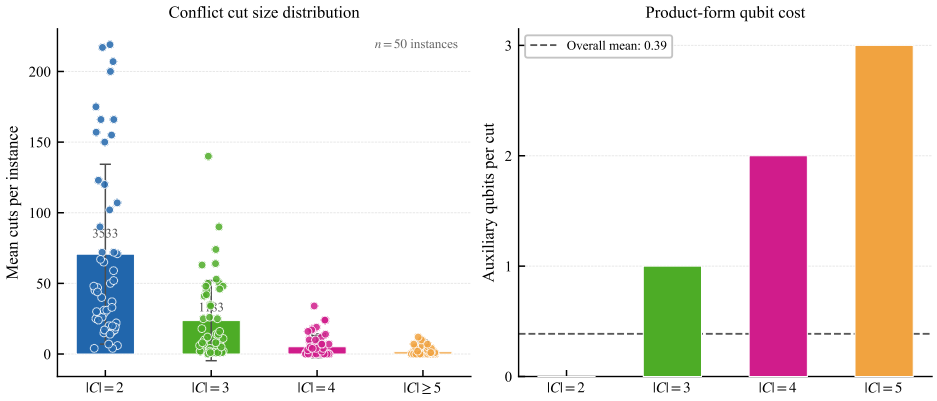

Fig. 2: Distribution of logic-based conflict cut sizes (left) and their associated QUBO encoding costs (right) across 50 medium instances.

where P is a penalty factor. Each optimality cut (6c) is converted to an equality by introducing a non-negative slack variable $S_{\text{opt}}^{(k')}$:

$$H_{\text{opt}}^{(k')} = \left(\Theta - \pi_0^{(k')} - \sum_{n,i} \pi_{n,i}^{(k')} R_{n,i}^z - S_{\text{opt}}^{(k')}\right)^2. \tag{9}$$

Similarly, conflict cuts (6d) are enforced independently for each EV:

$$H_{\text{cnf}}^{(k')} = \left(\sum_{i \in C_n^{(k')}} R_{n,i}^z + S_{\text{cnf}}^{(k')} - |C_n^{(k')}| + 1\right)^2, \tag{10}$$

with integer slack $S_{\text{cnf}}^{(k')} \in \{0, \ldots, |C_n^{(k')}| - 1\}$. For the dominant case $|C_n^{(k')}| = 2$ (Sect. 3.4) this collapses to the slack-free pairwise penalty $P\, R_{n,i}^z\, R_{n,j}^z$, requiring no auxiliary qubits.

4 Experimental Results

For benchmarking, we use the dataset of Ref. [VRL22], which provides 810 instances of varying sizes and configurations. We adapt the dataset to our bidirectional use case by replacing the non-negativity bound on $l_{n,t}$ with a symmetric range $-l_{\max} \le l_{n,t} \le l_{\max}$. The remaining parameters are set as follows:

- Trip-service reward: $\lambda_i = 2\, c_{\max}\, E_i^{\text{trip}}$
- Final-SoC requirement: $E_n^{\min} = 0.1 \cdot E_n^{\text{cap}}$
- Fuse limits: $P_{\max} = -P_{\min} = N_{\text{EV}} \cdot l_{\max}$.

Tab. 4: Per-instance results on medium instances ($T = 192$, $N = 5$, $R = 20$). Benders convergence gap $(UB - LB)/|UB| \times 100$ at termination, and Relat. gap = $(z^*_{\mathrm{B}} - z^*_{\mathrm{dir}})/|z^*_{\mathrm{dir}}| \times 100$.

	Direct MILP		Hybrid Benders						
Inst	Obj.	Time	Obj.	Time	Iters	Conf. cuts	Aux qubits	Bdrs gap (%)	Relat. gap (%)
01	-7,491.01	62	-7,487.82	101	31	22	1	1.00	0.04
02	-9,506.15	125	-9,503.36	75	33	64	14	0.94	0.03
03	-9,155.98	324	-9,149.74	600	42	32	6	0.99	0.07
04	-9,536.58	196	-9,529.30	83	35	38	1	0.96	0.08
05	-7,867.73	76	-7,864.21	50	29	24	5	0.98	0.04
Mean	-8,711.49	156.6	-8,706.88	181.9	34	36	5.4	0.97	0.05

The classical computations were performed on a Dell laptop with an Intel Core i7-1185G7 processor (4 cores, 8 logical threads) and 16 GB RAM.

First, we demonstrate the effectiveness of our logic-based cuts. Fig. 2 presents the distribution of conflict cut sizes across 50 medium instances. The left panel shows that the majority of cuts are size $|C_n| = 2$, while cuts with $|C_n| \geq 5$ are rare and cluster near zero. We also observed cuts of size $|C_n| = 1$. These arise when a single trip is infeasible on its own (e.g., when a trip starts so early in the horizon that the vehicle's initial SoC, E^0_n, cannot be raised to the required E^{trip}_i before t^{start}_i). The right panel illustrates the auxiliary qubit cost required for each cut size. Crucially, size-2 cuts require zero additional qubits for their QUBO encoding. In Table 4, we report the optimization results of 5 medium-sized instances. Note, that these instances are already far too large for pure quantum approaches [EWS26]. Overall, the hybrid approach terminates at a Benders convergence tolerance of $\leq 1\%$ on all instances, and successfully recovers the MILP optimum with a mean relative gap of 0.05%. For some instances, the computational time of the hybrid decomposition is larger than that of the direct solution, an expected result, since Benders advantages become more pronounced for large-scale problems. Interestingly, the qubit overhead for encoding conflict cuts is rather small. The largest count is 14 which is needed to encode 64 total cuts. Instances 1 and 4 require only 1 qubit for encoding their total cuts, which implies only 1 cut has size > 2.

As a proof of concept, we also executed a small-scale instance ($T = 3, N_{\mathrm{EV}} = 2, R = 2$) on real IBM quantum hardware. At each Benders iteration, the MP is solved on the IBM backend using QAOA. From the sampled bitstring distribution, we extract the top-k most frequent candidates and evaluate them in batch using the SP feasibility checks. For this reduced instance, the method converges within a single Benders iteration and returns the same optimal solution as the classical baseline. Table 5 summarizes this convergence against the Gurobi MILP solver.

Method	Iteration	Lower Bound	Upper Bound	Gap	Cuts Added
Hybrid Benders	0	−10000	10000	20000	4
	1	−350	−350	0	–
Gurobi MILP	–	−350	−350	0	–

Tab. 5: Comparison of the hybrid Benders approach (via QAOA) and the classical MILP baseline on a small-scale instance ($T = 3, N_{\mathrm{EV}} = 2, R = 2$).

5 Conclusion

In this paper, we have presented a hybrid quantum-classical framework for the joint optimization of large-scale bidirectional EV charging and trip assignment problems. By adopting a recursive SoC formulation, we improved the sparsity of the constraint matrix and consequently reduced memory and time requirements. Additionally, the introduction of conflict cuts to handle Benders feasibility cuts resulted in an efficient MP formulation with near-zero-qubit encoding overhead. The framework recovers MILP-optimal objectives on the benchmark instances within a small auxiliary-qubit budget, with end-to-end execution validated on IBM quantum hardware. Although our conflict cuts contribute significantly to the overall reduction of the qubit encoding cost, the total qubit budget remains dominated by Benders optimality cuts. Moving forward, we plan to address these optimality cuts and investigate cut-strengthening techniques.

Acknowledgements

Funded by the European Union under Horizon Europe Programme - Grant Agreement 101080086 — NeQST.
Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Climate, Infrastructure and Environment Executive Agency (CINEA). Neither the European Union nor the granting authority can be held responsible for them.

References

[AK16] Arslan, O.; Karaşan, O. E.: A Benders decomposition approach for the charging station location problem with plug-in hybrid electric vehicles. Transp. Res. B Methodol. 93, pp. 670–695, 2016.

[Be05] Benders, J. F.: Partitioning procedures for solving mixed-variables programming problems. Comput. Manag. Sci. 2 (1), Reprinted from Numerische Mathematik, 1962, pp. 3–19, 2005.

[ÇF19] Çalık, H.; Fortz, B.: A Benders decomposition method for locating stations in a one-way electric car sharing system under demand uncertainty. Transportation Research Part B: Methodological 125, pp. 121–150, 2019.

[De23] Deller, Y. et al.: Quantum approximate optimization algorithm for qudit systems. Physical Review A 107 (6), p. 062410, 2023.

[Du25] Du, J. et al.: Smart ride and delivery services with electric vehicles: Leveraging bidirectional charging for profit optimisation. Information Sciences, p. 122929, 2025.

[EWS26] Ekstrøm, L.; Wang, H.; Schmitt, S.: Comparing Qubit and Qudit Encodings for EV Charging and Trip Assignment Problems. arXiv:2605.10255, Accepted at GECCO 2026, 2026.

[FGG14] Farhi, E.; Goldstone, J.; Gutmann, S.: A quantum approximate optimization algorithm, 2014, arXiv: 1411.4028 [quant-ph].

[Ge72] Geoffrion, A. M.: Generalized benders decomposition. Journal of optimization theory and applications 10 (4), pp. 237–260, 1972.

[GH+25] Golsefidi, A. H.; Hipolito, F., et al.: Incremental expansion of large scale fixed and mobile charging infrastructure in stochastic environments: A novel graph-based Benders decomposition approach. Appl. Energy 380, p. 124985, 2025.

[HO03] Hooker, J. N.; Ottosson, G.: Logic-based Benders decomposition. Math. Program. 96 (1), pp. 33–60, 2003.

[In25] International Energy Agency: Global EV Outlook 2025, tech. rep., Available: https://www.iea.org/reports/global-ev-outlook-2025, IEA, 2025.

[Ku25] Kumar, P. et al.: A comprehensive review of vehicle-to-grid integration in electric vehicles: Powering the future. Energy Conversion and Management: X 25, p. 100864, 2025.

[Mü23] Müssig, D. et al.: Solving the Optimization Problem of Bidirectional EV-Charging on a Quantum Computer. Harzer Hochschultexte, 14th ser. NWK, pp. 94–101, 2023.

[Re25] Ren, K. et al.: The electric vehicle routing problem of a new mobile charging service. Transportation Safety and Environment 7 (1), tdae012, 2025.

[RS23] Rao, P. U.; Sodhi, B.: Hybrid quantum-classical solution for electric vehicle charger placement problem. Soft Comput. 27 (18), pp. 13003–13016, 2023.

[Ve23] Veshchezerova, M. et al.: A hybrid quantum-classical approach to the electric mobility problem. In: 2023 IEEE International Conference on Quantum Computing and Engineering (QCE). Vol. 1, IEEE, pp. 636–641, 2023.

[VRL22] Varga, J.; Raidl, G. R.; Limmer, S.: Computational methods for scheduling the charging and assignment of an on-site shared electric vehicle fleet. IEEE Access 10, pp. 105786–105806, 2022.