Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Scalability of Centralized and Decentralized Priority-Polling Coordination for Urban Electric Vehicle Charging: An Agent-Based SUMO Study of Muscat, Oman

Ramalingam Dharmalingam

ram@majancollege.edu.om

Majan College

Vaishnavi Dharmalingam

SASTRA University

Mahalakshmi Dharmalingam

A.V.C. College of Engineering

Research Article

Keywords:

Posted Date: November 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11794117/v1

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

Additional Declarations: No competing interests reported.

Scalability of Centralized and Decentralized Priority-Polling Coordination for Urban Electric Vehicle Charging: An Agent-Based SUMO Study of Muscat, Oman

Ramalingam Dharmalingam[1*], Vaishnavi Dharmalingam[2] and Mahalakshmi Dharmalingam[3]

[1*]Faculty of Information Technology, Majan University College, Muscat, Sultanate of Oman.
[2]Department of Computer Science and Engineering, Srinivasa Ramanujan Center, SASTRA Deemed to be University, Kumbakonam, 612001, Tamil Nadu, India.
[3]Department of Information Technology, A.V.C. College of Engineering, Mannampandal, Mayiladuthurai, 609305, Tamil Nadu, India.

*Corresponding author(s). E-mail(s): ram@majancollege.edu.om;
Contributing authors: vaishnavi@src.sastra.edu;
dmahaa1985@gmail.com;

Abstract

Coordinating where electric vehicles (EVs) charge determines whether an urban network absorbs demand or saturates. This study evaluates a Priority-Based Polling Agent-Based Model (PBP-ABM) in which EV agents poll charging-port agents and select a port by a composite score of travel time, expected queueing, service time and lane-change effort. Two centralized policies, in which a global controller holds the occupancy registry of every port, are compared with one decentralized policy, in which each EV observes only the vehicles physically present at a station, adds an anti-herding perturbation and re-polls after a waiting threshold. The policies were implemented in the Simulation of Urban MObility (SUMO) platform on an OpenStreetMap road network of Muscat, Oman, with two charging sites and six ports of three power classes. Fleet size was swept from 11 to 500 EVs at 60 and 120 EVs per hour, with five seeds per configuration (150 runs). Below capacity, all policies served every vehicle with flat access time. Above capacity, the centralized policies degraded sharply: mean

time to charging start rose from 115 s at 11 EVs to about 460 s at 500 EVs, with large seed-to-seed variance and a busiest-port occupancy of 85%. The decentralized policy held time to charging start at 180 s and peak occupancy at 75%, while leaving 3.6% of vehicles unserved. The divergence arises because a registry that counts only connected vehicles funnels late arrivals to slow ports, whereas local observation with re-polling redirects them to fast ports; architecture therefore matters only near saturation.

Keywords: electric vehicle charging, agent-based modelling, decentralized control, scalability, SUMO

1 Introduction

1.1 Background

With growing concerns about the environment and energy efficiency, the popularity of electric road transport has risen from niche to mainstream. Battery electric vehicles, plug-in hybrids and hybrid vehicles have the same characteristics in terms of charging pattern and infrastructure requirement, but they all require a network of charging points connected to the distribution grid. As distributed energy resources (DERs) which connect and disconnect dynamically at many grid nodes, EV batteries can cause power-quality and stability issues if they are not coordinated [1]. Placement and capacity allocation and scheduling of EV charging stations (EVCS) is therefore located in the intersection of transportation engineering, power systems and intelligent-agent design. In an urban setting, the difficulty is exacerbated by the aggregate driver behaviour, charging hotspot congestion, spatial heterogeneity in grid availability and cost trade-offs that influence the decisions of drivers. Agent-based modelling (ABM) is a suitable approach to model such situations, as it captures the behaviour of individual users of EVs, as well as the emergent load they impose on shared infrastructure: behaviour that is heterogeneous, autonomous and adaptive.

1.2 Problem statement

Charging stations need to be not only built, but used efficiently, as EV adoption reflects an increasing trend. Uncoordinated station choice is a problem that leads to a focus of demand at certain ports while others are unused, queues are longer and operating costs increase. The design question is a coordination algorithm that considers port capacity, queueing and varying demand and that continues to function as the fleet expands. Existing approaches address parts of this problem with reinforcement learning [2], multi-agent systems [3] or heuristic optimization [4], but comparisons of centralized and decentralized control are usually reported at a single fleet size, which hides the regime in which the two architectures actually diverge. Whether an architecture that performs well at low load continues to do so as the charging system approaches saturation is an empirical question that requires evaluation across scales.

1.3 Aim and objectives

This study proposes a Priority-Based Polling Agent-Based Model (PBP-ABM) for EV charging-port selection and evaluates centralized and decentralized coordination policies across fleet sizes spanning almost two orders of magnitude on a real urban road network. The specific objectives are to:

- formalise charging-port selection as a polling-based, multi-criteria decision in which EV agents and port agents exchange state;
- implement centralized and decentralized policies with identical scoring rules so that only the information architecture differs;
- evaluate the policies in the Simulation of Urban MObility (SUMO) platform with the Traffic Control Interface (TraCI) on a Muscat road network under arrival rates below and above charging capacity;
- sweep fleet size from 11 to 500 EVs with replicated random seeds and report service rate, time to charging start, traffic waiting time and per-port occupancy;
- identify the mechanism by which the architectures diverge and quantify the trade-off between service coverage and access time.

1.4 Scope

This is a simulation-based study. It takes into account a small sub-network of Muscat (3.6 km) with two charging stations, thus observing the saturation of the charging system, not the road network. Grid interaction, tariff dynamics and battery degradation are outside its scope. The EV agents are all identical, with identical vehicle dynamics and battery parameters, while heterogeneity is introduced in the origin–destination pairs, the departure times and the stochastic evolution of traffic.

2 Literature Review

2.1 EV charging and urban congestion

The placement of EVCS in urban areas has attracted considerable attention. Celik et al. [4] proposed an integrated model using genetic algorithms to minimise travel distance, waiting time and system cost. Nguyen et al. [2] combined agent-based simulation with reinforcement learning for station placement and reported a 53.28% reduction in mean waiting time. Khan et al. [3] abstracted individual EV and station behaviour as a multi-agent system to capture the stochastic interactions characteristic of charging environments. Demand-forecasting work includes Bellizio et al. [5] on public charge-point flexibility, Sanami and Mosalli [6] on multivariate long short-term memory (LSTM) forecasting at 15-minute resolution, and Koohfar et al. [7], who found transformer architectures to outperform ARIMA, SARIMA, RNN and LSTM baselines.

2.2 Centralized and decentralized decision-making

There is a stream of literature on the relative merits of centralized and decentralized coordination. Singh et al. [1] scheduled 200 EVs over one day and concluded that

centralized methods offer better load flattening whereas decentralized methods offer scalability. Dudkina et al. [8] applied mixed-integer quadratically constrained centralized coordination on motorways and reported a 50–60% reduction in queue waiting time. Casella et al. [9] developed an augmented-Lagrangian decentralized model predictive control for parking lots and tested it at 5 and 40 EVs. Paudel et al. [10] proposed a consensus-driven, fully decentralized protocol that preserves privacy and scales with fleet size. Wang et al. [11] addressed placement under demand uncertainty with a three-phase stochastic coverage model. A common feature of these studies is evaluation at one or two fleet sizes; the transition between the low-load regime, in which architecture is immaterial, and the saturated regime, in which it is decisive, is rarely mapped.

2.3 Microscopic traffic simulation for EV studies

The open-source SUMO simulator provides detailed vehicle dynamics, signal control and native charging-station and battery models [12]. Wang et al. [13] coupled SUMO and TraCI with a power-grid model for a district of Dresden. Surmann et al. [14] used SUMO's multi-agent interface for bidirectional charging in renewable-energy communities. Sayarshad [15] integrated power-aware routing and charging coordination in SUMO.

2.4 Research gap

The literature spans operational scheduling [3, 5–7, 16], station placement [2, 4, 17–25] and renewable integration [14, 15, 26]. Two gaps remain. First, port-level selection driven by an explicit polling protocol, in which the information available to the decision-maker is the design variable, has not been isolated from the scoring rule itself. Second, centralized and decentralized policies are seldom compared across a range of fleet sizes wide enough to expose the saturation regime, with replicated seeds that separate architectural effects from stochastic variation. The present work addresses both.

3 Research Methodology

3.1 Overview

The evaluation proceeds in five steps: (i) acquisition and editing of the Muscat road network from OpenStreetMap; (ii) configuration of two charging sites with three port classes each in SUMO; (iii) specification of the PBP-ABM decision logic; (iv) implementation of three coordination policies through TraCI; and (v) a factorial scaling experiment with replicated seeds.

3.2 Road network

The network was extracted from OpenStreetMap with the osmWebWizard tool for a selected area of Muscat (Figures 1–4). Disconnected branch roads were removed and the remaining carriageways made bidirectional in NETEDIT. The resulting car-only

network comprises 37 drivable edges totalling 3.6 km. Figure 5 shows the placement of the two charging sites, CS_1 on a five-lane edge and CS_2 on a four-lane edge.



Fig. 1 osmWebWizard interface showing the Muscat road-network selection

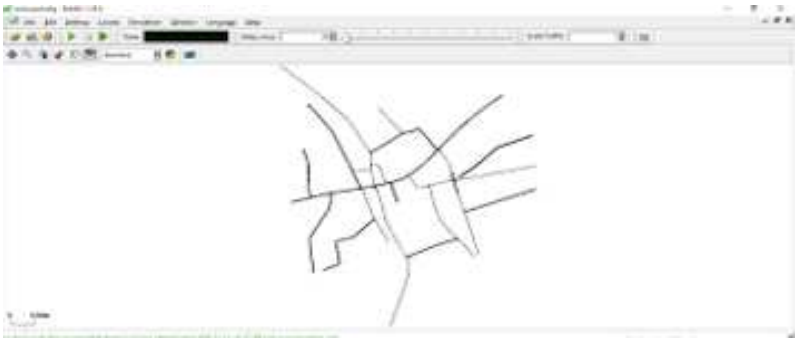


Fig. 2 Imported Muscat road network prior to editing

3.3 Charging infrastructure

Each site hosts one port of each class listed in Table 1. Ports are 15 m SUMO `chargingStation` elements on separate lanes of the same edge. The service time is the dwell imposed when a vehicle stops at the port; it decreases with power class. Infrastructure is held fixed throughout the experiment so that saturation emerges from fleet growth rather than from any change in capacity. Summing port service rates, the two sites can complete approximately 99 charging events per hour.

3.4 EV agents

All EVs use the vehicle type `ev` with maximum speed 25 m/s, acceleration 1.0 m/s^2, deceleration 4.5 m/s^2, and SUMO's battery device with 50 kWh capacity and 40% initial state of charge. Every EV requires one charging event before completing its trip.

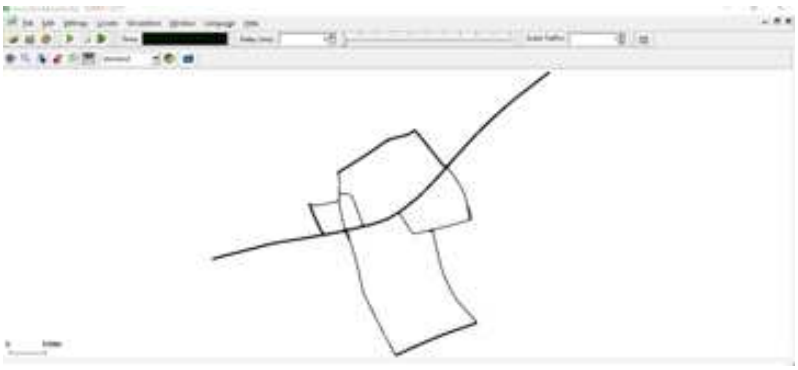


Fig. 3 Edited Muscat road network with bidirectional lanes

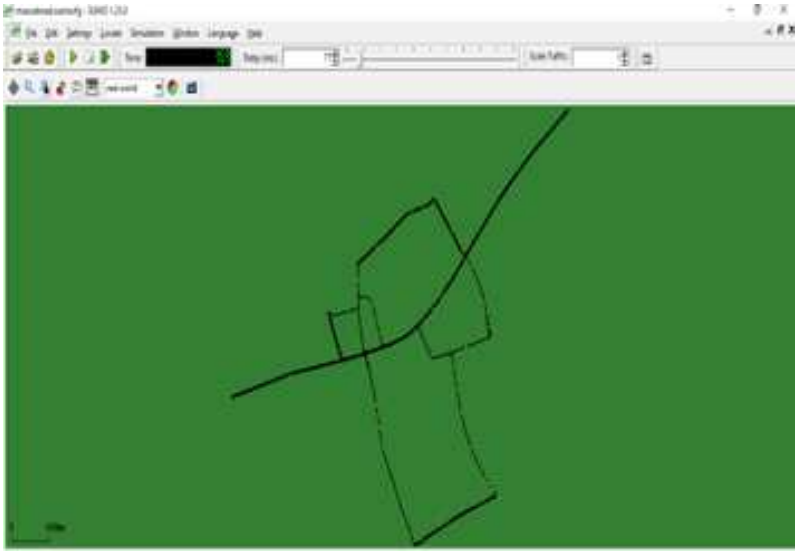


Fig. 4 SUMO simulation of the Muscat road network with active vehicle traffic

Origins and destinations are drawn from the twelve origin–destination (O–D) pairs of the pilot demand files; each trip is routed through the selected port and on to its destination (Equation (4)).

3.5 Priority-Based Polling Agent-Based Model

The PBP-ABM comprises four agent classes: EV agents, port agents, the road network and, in the centralized policies, a controller agent. Figure 6 illustrates the interdependency of the microscopic and macroscopic parameters that shape charging-system behaviour. Two priority levels are defined. Port priority reflects power class and current occupancy; EV priority reflects the state of the vehicle relative to each candidate port. An EV polls candidate ports for their occupancy state and selects the port with

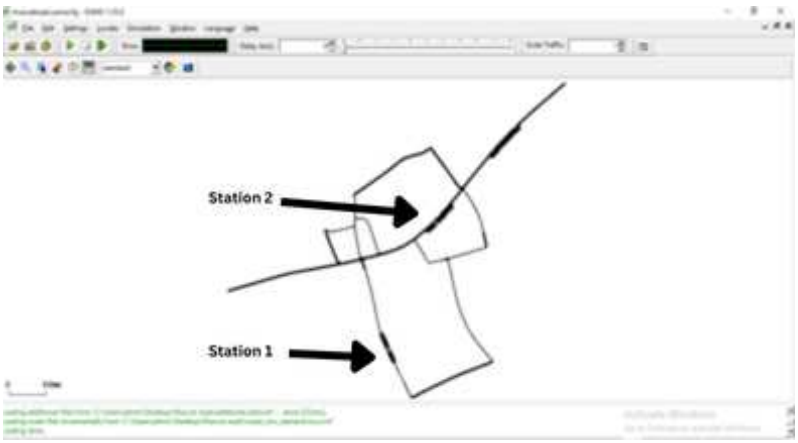


Fig. 5 Muscat road network with the placement of the two charging sites

Table 1 Charging-port configuration (identical at both sites)

Port class	Power (kW)	Service time c (s)	Service rate (EV/h)
Fast (F)	21	100	36.0
Medium (M)	11	400	9.0
Normal (N)	8	800	4.5

the lowest composite score,

$$S_{ij} = \alpha\, t_{ij} + \beta\, w_j + \gamma\, c_j + \delta\, p_{ij}, \tag{1}$$

where t_{ij} is the estimated travel time from the current position of EV i to port j, w_j is the expected waiting time at port j, c_j is the service time of the port class (Table 1), and p_{ij} is the lane-change penalty. The expected waiting time aggregates the residual energy demand of the vehicles counted at the port,

$$w_j = \sum_{k \in Q_j} \frac{E_k^{\mathrm{cap}} - E_k}{r}, \tag{2}$$

where Q_j is the set of vehicles the decision-maker knows to be at port j, E_k^{cap} and E_k are the battery capacity and current charge of vehicle k, and $r = 10$ is a nominal charging rate. The lane-change penalty is

$$p_{ij} = 5\,|\ell_j - \ell_i|, \tag{3}$$

with ℓ_j the lane index of the port and ℓ_i the current lane of the vehicle. Once a port is chosen, the route is rebuilt as

$$\begin{aligned} \text{Route}_i = \text{Route}(\text{current} \to \text{CS}_j) \\ \| \, \text{Route}(\text{CS}_j \to \text{destination}), \end{aligned} \tag{4}$$

and the vehicle is speed-limited to 80% of its maximum within 100 route edges of the port to facilitate the lane change. The weights are fixed at $\alpha = 2.0$, $\beta = 0.5$, $\gamma = 0.2$ and $\delta = 1.0$ in all policies; they were chosen during model development on the pilot scenarios so that travel time dominates while queueing and port class remain influential. Equation (1) is the only scoring rule used; the policies differ in what each decision-maker can observe when evaluating it.

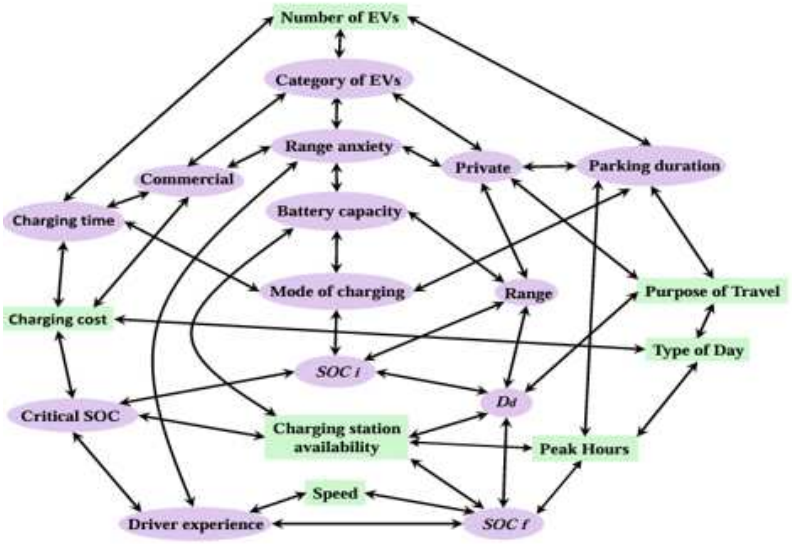


Fig. 6 Interdependency of EV charging-system parameters (adapted from Nanyang Technological University thesis, 2019)

3.6 Coordination policies

Three policies are compared (Table 2).

Centralized distance-priority (C-DP). A controller maintains a global registry of the vehicles that have begun charging at each port. For each site it estimates travel time to a reference port, selects the least-loaded port by Equation (2), and scores the site with $\alpha t + \beta w + \gamma c$. The EV is assigned to the best port of the best site. There is no re-assignment.

Centralized composite-cost (C-CC). As C-DP, but every port is scored individually with the full Equation (1) including the lane-change term, so that a port on a lane closer to the vehicle can outrank a nominally less loaded one.

Decentralized distance-priority (D-DP). There is no controller and no shared registry. Each EV estimates w_j from the vehicles physically present at port j at the moment of polling, obtained through the port agent's local state. A uniform random perturbation $\varepsilon \sim U(0, 5)$ is added to each site score to discourage simultaneous herding to the same port. If a vehicle's current waiting time exceeds 120 s and at least 100 s (a cooldown) have elapsed since its last decision, it re-polls and may re-route to a different port.

A fourth variant, in which decentralized agents scored ports with Equation (1) but read the same global registry as the controller, was also implemented. It reproduced C-CC exactly at every scale, because with identical information and an identical rule the decisions coincide. It is therefore not reported as a separate policy; the observation confirms that the differences reported below arise from the information architecture rather than from the scoring rule.

Table 2 Coordination policies compared

Policy	Decision-maker	Occupancy information Q_j	Lane term δp	Anti-herding ε	Re-polling
C-DP	Global controller	Registry of vehicles that have started charging	–	–	–
C-CC	Global controller	Registry of vehicles that have started charging	Yes	–	–
D-DP	Each EV	Vehicles physically at the port at polling time	–	$U(0, 5)$	After 120 s waiting, 100 s cooldown

All policies use the weights $\alpha = 2.0$, $\beta = 0.5$, $\gamma = 0.2$ and $\delta = 1.0$.

3.7 Model development

The decision logic was developed in three stages on a pilot scenario of 11–12 EVs: an energy-cost model that assigned EVs to the least-queued port; the multi-criteria score of Equation (1); and a normalised-weight variant. At pilot scale the stages were indistinguishable in service outcome, which motivated the scaling experiment reported here. Only the multi-criteria score is evaluated at scale, because it is the stage in which the polling information available to the decision-maker enters explicitly.

3.8 Scaling experiment

The experiment is a full factorial of fleet size $N \in \{11, 50, 100, 250, 500\}$, demand level (low, high) and policy (C-DP, C-CC, D-DP), with five random seeds per cell, giving 150 runs. Demand is defined as an arrival rate rather than a fixed window so that larger fleets do not simply compress into a burst: departure times are drawn uniformly over a window of length N/λ, with $\lambda = 60$ EV/h (low) and $\lambda = 120$ EV/h (high).

Relative to the aggregate service capacity of approximately 99 EV/h, the two levels place the charging system below and above saturation, with utilisation factors $\rho \approx 0.6$ and $\rho \approx 1.2$ respectively. For each seed, O–D pairs are sampled from the twelve pilot pairs and a demand file is written; the same demand file is used for all three policies at that seed, so that policy comparisons are paired. The SUMO random seed and the Python random state are both set from the run seed. A run ends when every vehicle has left the network or after the departure window plus 3600 s, whichever is first. Vehicles that SUMO removes from a standstill of more than 300 s (teleports) are counted as a congestion indicator.

3.9 Performance metrics

The following are recorded per run and reported as mean ± standard deviation over seeds.

- *Service rate*: percentage of vehicles that started a charging event.
- *Time to charging start*: mean duration from departure to charging start among vehicles that started charging. This is the primary access metric, since it includes both travel to the port and queueing at the port.
- *Traffic waiting time*: SUMO output of accumulated waiting time (speed below 0.1 m/s, excluding scheduled stopping time); reported as the mean and 95th percentile among served vehicles.
- *Port occupancy*: fraction of simulated time during which a port had a vehicle connected; reported as the mean over all six ports and as the maximum (busiest port).
- *Energy delivered*: accumulated energy charged at the charging station by the vehicle, obtained from SUMO's charging-station output; reported as the mean and 95th percentile among served vehicles.
- *Teleports* and, for D-DP, the number of *re-polling events*.

Differences between D-DP and C-DP at each fleet size were tested with Welch's t-test on the five paired seeds; p-values are reported without correction and should be interpreted as descriptive.

3.10 Implementation and reproducibility

The simulation was run headless with SUMO 1.27 and Python 3.12. Vehicle state was obtained through TraCI subscriptions once per step, and control commands were issued only on a change of state. The runner, the demand files for all 50 seed–fleet–demand combinations, the raw per-run results and the aggregation scripts are provided as supplementary material.

4 Results

4.1 Below charging capacity

Table 3 and Figure 7 summarise the low-demand runs. Every centralized run served all vehicles at every fleet size, and the mean time to charging start stabilised near

140 s from 100 EVs onward. Mean port occupancy rose from 0.19 at 11 EVs to 0.49 at 500 EVs, with the busiest port occupied 62% of the time, well short of saturation. The decentralized policy started charging sooner at every fleet size (78 s versus 136–142 s at 500 EVs; $p < 0.001$ from 250 EVs) but left 1–2% of vehicles unserved. Traffic waiting time was essentially flat across policies and fleet sizes (29–32 s).

Table 3 Low demand ($\lambda = 60$ EV/h): mean ± SD over five seeds

Policy	N	Served (%)	Time to charging start (s)	Traffic wait (s)	Mean occupancy	Peak occupancy
C-DP	11	100 ± 0	86 ± 32	28.2 ± 5.9	0.19 ± 0.01	0.43 ± 0.02
	50	100 ± 0	123 ± 36	29.4 ± 3.4	0.35 ± 0.06	0.51 ± 0.07
	100	100 ± 0	140 ± 27	32.4 ± 1.7	0.41 ± 0.02	0.54 ± 0.03
	250	100 ± 0	140 ± 7	30.9 ± 0.9	0.46 ± 0.02	0.58 ± 0.04
	500	100 ± 0	142 ± 8	32.0 ± 0.9	0.48 ± 0.01	0.60 ± 0.03
C-CC	11	100 ± 0	86 ± 32	28.2 ± 5.9	0.19 ± 0.01	0.43 ± 0.02
	50	100 ± 0	124 ± 32	29.5 ± 3.9	0.34 ± 0.05	0.49 ± 0.07
	100	100 ± 0	137 ± 24	32.2 ± 1.8	0.41 ± 0.02	0.54 ± 0.02
	250	100 ± 0	139 ± 8	30.8 ± 0.3	0.45 ± 0.02	0.58 ± 0.03
	500	100 ± 0	136 ± 6	32.1 ± 0.8	0.49 ± 0.01	0.62 ± 0.04
D-DP	11	100 ± 0	61 ± 23	24.9 ± 6.4	0.17 ± 0.05	0.37 ± 0.07
	50	98.8 ± 1.1	73 ± 15	29.5 ± 5.8	0.31 ± 0.03	0.53 ± 0.05
	100	98.0 ± 1.2	76 ± 10	30.2 ± 2.6	0.35 ± 0.02	0.57 ± 0.04
	250	98.3 ± 1.0	80 ± 3	29.8 ± 0.6	0.39 ± 0.02	0.60 ± 0.03
	500	98.8 ± 0.4	78 ± 3	29.4 ± 0.9	0.40 ± 0.02	0.61 ± 0.02

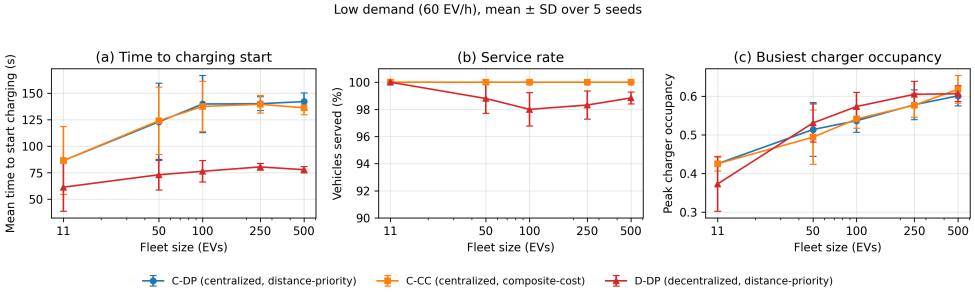


Fig. 7 Low demand ($\lambda = 60$ EV/h): (a) mean time to charging start, (b) service rate and (c) busiest-port occupancy versus fleet size. Points are means over five seeds; bars are one standard deviation

4.2 Above charging capacity

Table 4 and Figure 8 summarise the high-demand runs. Up to 50 EVs the three policies were practically indistinguishable in access time. From 100 EVs the centralized policies diverged from the decentralized one: mean time to charging start rose to 220–232 s at 250 EVs and to 457–471 s at 500 EVs, with a standard deviation across seeds of

roughly 240–300 s at the largest fleet, indicating that the outcome of a centralized run at saturation depended strongly on the particular arrival sequence. Peak port occupancy reached 0.85 at 250 EVs and plateaued thereafter, while mean occupancy also plateaued at 0.64–0.65: additional demand beyond 250 EVs was no longer converted into additional throughput. The decentralized policy kept mean time to charging start at 115 s at 250 EVs and 180 s at 500 EVs ($p = 0.031$ and $p = 0.051$ against C-DP), with peak occupancy 0.70 and 0.75 ($p = 0.021$, $p = 0.001$). Its service rate, however, fell to 97.3% at 250 EVs and 96.4% at 500 EVs ($p = 0.002$ in both cases), and the number of re-polling events grew from 27 to 157 per run. The waiting time for traffic was ordered in the same way but with lesser separation (34 s compared to 40–42 s at 500 EVs). The energy delivered was reduced by 9–13% at D-DP when compared with the 500 EVs case, which agrees with the unserved fraction and the reduced dwell at the fast ports.

Table 4 High demand ($\lambda = 120$ EV/h): mean ± SD over five seeds

Policy	N	Served (%)	Time to charging start (s)	Traffic wait (s)	Mean occupancy	Peak occupancy	Re-polls
C-DP	11	100 ± 0	115 ± 58	31.1 ± 7.2	0.21 ± 0.06	0.50 ± 0.09	–
	50	100 ± 0	126 ± 36	31.2 ± 3.9	0.43 ± 0.04	0.70 ± 0.09	–
	100	100 ± 0	165 ± 36	34.5 ± 2.5	0.53 ± 0.04	0.75 ± 0.07	–
	250	100 ± 0	232 ± 81	38.6 ± 4.2	0.62 ± 0.05	0.83 ± 0.08	–
	500	99.4 ± 0.6	471 ± 238	42.4 ± 3.3	0.65 ± 0.03	0.84 ± 0.03	–
C-CC	11	100 ± 0	115 ± 58	31.1 ± 7.2	0.21 ± 0.06	0.50 ± 0.09	–
	50	100 ± 0	126 ± 36	31.2 ± 3.9	0.43 ± 0.04	0.70 ± 0.09	–
	100	100 ± 0	156 ± 30	34.2 ± 1.9	0.53 ± 0.04	0.78 ± 0.09	–
	250	99.9 ± 0.2	220 ± 37	38.2 ± 1.2	0.65 ± 0.04	0.85 ± 0.06	–
	500	98.7 ± 2.1	457 ± 297	40.4 ± 2.3	0.64 ± 0.08	0.86 ± 0.07	–
D-DP	11	100 ± 0	67 ± 9	30.0 ± 7.4	0.23 ± 0.03	0.52 ± 0.07	0
	50	97.6 ± 2.2	99 ± 20	32.4 ± 4.0	0.43 ± 0.05	0.60 ± 0.03	6 ± 4
	100	98.8 ± 1.3	106 ± 12	32.1 ± 3.2	0.52 ± 0.02	0.64 ± 0.06	9 ± 2
	250	97.3 ± 0.8	115 ± 15	34.0 ± 1.6	0.58 ± 0.03	0.70 ± 0.04	27 ± 8
	500	96.4 ± 1.2	180 ± 50	34.4 ± 2.0	0.60 ± 0.02	0.75 ± 0.01	157 ± 68

Re-polling applies to D-DP only.

4.3 Per-port occupancy

Figure 9 decomposes occupancy by port at 500 EVs under high demand. The centralized policies loaded the Normal ports most heavily (0.79–0.82 at both sites), the Medium ports next (0.67–0.74) and the Fast ports least (0.42–0.45). The decentralized policy inverted this at site 1 (Fast 0.63, Medium 0.75, Normal 0.66) and reduced the load on the Normal port of site 2 to 0.46. The same inversion, in weaker form, was present under low demand, where D-DP left the site-2 Normal port almost idle (0.09) while raising use of the site-1 Fast port to 0.50.

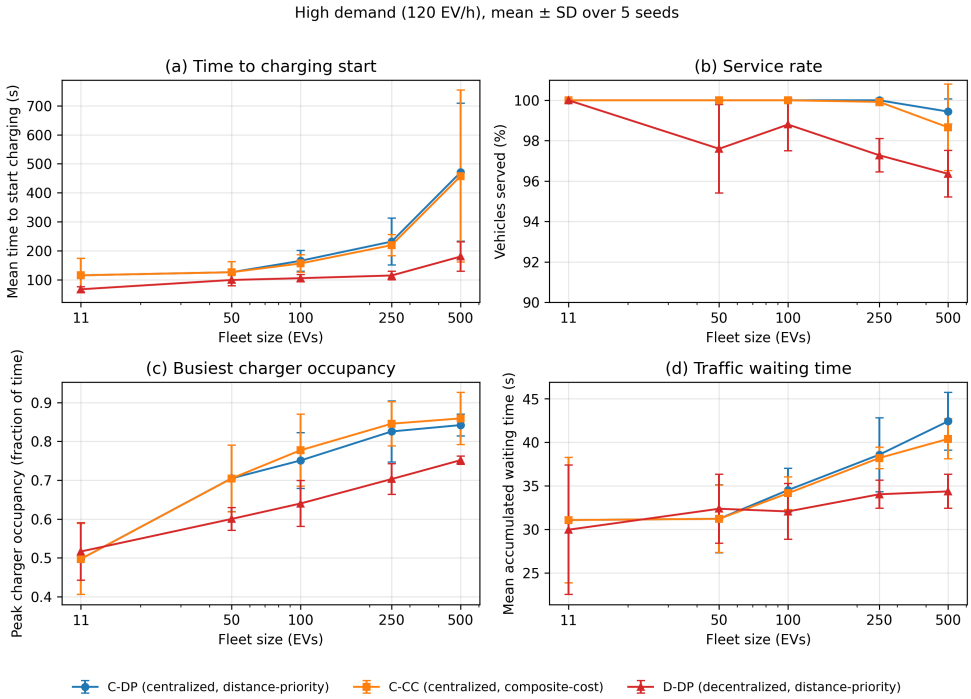


Fig. 8 High demand ($\lambda = 120$ EV/h): (a) mean time to charging start, (b) service rate, (c) busiest-port occupancy and (d) mean accumulated traffic waiting time versus fleet size. Points are means over five seeds; bars are one standard deviation

4.4 Network congestion

Teleports increased with fleet size under all policies (Figure 10(a)), from about 9 per run at 11 EVs to 341–355 under the centralized policies and 219 under D-DP at 500 EVs. Because a teleported vehicle is removed from its queue and re-inserted downstream, the lower count under D-DP indicates shorter queues on the charging edges rather than a different road-network load, since the demand files are shared across policies. Re-polling events under D-DP (Figure 10(b)) rose steeply between 250 and 500 EVs, marking the point at which local re-decision became the dominant behaviour.

5 Discussion

5.1 Architecture matters only near saturation

Below charging capacity, and up to about 50 EVs above it, the three policies produced statistically indistinguishable access times and identical service rates. This is the regime in which most single-scenario comparisons in the literature are conducted, including the pilot stage of the present work, and it explains why such comparisons often report that the control architecture makes little difference. The divergence

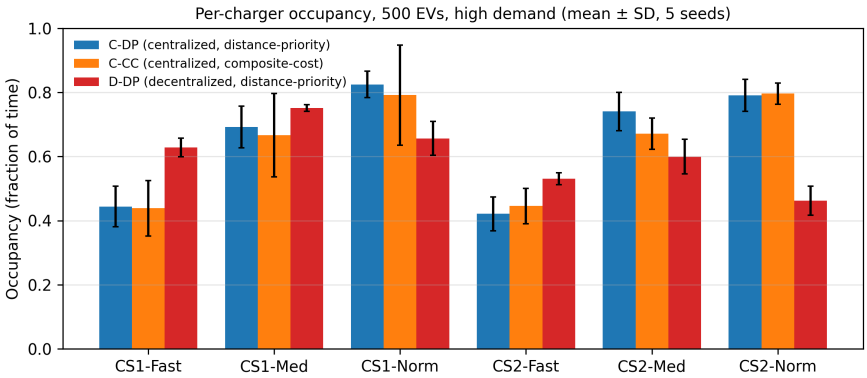


Fig. 9 Per-port occupancy at 500 EVs under high demand (mean over five seeds; bars are one standard deviation)

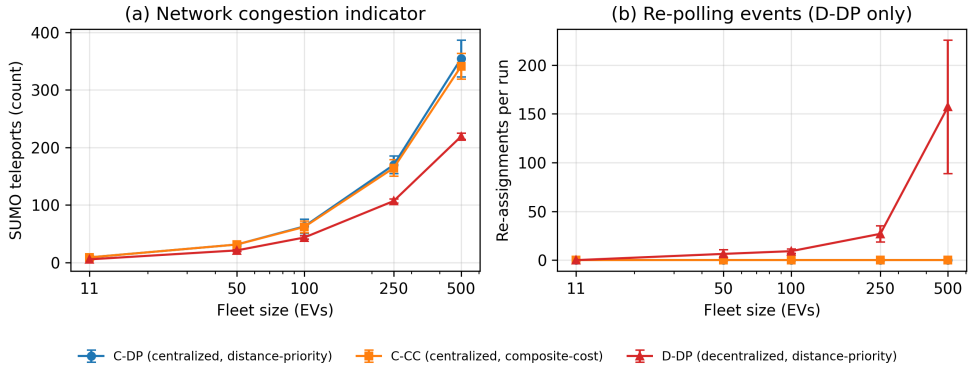


Fig. 10 High demand: (a) SUMO teleports per run as a congestion indicator and (b) re-polling events per run under D-DP, versus fleet size

appears only when mean port occupancy exceeds roughly 0.6, and it grows with fleet size. Any claim about the relative merit of centralized and decentralized EV charging coordination therefore needs to state the utilisation regime to which it applies.

5.2 Mechanism of divergence

The per-port occupancy pattern identifies the mechanism. The centralized registry records a vehicle only once it has begun charging, so a port that has just been chosen by several en-route vehicles still appears empty to the controller. At high arrival rates this stale view sends successive arrivals to the same port, and because the γc term makes fast ports attractive until their registry fills, the overflow lands on the slow Normal ports, where an 800 s service time converts a short queue into a long delay. The large across-seed variance of centralized access time at 500 EVs is the imprint of this cascade: whether it occurs or not depends on the order in which the seeds arrive. The decentralized policy models a vehicle's physical presence, including some vehicles

that are waiting in line, and has a re-polling rule, in which a vehicle that waited for 120 s or more can give up a slow port and switch to a fast one. The reasons why some of the re-polled vehicles do not stop are that they arrive at the edge of the new port after the stop position has been passed and cannot stop (the majority of the 3–4% unserved). The two centralized policies were the same except for the inclusion of the lane-change penalty in the score; they were identical at 11 and 50 EVs and within one standard deviation of each other thereafter.

5.3 Relation to prior work

The finding that centralized coordination loses its advantage at scale is consistent with Singh et al. [1], who attributed scalability to decentralized methods, and with the emphasis on decentralised protocols in Paudel et al. [10] and Casella et al. [9]. It qualifies the large queue-time reductions reported for centralized coordination by Dudkina et al. [8]: those gains rely on the controller having an accurate occupancy picture, which a registry that lags physical arrivals does not provide. The results also suggest that the polling information available to the decision-maker, rather than the sophistication of the scoring rule, is the lever that determines behaviour under load, which is a design consideration for the multi-agent formulations of Khan et al. [3].

5.4 Implications

Three points are relevant for an operator. First, a centralized assignment scheme ought to count assigned cars that are moving, not just connected cars; in this case, the stale-registry cascade should be avoidable, and testing it is a logical next step. Secondly, a re-polling mechanism combined with a waiting-threshold provides most of the benefit of decentralisation from access-time perspective, and needs to be complemented with a stop-position check: otherwise, re-routed vehicles may overshoot the port. Third, on a network of this scale the road will start to become congested when the demand is roughly 250 EVs per hour, at which point coordination is not sufficient; infrastructure expansion is also necessary.

5.5 Limitations

The road network is a 3.6 km extract with two charging sites; the results characterise saturation of a small charging system and should not be extrapolated to city-wide deployment without a larger network and more sites. Demand is synthetic, drawn from twelve O–D pairs, and all EVs share battery parameters and a 40% initial state of charge; user heterogeneity in urgency, willingness to pay or destination flexibility is not modelled. Charging is represented by a fixed dwell per port class rather than by a state-of-charge target, and tariffs, grid constraints and battery degradation are excluded. The score weights were fixed rather than tuned per policy, so the comparison is between information architectures under one rule, not between optimally tuned policies. Five seeds per cell limit the power of the statistical comparisons, and the reported p-values are uncorrected for the number of fleet sizes examined. Finally, SUMO teleports remove vehicles from standstills after 300 s; they are reported rather

than suppressed, but they mean that the congestion at 500 EVs is partly resolved by the simulator rather than by the policies.

6 Conclusions

This study evaluated a Priority-Based Polling Agent-Based Model for EV charging-port selection on a Muscat road network across fleet sizes from 11 to 500 EVs, at arrival rates below and above charging capacity, with five replicated seeds per configuration. Three coordination policies sharing one scoring rule were compared: two centralized assignment policies with a global occupancy registry and one decentralized policy with local observation, anti-herding perturbation and re-polling. The main findings are as follows.

- Below charging capacity, the architecture is immaterial: all policies serve every vehicle, access time is flat in fleet size, and peak port occupancy stays below 0.62.
- Above capacity, centralized assignment degrades from 100 EVs onward; mean time to charging start reaches 457–471 s at 500 EVs with very high across-seed variance, and throughput plateaus at a mean occupancy of about 0.65.
- Decentralized polling with re-decision holds access time to 180 s and peak occupancy to 0.75 at 500 EVs, at the cost of 3.6% unserved vehicles.
- The divergence is caused by the information available to the decision-maker: a registry that counts only connected vehicles funnels late arrivals to slow ports, whereas observation of physical presence redistributes them.
- The pilot-scale evaluation of 11–12 EVs, on which the model was developed, could not have revealed any of these effects.

Future work will extend the network to multiple districts of Muscat with additional charging sites, introduce heterogeneous initial state of charge and state-of-charge charging targets, test a centralized registry that counts assigned vehicles, and examine a hybrid architecture in which a lightweight coordinator publishes port state to locally deciding agents. Learning-based adaptation of the score weights and vehicle-to-grid interaction remain longer-term directions.

Supplementary information. The simulation runner, the SUMO network and charging-station files, the 50 generated demand files, the raw per-run results (150 runs) and the aggregation scripts are provided as supplementary material.

Acknowledgements. This research aligns with Oman Vision 2040, which identifies sustainable transportation and clean energy adoption as core pillars of national development. The authors acknowledge that the simulation case study, based on the urban road network of Muscat, Sultanate of Oman, contributes toward the Vision's broader sustainability goals of reducing carbon emissions, promoting electric mobility, and supporting the transition to a diversified, environmentally sustainable energy infrastructure.

Declarations

- **Funding:** This research received no external funding.

- **Conflict of interest/Competing interests:** The authors declare that they have no conflicts of interest to this work.
- **Ethics approval and consent to participate:** This study does not contain any studies with human or animal subjects performed by any of the authors.
- **Consent for publication:** Not applicable.
- **Data availability:** The simulation runner, the SUMO network and charging-station files, the 50 generated demand files, the raw per-run results (150 runs) and the aggregation scripts are provided as supplementary material and will be deposited in a public repository upon acceptance.
- **Materials availability:** Not applicable.
- **Code availability:** The simulation and aggregation code is provided as supplementary material and will be deposited in a public repository upon acceptance.
- **Author contribution: Ramalingam Dharmalingam:** Conceptualization, Methodology, Software, Investigation, Formal analysis, Writing – original draft, Writing – review & editing. **Vaishnavi Dharmalingam:** Methodology, Software, Investigation, Formal analysis, Writing – original draft, Writing – review & editing. **Mahalakshmi Dharmalingam:** Conceptualization, Methodology, Formal analysis, Writing – original draft, Writing – review & editing.

References

[1] Singh, S.K., Yang, L., Hao, M., Voulgaris, P.G.: Centralized and decentralized schemes for optimal scheduling of electric vehicles. Indian Journal of Science and Technology **14**(19), 1554–1564 (2021) https://doi.org/10.17485/IJST/v14i19.502

[2] Nguyen, M.-D., Le, D.D., Nguyen, P.L.: Optimizing electric vehicle charging station placement using reinforcement learning and agent-based simulations. Preprint at https://doi.org/10.48550/arXiv.2511.01218 (2025)

[3] Khan, J.S., Navaratne, M.A.U.S., Ekanayake, J.B.: Modeling electric vehicle charging station behavior using multiagent system. In: Sheremet, I.A. (ed.) Multi-Agent Technologies and Machine Learning. IntechOpen, London (2023). https://doi.org/10.5772/intechopen.105613

[4] Çelik, S., Ok, Ş.: Electric vehicle charging stations: Model, algorithm, simulation, location, and capacity planning. Heliyon **10**, 29153 (2024) https://doi.org/10.1016/j.heliyon.2024.e29153

[5] Bellizio, F., Dijkstra, B., Fertig, A., Dijk, J., Heer, P.: Machine learning approaches for the prediction of public EV charge point flexibility. Sustainable Energy, Grids and Networks, 101657 (2025) https://doi.org/10.1016/j.segan.2025.101657

[6] Sanami, S., Mosalli, H.: Demand forecasting for electric vehicle charging stations using multivariate time-series analysis. Preprint at https://doi.org/10.48550/arXiv.2502.16365 (2025)

[7] Koohfar, S., Woldemariam, W., Kumar, A.: Prediction of electric vehicles charging demand: A transformer-based deep learning approach. Sustainability **15**(3), 2105 (2023) https://doi.org/10.3390/su15032105

[8] Dudkina, E., Scarpelli, C., Apicella, V., Ceraolo, M., Crisostomi, E.: Optimised centralised charging of electric vehicles along motorways. Sustainability **17**(12), 5668 (2025) https://doi.org/10.3390/su17125668

[9] Casella, V., Ferro, G., Robba, M.: A decentralized optimization approach to the power management of electric vehicles parking lots. Sustainable Energy, Grids and Networks **38**, 101301 (2024) https://doi.org/10.1016/j.segan.2024.101301

[10] Paudel, A., Hussain, S.A., Sadiq, R., Zareipour, H., Hewage, K.: Decentralized cooperative approach for electric vehicle charging. Journal of Cleaner Production **364**, 132590 (2022) https://doi.org/10.1016/j.jclepro.2022.132590

[11] Wang, J., Kaushik, H.D., Jacob, R.A., Zhang, J.: Spatiotemporal planning of electric vehicle charging infrastructure: Demand estimation and grid-aware optimization under uncertainty. iScience **28**(9), 113368 (2025) https://doi.org/10.1016/j.isci.2025.113368

[12] Lopez, P.A., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.-P., Hilbrich, R., Lücken, L., Rummel, J., Wagner, P., Wießner, E.: Microscopic traffic simulation using SUMO. In: 2018 21st International Conference on Intelligent Transportation Systems (ITSC), pp. 2575–2582. IEEE, Maui, HI, USA (2018). https://doi.org/10.1109/ITSC.2018.8569938

[13] Wang, S., Haider, S.I., Shen, S., Motazedian, F., Radeke, R., Fitzek, F.H.P.: Simulation architecture for electric vehicle charging optimization in Dresden's Ostra District. In: Proceedings of the 14th International Conference on Smart Cities and Green ICT Systems (SMARTGREENS 2025), pp. 56–65. SCITEPRESS, Setúbal (2025). https://doi.org/10.5220/0013445000003953

[14] Surmann, A., Walia, R., Kohrs, R.: Agent-based bidirectional charging algorithms for battery electric vehicles in renewable energy communities. Energy Informatics **3**(S1) (2020) https://doi.org/10.1186/s42162-020-00122-8

[15] Sayarshad, H.R.: Optimization of electric charging infrastructure: Integrated model for routing and charging coordination with power-aware operations. npj Sustainable Mobility and Transport **1** (2024) https://doi.org/10.1038/s44333-024-00004-6

[16] Suanpang, P., Jamjuntr, P.: Optimizing electric vehicle charging recommendation in smart cities: A multi-agent reinforcement learning approach. World Electric Vehicle Journal **15**(2), 67 (2024) https://doi.org/10.3390/wevj15020067

[17] Alanazi, F., Alshammari, T.O., Azam, A.: Optimal charging station placement

and scheduling for electric vehicles in smart cities. Sustainability **15**(22), 16030 (2023) https://doi.org/10.3390/su152216030

[18] Iravani, H.: A multicriteria GIS-based decision-making approach for locating electric vehicle charging stations. Transportation Engineering **9**, 100135 (2022) https://doi.org/10.1016/j.treng.2022.100135

[19] Calvo-Jurado, C., Ceballos-Martínez, J.M., García-Merino, J.C., Muñoz-Solano, M., Sánchez-Herrera, F.J.: Optimal location of electric vehicle charging stations using proximity diagrams. Sustainable Cities and Society **113**, 105719 (2024) https://doi.org/10.1016/j.scs.2024.105719

[20] Ahmad, F., Iqbal, A., Ashraf, I., Marzband, M., Khan, I.: Optimal location of electric vehicle charging station and its impact on distribution network: A review. Energy Reports **8**, 2314–2333 (2022) https://doi.org/10.1016/j.egyr.2022.01.180

[21] Kaya, Ö., Alemdar, K.D., Campisi, T., Tortum, A., Kayaci Çodur, M.: The development of decarbonisation strategies: A three-step methodology for the suitable analysis of current EVCS locations applied to Istanbul, Turkey. Energies **14**(10), 2756 (2021) https://doi.org/10.3390/en14102756

[22] Gong, D., Tang, M., Buchmeister, B., Zhang, H.: Solving location problem for electric vehicle charging stations—A sharing charging model. IEEE Access **7**, 137482–137495 (2019) https://doi.org/10.1109/ACCESS.2019.2943079

[23] Karolemeas, C., Tsigdinos, S., Tzouras, P.G., Nikitas, A., Bakogiannis, E.: Determining electric vehicle charging station location suitability: A qualitative study of Greek stakeholders employing thematic analysis and analytical hierarchy process. Sustainability **13**(4), 2298 (2021) https://doi.org/10.3390/su13042298

[24] Zhao, H., Gao, J., Cheng, X.: Electric vehicle solar charging station siting study based on GIS and multi-criteria decision-making: A case study of China. Sustainability **15**(14), 10967 (2023) https://doi.org/10.3390/su151410967

[25] Guo, S., Liang, F.: The planning of electric vehicle charging stations in the urban area. In: Proceedings of the 2nd International Conference on Electronic and Mechanical Engineering and Information Technology (EMEIT 2012). Atlantis Press, Paris (2012). https://doi.org/10.2991/emeit.2012.356

[26] Kumar, B.A., Jyothi, B., Singh, A.R., Bajaj, M., Rathore, R.S., Tuka, M.B.: Hybrid genetic algorithm-simulated annealing based electric vehicle charging station placement for optimizing distribution network resilience. Scientific Reports **14**, 7637 (2024) https://doi.org/10.1038/s41598-024-58024-8