Delft University of Technology

Document Version
Final published version

Licence
CC BY

Citation (APA)
Shams Ashkezari, L., Lagae, B., Yorke-Smith, N., Saeednia, M., Bauer, P., & Chandra Mouli, G. R. (2026). Planning highway charging networks for heavy-duty electric vehicles: From network design to station-level charging profiles. *Applied Energy*, *426*, Article 128590. https://doi.org/10.1016/j.apenergy.2026.128590

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

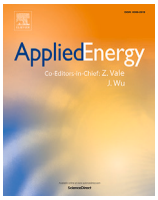
Applied Energy 426 (2026) 128590



Contents lists available at ScienceDirect

Applied Energy

journal homepage: www.elsevier.com/locate/apen



Planning highway charging networks for heavy-duty electric vehicles: From network design to station-level charging profiles☆

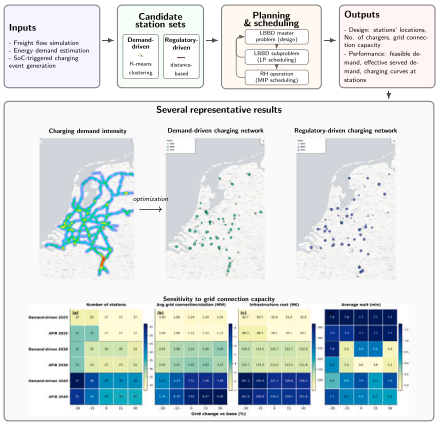
Leila Shams Ashkezari[a,*], Bart Lagae[a], Neil Yorke-Smith[b], Mahnam Saeednia[c], Pavol Bauer[a], Gautham Ram Chandra Mouli[a]

[a] Department of Electrical Sustainable Energy, Delft University of Technology, Delft, the Netherlands
[b] STAR Lab, Delft University of Technology, Delft, the Netherlands
[c] Department of Transport and Planning, Delft University of Technology, Delft, the Netherlands

HIGHLIGHTS

- Demand-driven planning achieves full corridor feasibility, while regulation-driven siting leaves small persistent gaps.
- Regulation- and demand-driven plans overlap at only 14–31% of stations, but shared sites carry more capacity.
- By 2040, planning shifts from expanding spatial coverage to densifying capacity at high-demand corridors.
- Aggregate charging profiles remain similar; key differences are spatial siting and effective demand coverage.
- CC–CV charging shows 1 MW chargers better match practical en-route stop-duration constraints than 350 kW.

GRAPHICAL ABSTRACT



ARTICLE INFO

Keywords:
Charging demand estimation
Charging network planning
Heavy-duty electric vehicles (HDEVs)
Long-haul freight electrification
Logic-based Benders decomposition (LBBD)
Megawatt charging stations (MCS)

ABSTRACT

As freight electrification progresses, the planning of high-power charging infrastructure becomes a critical challenge due to the interaction between time-constrained operations, spatially uneven demand, and limited grid-connection capacity. Current European Union regulation specifies corridor coverage requirements, but does not explicitly account for local demand intensity, station sizing, or grid constraints. This article presents an integrated optimization framework for planning highway charging infrastructure for heavy-duty electric vehicles (HDEVs). The framework combines freight-tour simulation, route- and energy-aware charging event generation, candidate-site identification, and optimization-based station siting, sizing, and event-level scheduling. In addition, a constant-current/constant-voltage (CC–CV) post-processing step is used to obtain more realistic charging durations and aggregate charging profiles. The framework is applied to the Netherlands for 2025, 2030, and 2040 under both demand-driven and regulatory-driven planning paradigms, considering charger ratings of 350 kW and 1 MW. The results show that demand-driven planning achieves full event feasibility on the modelled corridor network, whereas the regulatory-driven planning exhibits a small but persistent feasibility gap due to its more geographically constrained candidate set. Across scenario years, infrastructure expansion follows a two stage

☆ This work was supported by the Dutch Ministry of Economic Affairs and Climate Policy under the project “Plug-in”.
* Corresponding author.
Email address: l.shamsashkezari@tudelft.nl (L. Shams Ashkezari).

https://doi.org/10.1016/j.apenergy.2026.128590
Received 25 March 2026; Received in revised form 13 July 2026; Accepted 29 July 2026
Available online 6 August 2026
0306-2619/© 2026 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

pattern, shifting from initial spatial coverage to capacity densification at persistent freight corridors. Sensitivity analysis further shows that 1 MW networks are substantially more exposed than 350 kW networks to grid limits and vehicle-side charging assumptions. Overall, the proposed framework provides a practical decision support tool for designing scalable, cost-effective, and operationally realistic HDEVs charging networks.

1. Introduction

Transport-related emissions currently account for roughly 25% of the European Union's (EU) total greenhouse gas emissions, and they have been rising in recent years. To address this, the European Commission has proposed measures across climate, energy, mobility, and taxation policies aiming for a minimum 55% reduction in net emissions by 2030, and climate neutrality for Europe by 2050. Achieving this target requires a 90% reduction in transport-related emissions [1]. Road vehicles contribute more than 75% of CO_2 emissions in the transport sector [2], with freight responsible for over 30%, a share expected to grow [3]. Electrification of road transport is therefore a critical step toward decarbonization.

Heavy-duty electric vehicles (HDEVs) are still in the early stages of adoption, making it essential to estimate future demand and plan corresponding charging infrastructure in advance. In contrast to light-duty electric vehicles, HDEVs carry much larger batteries—typically 300–1000 kWh—to meet higher energy and range requirements. They also have distinctive operational characteristics, operating on strict schedules that limit charging opportunities and require megawatt scale charging, posing unique infrastructure challenges. The EU has set baseline requirements for HDEV charging stations along major corridors, but these remain high-level, suitable primarily for early-stage deployment. Current regulations indicate approximate station placement along core roads but do not define capacity requirements or consider operational factors such as traffic density, peak demand, dwell times, grid constraints, and redundancy. Without detailed planning, infrastructure may be misaligned with real world needs, resulting in bottlenecks, under-utilized assets, and higher costs.

While prior studies provide valuable insights into HDEV charging demand, station placement, and local grid impacts, they are often limited in scope by corridor-specific analyses, aggregated demand representations, or simplified operational assumptions. Section 2 reviews the literature and highlights the remaining research gaps. In response, this article presents an integrated national scale framework for planning a public fast charging network for HDEVs. The framework combines freight-activity-based demand modelling, candidate site generation under both regulatory- and demand-driven paradigms, and optimization-based charging network design with operational feasibility assessment.

The remainder of the article is organized as follows. Section 3 presents the methodology, including truck-trip simulation, routing and energy-demand estimation, candidate charging station identification, and the optimization framework for station placement and sizing. Section 4 describes the problem definition, study inputs, and modelling assumptions. Section 5 presents and discusses the results, including the comparison between demand-driven and regulatory-driven planning, sensitivity analyses, and operational characteristics of the optimized charging networks. Finally, Section 6 concludes the article and outlines directions for future work.

2. Literature review

Research on charging demand and infrastructure deployment has advanced considerably for passenger electric vehicles, providing detailed insights into user behaviour, spatial charging patterns, and grid interactions. Comparable studies for HDEVs, particularly long-haul operations that rely on public en-route charging, remain less mature and are spread across partially-disconnected modelling traditions. Review studies already point to several key challenges: route- and path-dependent energy estimation is essential for realistic demand reconstruction [4], while the connection between charging demand estimation, infrastructure siting, and grid feasibility remains weak in much of the existing literature [5]. In parallel, large scale simulation frameworks have demonstrated the value of trajectory-rich freight representations for national charging infrastructure assessment and policy analysis [6]. Together, these studies suggest the need for integrated planning approaches that combine operationally grounded demand reconstruction, infrastructure optimization, and feasibility assessment under evolving electrification scenarios. This need is especially relevant because recent studies have quantified future charging needs and assessed regulatory-driven sufficiency at an aggregate level [7], while direct comparisons between regulatory-driven and demand-driven network designs under explicit operational feasibility remain limited. The literature reviewed below is organized into three closely related streams: (i) charging demand estimation and operational modelling, (ii) infrastructure placement, corridor optimization, and demand alignment, and (iii) grid integration and hosting capacity assessment.

2.1. Charging demand estimation and operational modelling

Charging demand estimation for HDEVs has been addressed at different levels of detail, ranging from national scale electricity demand assessments to operational studies of charging behaviour and station utilization. At large spatial scales, the main objective is typically to translate freight activity into geographically and temporally detailed charging demand. For example, Diks et al. [5] combined truck energy models with Freight Analysis Framework data to estimate national scale charging load profiles for the United States. Tong et al. [8] similarly derived spatially resolved charging needs for long-haul truck electrification, but relied on simplifying assumptions regarding charging choice, station access, and detours. These studies are valuable for identifying system-level demand patterns, but they do not explicitly resolve how infrastructure placement and local operational constraints affect the resulting charging load.

A second stream of work focuses on converting trip-based freight activity into electricity demand using operational assumptions on vehicle range, dwell opportunities, and charging behaviour. Liimatainen et al. [9] examined the feasibility of freight electrification for medium- and long-haul operations, highlighting the importance of dwell-time availability and range limitations. More recent studies have incorporated trip-chain and stop-duration information to improve the temporal realism of charging energy and power estimates. For instance, Shoman et al. [10] used trip-chain-based representations to better capture when and where charging demand arises. However, these approaches generally stop at demand estimation and do not optimize the charging network itself.

A third line of research considers corridor-level infrastructure requirements. Sauter et al. [11] and Rajon Bernard et al. [12] estimated charging needs along freight corridors by combining traffic intensity with simplified service or queueing representations. Such studies are useful for corridor planning and early infrastructure sizing, but they usually treat charging demand as exogenous and do not capture how alternative site locations, station capacities, or detour feasibility reshape charging behaviour across the network.

Finally, a growing literature addresses the operational realism of HDEV charging by using telematics, facility-level data, or explicit queueing models. Borlaug et al. [13] and Giuliano et al. [14] studied real operational settings such as depots and port drayage, showing that charging feasibility depends strongly on duty cycles, dwell times, and the

availability of mid-shift charging. Walz et al. [15] synthesized station-level arrival and dwell-time patterns to generate time series for station sizing and grid-impact studies. Eriksson et al. [16] showed that queueing and peak power at HDEV charging sites are strongly influenced by stochastic arrivals and limited service capacity, while Karlsson et al. [17] used an agent-based formulation to study competition between charging sites and queue formation in long-haul applications. These studies make it clear that realistic charging demand cannot be separated from operational constraints, but they are often conditional on a predefined network or do not jointly optimize infrastructure placement, capacity, and event-level charging feasibility.

Overall, the literature has made substantial progress in representing freight-driven charging demand, temporal charging behaviour, and station-level operational effects. However, many studies either estimate charging demand without optimizing the infrastructure, or analyze operations assuming an existing charging network. This leaves a gap for integrated frameworks that link freight activity, candidate-site generation, infrastructure design, and operational scheduling feasibility within a single planning approach.

2.2. Infrastructure placement, corridor optimization, and demand alignment

A substantial part of the literature on HDEV charging infrastructure focuses on corridor planning, where the main objective is to ensure spatial coverage while maintaining acceptable service levels. In this stream, station deployment is typically guided by spacing requirements, traffic exposure, and simplified queueing or throughput assumptions. Sauter et al. [11], Speth et al. [18,19], Hall and Lutsey [20], and Rajon Bernard et al. [12] all follow this general logic: they estimate the number and approximate placement of charging facilities needed to satisfy corridor demand under service-oriented assumptions. These studies provide valuable guidance for early-stage infrastructure roll-out and for translating policy targets into approximate network requirements, but they usually rely on exogenous demand inputs and simplified assumptions on charging choice, detours, and time-varying congestion.

A second line of work formulates infrastructure deployment as an explicit optimization problem. These studies determine station locations and capacities by balancing cost, coverage, and operational constraints over highway networks or large regions. For example, ORAGENT [21] and related corridor scale formulations such as Siekmann et al. [22] use optimization approaches to support large scale deployment planning for alternative-fuel infrastructure. Related coverage-oriented models also investigate how existing logistics assets or existing facilities can be leveraged to extend effective charging reach and reduce deployment costs, as in Hurtado-Beltrán et al. [23]. Compared with purely rule-based corridor studies, these optimization-based approaches offer greater flexibility in analyzing cost–coverage trade-offs, but many still treat demand in a largely static form and only partially capture operational interactions such as event timing, queuing, and dynamic detour feasibility.

More recent work has moved toward tighter coupling between station placement and local service capacity. In particular, studies that incorporate explicit throughput constraints, limited simultaneous charging capability, or parking-capacity restrictions show that optimal layouts can differ markedly from those obtained under purely distance-based coverage rules. Speth et al. [24], for example, show that accounting for service-capacity limitations changes both the number and location of preferred sites. In parallel, freight-oriented charging placement problems continue to develop exact and heuristic methods for location selection under operational constraints, as illustrated by Dimitriou et al. [25]. These contributions move closer to practical planning needs, but they still generally do not integrate freight-demand generation, candidate-site construction, infrastructure design, and explicit event-level scheduling feasibility within a unified framework.

As optimization models become more detailed, computational tractability becomes a central issue. Integrating discrete placement decisions with integer numbers of charger sizing and time-indexed event assignments under capacity and waiting constraints leads to very large mixed-integer models, especially at a national scale. For this reason, several studies adopt multi-stage or decomposition solution strategies. In the HDEV context, Zahedi et al. [26], for example, model phased or multi-period deployment under explicit grid capacity constraints to support staged investment planning. Such studies represent important progress toward realistic infrastructure roll-out, but the operational layer is still commonly represented through aggregate service approximations, such as queueing-based sizing or average utilization measures, rather than through an explicit time-indexed scheduling feasibility model.

Overall, the infrastructure placement literature has progressed from rule-based corridor coverage toward increasingly detailed optimization and service-capacity formulations. However, an important gap remains between coverage-oriented network design and operations-based planning. Many existing models either optimize infrastructure planning using simplified or static demand representations, or incorporate operational constraints without explicitly linking them to tour-based freight activity and charging event generation. As a result, it remains unclear whether the resulting layouts are feasible once discrete event assignments, charger availability over time, and waiting constraints are enforced simultaneously. This motivates the integrated planning approach developed in this article, where station placement and sizing are evaluated jointly with charging event feasibility and operational scheduling constraints.

2.3. Grid integration and hosting capacity

Grid integration studies examine how high-power charging demand interacts with existing electricity networks and where reinforcement or flexibility measures become necessary. A first group of studies focuses on the direct impact of HDEV charging on local network components. For example, Then et al. [27] showed that clusters of high-power chargers can overload typical substations, highlighting the need for dedicated feeders or upstream grid reinforcement. Similarly, Borlaug et al. [28] showed that unmanaged depot charging can exceed local grid limits when vehicle charging is highly synchronized, emphasizing the importance of charging coordination even when infrastructure is already available. At the station level, Mishra et al. [29] quantified voltage and thermal constraints caused by HDEV charging and evaluated mitigation options such as local battery storage and managed charging.

A second line of work considers grid constraints more directly within infrastructure planning and investment decisions. Ye et al. [30] and Zahedi et al. [26], for example, incorporate grid capacity considerations or phased expansion logic into charging infrastructure planning. These studies move beyond pure grid impact assessment by recognizing that infrastructure deployment and grid reinforcement decisions are interdependent.

Overall, the grid integration literature shows that hosting capacity can become a binding constraint for HDEV charging deployment and that local flexibility or staged reinforcement may be required to enable large scale roll-out. However, most existing studies remain localized or scenario-specific, and they are typically driven by exogenous charging profiles rather than charging demand that emerges endogenously from jointly-optimized station placement, charger sizing, and event assignment. This leaves a gap for planning frameworks that connect national scale charging network design with operationally grounded charging profiles that can subsequently be used for grid-impact assessment.

2.4. Gap and contribution of this study

Despite recent progress, four methodological gaps remain in national scale planning of public fast charging networks for long-haul HDEVs:

(i) *Limited integration of infrastructure design and operational feasibility:* Demand-estimation studies often assume a fixed charging

network and therefore cannot capture how infrastructure placement and capacity reshape charging assignments, local congestion, and station-level power peaks [5,8]. Conversely, many corridor planning studies optimize layouts using static or aggregated demand representations, with limited explicit treatment of charger availability and waiting constraints over time [11,19,21]. In addition, charger sizing is often derived in a subsequent step using queueing or utilization rules rather than being determined jointly with event-level service feasibility [12]. Although capacity-aware and phased planning is beginning to emerge [26], integrated large scale formulations that jointly optimize charging networks and integer charger numbers while verifying scheduling feasibility, remain limited.

(ii) *Lack of scalable methods for integrated national scale planning:* When event-level operational feasibility is modelled explicitly, fully integrated formulations become computationally difficult due to the large number of event, station and time-coupled decisions. As a result, many national scale studies rely on simplified service representations or heuristic methods. Scalable exact approaches that preserve both strategic design decisions and operational scheduling realism are still scarce.

(iii) *Lack of comparison of regulation-driven and demand-driven designs:* Existing studies have assessed future truck-charging needs and the extent to which regulatory targets may cover them [7]. However, to the best of the authors' knowledge, a direct side-by-side comparison between adaptive, demand-driven charging networks and regulation-driven ones under multiple electrification scenarios and explicit operational feasibility constraints has not yet been reported. As a result, it is still unclear how compliance-oriented layouts compare with demand-driven designs in terms of effective demand satisfaction, infrastructure cost, and operational robustness.

3. Methodology

This section presents the integrated planning framework developed in this study. The framework links freight-activity simulation, charging event and energy-demand estimation, candidate-site generation, infrastructure placement and sizing, and post-processed station-level charging profile reconstruction using a constant-current/constant-voltage (CC–CV) charging profile, as illustrated in Fig. 1. Together, these components enable the translation of freight operations into operationally feasible charging network designs and corresponding station-level load profiles. The following subsections describe each component of the framework in the order shown in the figure.

3.1. Truck trip simulation

Estimating charging demand for HDEVs requires a detailed representation of freight movements, including vehicle departure times, tour structures, stop locations, and vehicle utilization patterns. These operational characteristics determine when vehicles are active on the road network, how far they travel between stops, and where charging opportunities may arise. Therefore, a high-resolution representation of truck tours is used as the starting point for the charging demand estimation framework.

In this work, freight movements are generated using an agent-based freight transport simulator, the Multi-Agent Simulation System for Goods Transport (MASS-GT) [31]. The model represents logistics operations through interacting agents such as producers/receivers, shippers, and carriers. Shipment generation, vehicle selection, and tour formation are modelled using discrete choice and optimization components that capture logistics decision-making processes, including delivery sequencing, vehicle allocation, and routing constraints. The simulator produces high-resolution tour records containing origin and destination

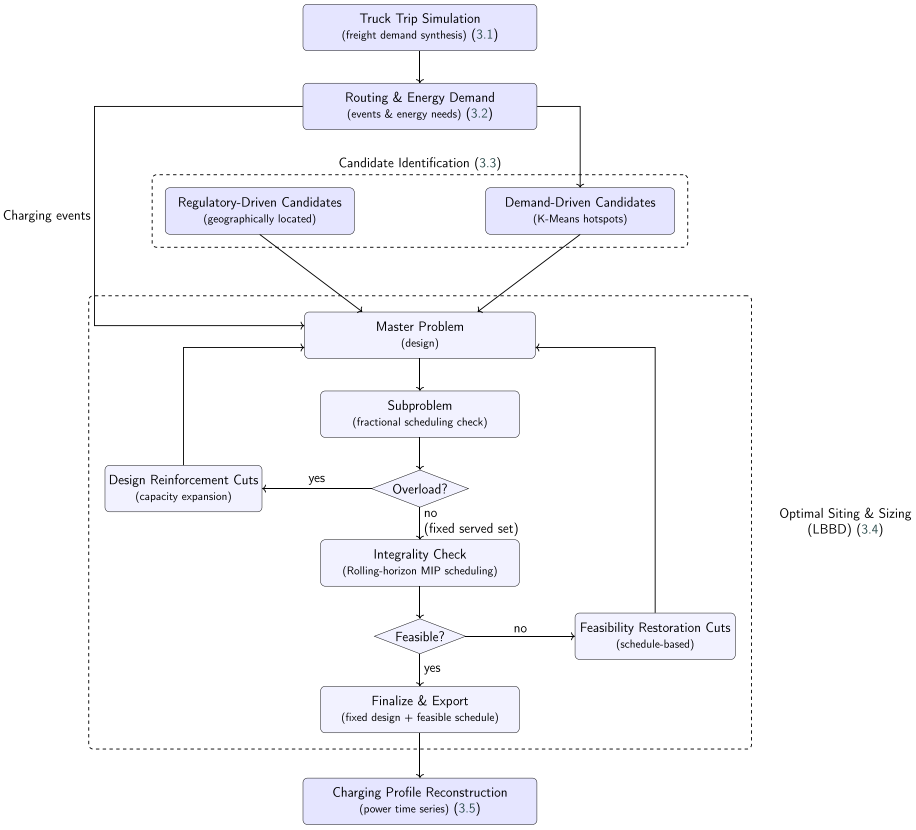


Fig. 1. Overview of the charging infrastructure planning workflow.

Table 1
Vehicle class specifications used to represent electrified truck categories [38–42].

ID	Vehicle type	Battery capacity (kWh)	Range (km)	Consumption rate (kWh/km)
0	Truck (small)	400	454.5	0.88
1	Truck (medium)	500	568.2	0.88
2	Truck (large)	600	681.8	0.88
3	Truck + Trailer (small)	400	333.3	1.20
4	Truck + Trailer (large)	600	500	1.20
5	Tractor + Trailer	600	500	1.20

coordinates, vehicle class, departure time, and total vehicle kilometers traveled. These simulated tours provide the basis for reconstructing spatiotemporal truck activity and identifying potential charging events.

Compared with aggregate flow models or static origin–destination matrices, this tour-based representation captures operational heterogeneity such as multi-stop tours, fragmented tour chains, and variable vehicle utilization. These characteristics strongly influence charging demand because they determine vehicle dwell times, route lengths, and potential charging opportunities along transport corridors.

Because the simulator represents conventional truck operations and does not explicitly provide detailed driving trajectories, additional processing is required to translate simulated tours into HDEV energy demand and charging events. Each tour is therefore mapped onto a road network route using a routing engine (GraphHopper [32]), as described in Section 3.2.

3.2. Routing and energy demand estimation

Each tour of the simulator is assigned to a predefined vehicle class v, characterized by a base battery capacity E_v^{bat} and a specific energy consumption rate η_v (Table 1). To represent parameter uncertainty while preserving reproducibility, tour level parameters are sampled once per tour and held constant across all trips within that tour. Specifically, battery capacity and consumption are sampled from truncated normal distributions centered at the base values with a prescribed coefficient of variation. The resulting tour specific parameters are denoted E_r^{bat} and η_r for tour $r \in \mathcal{R}$.

Step 1: routing on the modelled road network and fixed distance resampling. For each trip k within tour r, a driven path on the modelled road network is obtained from a routing engine. The trip endpoints are mapped to nearby routable road locations to ensure feasibility on the modelled network. Trips whose origin and/or destination lie outside the modelled study area are retained by computing the within-study-area portion of the route and representing the out-of-network parts through boundary connectors. The out-of-network distance is represented by a buffer distance $s_{r,k}^{\text{out}} \geq 0$ (applied separately to pre- and post-network segments when applicable), together with an associated travel time approximation $t_{r,k}^{\text{out}} \geq 0$.

To control output size and impose a uniform spatial resolution, each within-study-area route is resampled to a fixed spacing Δs. The resampled route of trip (r,k) is represented by a sequence of points $p = 0, \ldots, N_{r,k}$ with geographic coordinates, cumulative distance $s_{r,k,p}$, and baseline cumulative travel time $\tilde{t}_{r,k,p}$ provided by the routing engine.

Step 2: state of charge (SoC) simulation with tour consistent timing, breaks, and charging. Scheduled departure times are provided by the tour simulator. Realized departure times are obtained by chaining trips within each tour. For the first trip,

$$t_{r,0}^{\text{dep}} = t_{r,0}^{\text{sched}} + t_{r,0}^{\text{out,pre}}$$

and for subsequent trips $k > 0$,

$$t_{r,k}^{\text{dep}} = t_{r,k-1}^{\text{arr}} + t_{r,k}^{\text{unload}} + t_{r,k}^{\text{out,pre}}$$

Here, $t_{r,k}^{\text{arr}}$ denotes the realized arrival time of trip k after accounting for driving, breaks, and any charging performed during the trip. The term $t_{r,k}^{\text{unload}}$ denotes an inferred loading/unloading duration between trips, obtained from the gap between the previous realized arrival and the next scheduled departure in the original schedule. The term $t_{r,k}^{\text{out,pre}} \geq 0$ represents a pre-network travel time approximation applied when the trip origin lies outside the modelled study area (zero otherwise).

The battery SoC is simulated along the resampled route points. For trip (r,k), the SoC depletion between consecutive points $(p-1)$ and p is computed from the distance increment:

$$\Delta \text{SoC}_{r,k,p} = \frac{\Delta s\, \eta_r}{E_r^{\text{bat}}} \times 100, \qquad \text{SoC}_{r,k,p} = \text{SoC}_{r,k,p-1} - \Delta \text{SoC}_{r,k,p} \tag{1}$$

Driving time is tracked in order to approximate regulatory rest requirements. When cumulative driving since the last qualifying reset exceeds a maximum continuous driving duration H^{drive}, a rest break of duration H^{break} is inserted into the timeline. Dwell times between trips exceeding a minimum qualifying duration reset the driving time counter.

Step 3: charging trigger, energy request, and event definition. A charging decision is evaluated at each route point p. A charging session is triggered when $\text{SoC}_{r,k,p} < \text{SoC}^{\text{trig}}$ and the remaining distance cannot be completed while maintaining a reserve SoC^{res}. The remaining distance from point p is defined as

$$s_{r,k,p}^{\text{rem}} = \left(s_{r,k}^{\text{trip}} - s_{r,k,p}\right) + s_{r,k}^{\text{out,post}} \tag{2}$$

where $s_{r,k}^{\text{trip}}$ is the total within network distance of trip (r,k) and $s_{r,k}^{\text{out,post}} \geq 0$ is an out of network buffer distance when the destination lies outside the modelled study area.

The minimum SoC required to complete the remaining distance is

$$\text{SoC}_{r,k,p}^{\text{need}} = \text{SoC}^{\text{res}} + \frac{s_{r,k,p}^{\text{rem}}\, \eta_r}{E_r^{\text{bat}}} \times 100 \tag{3}$$

and the target SoC is capped at $\text{SoC}^{\max}$ to avoid high SoC charging where power tapering can increase dwell time:

$$\text{SoC}_{r,k,p}^{\text{target}} = \min\left(\text{SoC}^{\max}, \text{SoC}_{r,k,p}^{\text{need}}\right) \tag{4}$$

Whenever $\text{SoC}_{r,k,p}^{\text{need}} > \text{SoC}_{r,k,p}$, the trigger at (r,k,p) defines a charging event $e \in \mathcal{E}$. The requested energy of event e is

$$E_e = \frac{\left(\text{SoC}_{r,k,p}^{\text{target}} - \text{SoC}_{r,k,p}\right) E_r^{\text{bat}}}{100} \tag{5}$$

For trips entering the modelled study area from outside, the initial SoC at the first within network route point is set to an entry value SoC^{in}, representing the possibility of charging prior to entering the modelled region.

Each charging event $e \in \mathcal{E}$ is recorded with (i) location, (ii) baseline event time $t_{e,0}$ corresponding to the arrival time at the trigger point prior to charging, and (iii) energy request E_e. The resulting per point trajectories (time and SoC) and the event set $\mathcal{E}$ provide the demand input to candidate generation and the LBBD-based placement and sizing model.

3.3. Candidate charging station identification

Candidate charging station locations are generated using two alternative approaches on a nationwide highway network. The first approach is

demand-driven and derives candidates from simulated charging events and traffic patterns. The second approach is regulation-driven and constructs a benchmark candidate set based on the Alternative Fuels Infrastructure Regulation (AFIR) minimum spacing requirements on the Trans-European Transport Network (TEN-T). AFIR, adopted by the EU in 2024, targets accelerated deployment of alternative-fuel infrastructure across Europe [33,34]. In this study, AFIR minimum requirements are used only to define the benchmark candidate set within the same national network. We thus do not restrict the analysis to TEN-T corridors.

Demand-driven candidates. The demand-driven candidate set is constructed by applying the k-means clustering algorithm to the spatial coordinates of the ideal charging events identified in Section 3.2. The centroid of each cluster is interpreted as a representative demand hotspot and included in the candidate charging station set $\mathcal{I}$. The purpose of using k-means is to group geographically close charging events into a limited number of representative locations, thereby reducing the dimensionality of the candidate placement problem while preserving the main spatial distribution of charging demand. In this way, areas with a higher concentration of ideal charging events are more likely to be represented by one or more centroids in the candidate set. Although the resulting centroids are not explicitly projected back onto the nearest road network node after clustering, the charging event locations are already projected to the road network nodes in the preprocessing stage. Therefore, the centroids are derived from road-aligned demand points and are generally located close to the motorway network rather than arbitrarily off-route. In addition, candidate feasibility is checked in the subsequent optimization stage, so locations that do not satisfy the route-based feasibility checks are not retained as feasible charging candidates. The number of clusters was selected through several rounds of trial and error to ensure that, after these steps, the candidate set remains sufficiently rich to represent the spatial distribution of charging demand while avoiding an unnecessarily large number of redundant locations that would increase the computational burden of the optimization.

Regulation-driven candidates. A regulation-based candidate set is constructed to represent a regulation-driven benchmark for corridor charging deployment. Since AFIR specifies corridor spacing requirements rather than exact station sites, candidate locations are generated on the road graph by solving a placement problem that minimizes the number of selected sites subject to a spacing-based coverage constraint. In the simulations, a uniform spacing threshold of $D^{\max} = 60$ km is adopted to generate the AFIR-based candidate set. This provides a simplified AFIR-compliant benchmark layout, after which the main planning optimization determines which of these candidates are actually developed and how much charging capacity is installed at each selected site. For brevity, this regulation-driven benchmark is referred to as AFIR in the remainder of the article, including the results section. The mathematical formulation used to generate the AFIR-based candidate set is summarized in Appendix A.

Event–station feasibility precomputation and pruning. After constructing the candidate set $\mathcal{I}$, event–station feasibility sets $I_e \subseteq \mathcal{I}$ are precomputed for each charging event $e \in \mathcal{E}$ to keep the planning model tractable.

A candidate station i is included in I_e if it can serve event e with an allowable detour from the driven trajectory. Operationally, this is enforced by:

- Detour proximity: i must lie within a predefined distance threshold of the event's driven path (i.e., the vehicle is allowed to deviate to a nearby station rather than charging exactly at the event anchor point).
- Directional/along-route consistency: i must be consistent with the event position along the driven route (using the route's cumulative distance/time coordinate), so that serving the event by i represents a plausible deviation along the trip progression rather than an unrelated location.

To further limit model size, the feasible set can be capped by retaining only the most relevant feasible candidates per event (top-k pruning). Finally, any candidate station that does not appear in at least one feasibility set (i.e., $i \notin I_e$ for all $e \in \mathcal{E}$) is removed from $\mathcal{I}$ prior to optimization.

3.4. Optimal charging station placement and sizing

A fully integrated time-indexed mixed-integer linear programming model would lead to a very large binary assignment structure, growing with the number of charging events, feasible stations per event, and admissible time slots, in addition to the combinatorial station-opening and charger-sizing decisions. To maintain tractability, the problem is solved using LBBD. The model couples (i) a design layer that selects charging station locations from a large candidate set and determines integer charger capacities subject to per-station grid-connection limits, and (ii) an operational layer that assigns charging events to feasible stations and admissible start times while enforcing time-dependent station capacity constraints and waiting time limits. To accelerate the iterative feasibility assessment, the scheduling subproblem is first evaluated in a relaxed LP form with station overload slack variables, which provides a fast infeasibility measure for cut generation. Once the design passes the operational feasibility check, a final rolling-horizon (RH) mixed-integer programming (MIP) model is solved for the fixed design to recover a fully discrete operational schedule.

Master problem (design). The master decides which stations are built, how many chargers are installed, and which events are selected to be served. Binary variable x_i indicates whether a station is built at i, integer variable c_i denotes the number of chargers installed at i, and binary variable y_e indicates whether event e is selected to be served. The master problem minimizes infrastructure cost,

$$\min \sum_{i \in \mathcal{I}} \left(C^s x_i + C^c c_i + C^g (P^{\mathrm{ch}} c_i) \right) \tag{6}$$

where P^{ch} is the charger power (kW) and C^s, C^c and C^g are the station, charger and grid capacity cost coefficients, respectively. Capacity is linked to station opening and bounded by the grid implied limit as

$$\underline{c}\, x_i \le c_i \le \bar{c}\, x_i, \qquad \forall i \in \mathcal{I} \tag{7}$$

Constraint (7) enforces $c_i = 0$ when $x_i = 0$ and, when a station is opened, restricts the installed chargers and grid connection capacity to lie between a scenario dependent minimum $\underline{c}$ and maximum $\bar{c}$. The lower bound $\underline{c} \ge 1$ is used to guide the design search and improve computational tractability.

A minimum service level is enforced only over events that have at least one feasible candidate, $\mathcal{E}^{\mathrm{feas}} = \{e \in \mathcal{E} : |I_e| > 0\}$,

$$\sum_{e \in \mathcal{E}^{\mathrm{feas}}} y_e \ge \alpha\, |\mathcal{E}^{\mathrm{feas}}| \tag{8}$$

An availability constraint prevents selecting an event as served without placing any local capacity,

$$\sum_{i \in \tilde{I}_e} c_i \ge y_e, \qquad \forall e \in \mathcal{E}^{\mathrm{feas}} \tag{9}$$

Constraint (9) is necessary but not sufficient for time-feasible operation; time feasibility is evaluated by the subproblem.

Subproblem (scheduling check with station overload slack). Given a master solution (x^*, c^*, y^*), the set of open stations is $\mathcal{I}^{\mathrm{open}} = \{i \in \mathcal{I} : c_i^* \ge 1\}$ and the set of served events is $\mathcal{E}^{\mathrm{srv}} = \{e \in \mathcal{E} : y_e^* = 1\}$. Operational feasibility is checked through an LP that allows fractional assignment. The LP uses the same capped feasibility sets as the master: for each event e, assignments are only allowed to stations in $\tilde{I}_e \cap \mathcal{I}^{\mathrm{open}}$, where $\tilde{I}_e \subseteq I_e$

contains the K nearest feasible candidates retained after preprocessing. Continuous variables $a_{e,t,i} \in [0,1]$ represent the fraction of event e assigned to station $i \in \tilde{I}_e \cap \mathcal{I}^{\text{open}}$ with start time $t \in \mathcal{T}_e$. Nonnegative slack variables $\delta_i \geq 0$ represent additional chargers required at station i to satisfy time-indexed capacity in the relaxed model.

The LP minimizes early charging penalties and overload slack (δ_i),

$$\min \sum_{e \in \mathcal{E}^{\text{srv}}} \sum_{i \in \tilde{I}_e \cap \mathcal{I}^{\text{open}}} \sum_{t \in \mathcal{T}_e} \phi_{e,i}\, a_{e,t,i} \;+\; \kappa \sum_{i \in \mathcal{I}^{\text{open}}} \delta_i \tag{10}$$

where $\phi_{e,i} \geq 0$ is a fixed early charging penalty. It is used to discourage premature charging before the preferred SoC-based charging trigger. Without this term, the model could compensate for limited infrastructure by assigning vehicles to charge earlier than operationally necessary, potentially creating multiple unnecessary charging stops. The penalty therefore helps prevent unrealistic charging behaviour that is induced by infrastructure scarcity rather than vehicle need. Thus, $\phi_{e,i} > 0$ only when charging occurs earlier than the preferred threshold; otherwise, $\phi_{e,i} = 0$. κ is the per charger slack cost (approximated by the marginal cost of one additional charger plus the associated grid capacity cost).

Each event is assigned according to the served decision,

$$\sum_{i \in \tilde{I}_e \cap \mathcal{I}^{\text{open}}} \sum_{t \in \mathcal{T}_e} a_{e,t,i} = 1, \qquad \forall e \in \mathcal{E}^{\text{srv}} \tag{11}$$

Charger capacity is enforced for each station i and time slot τ using a precomputed occupancy indicator $\chi_{e,t,\tau} \in \{0,1\}$. Charging time is represented through a constant power approximation: for each event e with required energy E_e (kWh), the effective charging power is

$$P_e^{\text{eff}} \;=\; \min\{\, P^{\text{ch}},\, P_v^{\max} \} \tag{12}$$

where $P_v^{\max}$ is the vehicle specific maximum charging power. The resulting charging duration is rounded up to the time slot length Δt (minutes) as

$$L_e \;=\; \Delta t \cdot \max\left\{ 1, \left\lceil \frac{60\, E_e}{P_e^{\text{eff}}\, \Delta t} \right\rceil \right\} \tag{13}$$

Then, for any candidate start time $t \in \mathcal{T}_e$, $\chi_{e,t,\tau} = 1$ if and only if the charger is occupied at time slot τ, i.e.,

$$\chi_{e,t,\tau} = 1 \quad \Longleftrightarrow \quad \tau \in \{t,\, t+\Delta t,\, \ldots,\, t + L_e - \Delta t\} \tag{14}$$

Accordingly, the time indexed capacity constraints become

$$\sum_{e \in \mathcal{E}^{\text{srv}}} \sum_{t \in \mathcal{T}_e} \chi_{e,t,\tau}\, a_{e,t,i} \leq c_i^* + \delta_i, \qquad \forall i \in \mathcal{I}^{\text{open}},\ \forall \tau \in \mathcal{T} \tag{15}$$

Tour level waiting budgets are imposed as hard constraints using the start time deviation from the event baseline start $t_{e,0}$ (in minutes),

$$\sum_{e \in \mathcal{E}_r \cap \mathcal{E}^{\text{srv}}} \sum_{i \in \tilde{I}_e \cap \mathcal{I}^{\text{open}}} \sum_{t \in \mathcal{T}_e} (t - t_{e,0})\, a_{e,t,i} \leq W_r^{\max}, \qquad \forall r \in \mathcal{R} \tag{16}$$

Design reinforcement cuts. The LP scheduling model in (10)–(16) acts as a diagnostic relaxation that evaluates whether the infrastructure design proposed by the master provides sufficient aggregate charging capacity. The overload slack variables $\delta_i \geq 0$ introduced in the time-indexed capacity constraints quantify the additional number of chargers required at station i to accommodate the currently selected served events under fractional assignment. Hence, a positive optimal value $\delta_i^* > 0$ indicates structural under-capacity of the infrastructure rather than a combinatorial scheduling issue.

Let $\mathcal{I}^+ = \{i \in \mathcal{I}^{\text{open}} : \delta_i^* > 0\}$ denote the set of overloaded stations. For each $i_0 \in \mathcal{I}^+$, an integer reinforcement increment is computed as

$$\Delta_{i_0} = \left\lceil \delta_{i_0}^* \right\rceil \tag{17}$$

which is scaled and capped to ensure gradual design evolution across LBBD iterations. This increment represents the estimated number of additional chargers required to eliminate overload in the relaxed scheduling model.

Two complementary design reinforcement cuts are added to the master problem in order to increase charging capacity while maintaining stable convergence:

- Centre expansion cut (local reinforcement): Capacity is increased directly at the overloaded station,

$$c_{i_0} \geq c_{i_0}^* + \Delta_{i_0}^{\text{cen}} \tag{18}$$

  where $\Delta_{i_0}^{\text{cen}} \leq \Delta_{i_0}$ allocates a portion of the required reinforcement locally. This cut addresses congestion concentrated at a specific station.
- Regional expansion cut (spatial reinforcement): Additional capacity is simultaneously encouraged within a neighborhood $U(i_0)$ of nearby candidate stations,

$$\sum_{i \in U(i_0)} c_i \geq \sum_{i \in U(i_0)} c_i^* + \Delta_{i_0}^{\text{reg}} \tag{19}$$

  where $\Delta_{i_0}^{\text{reg}} \leq \Delta_{i_0}$ distributes reinforcement across alternative locations. This prevents overload from simply shifting to adjacent stations after local expansion.

The parameters $\Delta_{i_0}^{\text{cen}}$ and $\Delta_{i_0}^{\text{reg}}$ are derived from the same overload estimate Δ_{i_0} and are chosen so that capacity increases progressively across iterations rather than through a single large adjustment. After adding these cuts, the master problem is re-optimized.

Operational feasibility verification. The LP scheduling model in Eqs. (10)–(16) certifies feasibility only for a relaxed assignment in which charging events may be fractionally distributed across stations and start times. Therefore, each incumbent master solution (x^*, c^*, y^*) is subsequently validated through a discrete operational scheduling check enforcing binary assignments.

The integer scheduling model adopts the same structure as the LP formulation, using the capped feasibility sets $\tilde{I}_e$, the assignment constraints (11), the time-indexed capacity formulation (15), and the tour-level waiting constraints (16), with the only modification that assignment variables become binary, $a_{e,t,i} \in \{0,1\}$, and overload slack variables are removed ($\delta_i = 0$). Thus, the integer model verifies whether the installed capacities c_i^* admit a physically realizable charging schedule.

Because nationwide instances involve a very large number of charging events, solving a single monolithic integer program is computationally intractable. Operational feasibility is therefore evaluated using an RH decomposition.

Let $\mathcal{E}^{\text{srv}} = \{e \in \mathcal{E} : y_e^* = 1\}$ denote the served events selected by the master. Events are ordered by arrival time and partitioned into sequential, partially overlapping temporal windows

$$\mathcal{E}^{(1)}, \ldots, \mathcal{E}^{(W)}, \qquad \bigcup_{w=1}^{W} \mathcal{E}^{(w)} = \mathcal{E}^{\text{srv}}$$

For each window w, the integer scheduling model is solved only over events in $\mathcal{E}^{(w)}$, while assignments determined in earlier windows remain fixed. Time-indexed charger capacity constraints are enforced by augmenting (15) with previously committed charger occupancy,

$$\sum_{e \in \mathcal{E}^{(w)}} \sum_{t \in \mathcal{T}_e} \chi_{e,t,\tau} a_{e,t,i} + \Gamma_{i,\tau}^{(w-1)} \leq c_i^*, \qquad \forall i, \tau \tag{20}$$

where $\Gamma_{i,\tau}^{(w-1)}$ represents cumulative charger usage fixed by preceding windows. After solving window w, feasible assignments are frozen and propagated forward through $\Gamma_{i,\tau}^{(w)}$.

If all RH windows are feasible, the constructed binary assignment defines a globally feasible charging schedule consistent with the master design. Otherwise, the failing windows identify integrality-driven conflicts that trigger feasibility restoration cuts in the subsequent LBBD iteration.

Feasibility restoration cuts. If at least one RH window is infeasible despite zero LP overload, the infeasibility originates from integrality effects or tight temporal coupling rather than insufficient aggregate charging capacity. In this situation, infrastructure expansion is not immediately enforced. Instead, a hierarchical feasibility restoration procedure is applied to recover a discrete schedule while minimally modifying the served event configuration. The overall decision logic is illustrated in Fig. 2.

Let $\mathcal{F} \subseteq \mathcal{E}^{\text{srv}}$ denote the set of events associated with failed RH windows. The restoration strategy proceeds through progressively stronger actions:

- Local RH repair (temporary drops): If a limited drop budget is available, the RH scheduler is first retried after temporarily removing a small subset of problematic events. Priority is given to (i) events assigned to stations without remaining grid connection headroom and (ii) events repeatedly involved in scheduling conflicts. These removals are local to the RH attempt and do not modify the master decision variables y_e. If the subsequent RH execution becomes feasible, the algorithm proceeds directly to schedule finalization (left branch of Fig. 2).
- Additional design reinforcement cuts: If RH infeasibility persists after local repair, the algorithm evaluates whether failing events remain connected to stations with available grid connection headroom. When such headroom exists, the infeasibility is interpreted as hidden capacity insufficiency not revealed by the LP relaxation. Additional design reinforcement cuts are therefore generated and the master problem is re-optimized. Temporary RH drops are discarded before returning to the master so that tentative scheduling adjustments are not permanently enforced.
- Persistent infeasibility handling: If no capacity headroom exists, infeasibility cannot be resolved through infrastructure expansion. Provided that the coverage requirement in Eq. (8) remains satisfied, a small subset $D \subseteq \mathcal{F}$ of persistently infeasible events is permanently removed from the served set through an events drop cut, forcing the master to select an alternative feasible configuration in subsequent iterations.

$$\sum_{e \in D} y_e \leq |D| - 1 \tag{21}$$

- Partial scheduling fallback: If none of the above restoration actions yields a fully feasible RH schedule, the algorithm does not terminate the LBBD iteration. Instead, a partial scheduling solution is accepted in which only the subset of events successfully assigned during the RH process is retained. Events that remain unscheduled are excluded from the finalized operational plan while preserving the infrastructure design determined by the master problem. This fallback mechanism ensures robustness of the planning framework by preventing algorithmic deadlock in highly congested scenarios.

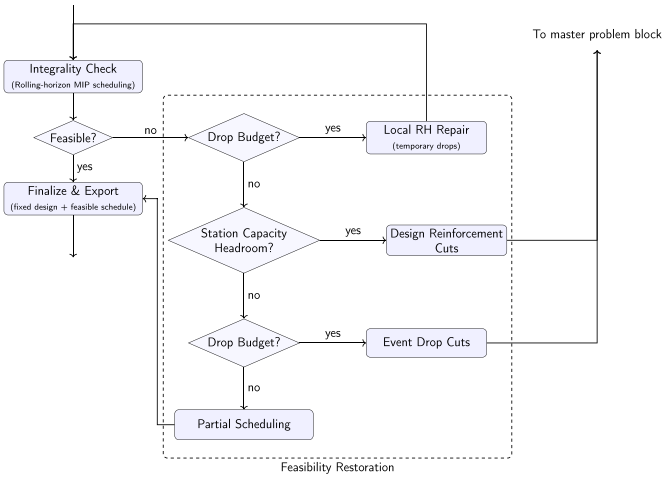


Fig. 2. Feasibility restoration policy applied after RH infeasibility.

Together, these feasibility restoration steps distinguish transient scheduling conflicts from structural design deficiencies and ensure that infrastructure expansion is invoked only when operational repair is insufficient, while guaranteeing that a feasible (possibly partial) operational schedule can always be constructed.

3.5. Charging demand profile reconstruction

The planning model determines both infrastructure design and operational charging decisions. Specifically, the optimization decides (i) which candidate stations are opened and how many chargers are installed, and (ii) for each charging event, whether it is served together with its assigned station and charging start time. These decisions define when and where charging takes place across the network. To evaluate the resulting infrastructure design in terms of station utilization, temporal load patterns, and grid impact, the optimized schedule must be translated into station-level power and energy time series. This subsection describes the reconstruction procedure used to obtain these spatiotemporal charging demand profiles.

Constant power charging profile. Within the LBBD optimization framework, charger occupancy is modelled using a constant power approximation (12). This representation allows the scheduling subproblem to remain linear and computationally tractable while capturing the key temporal interactions between charging events.

While this constant power representation provides a convenient abstraction for optimization, it neglects the reduction in charging power that occurs at high SoC levels during the constant-voltage (CV) phase of battery charging and therefore tends to underestimate session duration.

CC–CV charging profile. To obtain more realistic departure times and load profiles for the evaluation stage, the optimized charging schedules are post-processed using a CC–CV charging profile. This step is introduced to avoid the common simplification in charging infrastructure planning studies where vehicles are assumed to charge at constant rated power throughout the entire charging session. Such a constant power assumption can underestimate charging duration, charger occupation time, waiting time, and consequently charging delay, particularly at higher SOC levels where charging power is typically reduced.

Because publicly available measured power–SOC charging curves for HDEVs, especially under high power MCS charging, remain limited, the CC–CV profile used in this study is derived from measured passenger vehicle fast charging data [35]. The measured C-rate curve is represented by a piecewise-linear (PWL) approximation and applied only in the post-processing stage to estimate effective charging duration, delivered energy, departure time, and the resulting load profile.

To assess the uncertainty introduced by using a passenger vehicle proxy, an additional OEM-calibrated HDEV curve is constructed from publicly available manufacturer information on battery capacity, maximum charging power, and reported 20–80% MCS charging time [36,37]. The considered truck has a reported battery capacity of 621 kWh and megawatt charging capability of up to approximately 1 MW. The reported 20–80% charging time of approximately 30min is used as the calibration target. Since detailed HDEV charging profiles are not publicly available, the exact power–SOC shape remains uncertain; however, the main purpose here is to evaluate the effect of CV stage tapering on charging time and energy delivery rather than to model the detailed internal charging control strategy.

Fig. 3 compares the measured curve, the PWL fit, the OEM-calibrated HDEV estimate, and the constant 2C reference. For a 621 kWh battery, charging from 20 to 80% SOC corresponds to 372.6 kWh of delivered energy. The fitted measured profile requires approximately 24min to

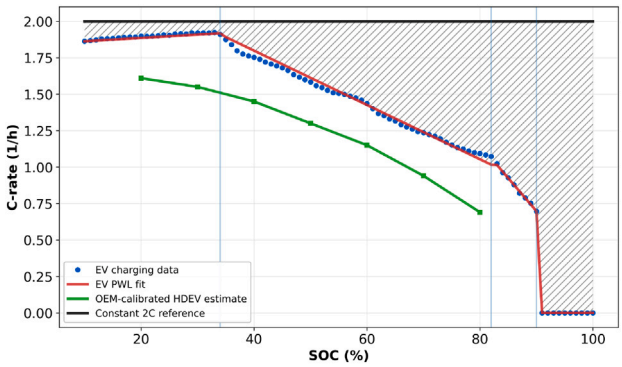


Fig. 3. Comparison of the measured personal vehicle CC–CV charging curve, an OEM-calibrated HDEV estimate, and a constant 2C charging reference.

deliver this energy, while the OEM-calibrated HDEV estimate requires 30min. In comparison, the constant 2C reference requires only 18min. The comparison indicates that the adopted measured profile provides a reasonable proxy for evaluating charging delay effects, while the OEM-calibrated HDEV curve provides a sensitivity check on this assumption.

4. Model setup and input data

This section summarizes the main data inputs, scenario assumptions, and numerical parameters used to instantiate the proposed planning framework. It describes the study area, freight tour dataset, vehicle class electrification mapping, scenario-year assumptions, infrastructure cost parameters, and key operational constraints. Together, these inputs provide the basis for generating charging demand, identifying candidate station locations, and optimizing the placement, sizing, and scheduling of HDEV charging infrastructure.

4.1. Study area and freight tour dataset

The framework is instantiated using a national scale application for the Netherlands. Freight activity inputs are obtained from the MASS-GT tour dataset described in Section 3.1. The dataset provides simulated truck tours on the Dutch road network with origin and destination coordinates, departure times, vehicle classes, and tour-level travel distances. These tours form the basis for constructing spatiotemporally resolved charging demand inputs along transport corridors. In total, the dataset comprises 3,062,056 trips and 1,286,080 unique carrier–tour combinations for a day.

4.2. Vehicle class electrification mapping

MASS-GT vehicle classes represent conventional diesel operations and therefore require an electrification mapping to specify representative battery-electric truck characteristics. For each vehicle class v, a baseline battery capacity E_v^{bat}, driving range, and energy consumption rate η_v are assigned using publicly-available manufacturer specifications and recent HDEV studies. Table 1 summarizes the class-level parameters used throughout the demand estimation and scheduling models. Tour-specific battery capacity and energy consumption parameters are sampled once per tour from truncated normal distributions centered at the class values, with a coefficient of variation of 10% for battery capacity and 5% for the consumption rate. The sampled values remain fixed for all trips within the same tour.

4.3. Electrification scenarios

To analyze the evolution of charging infrastructure needs over time, three scenario years are considered, corresponding to different stages of HDEV adoption. A 2025 baseline assumes an adoption rate of approximately 0.75% based on current market trends and vehicle registration statistics [43]. A medium-term outlook for 2030 assumes a 7.5% adoption rate under a policy-driven electrification trajectory. A longer-term scenario for 2040 assumes a substantially higher adoption level (approximately 27.5%), reflecting accelerated deployment of zero-emission trucks as regulatory targets and fleet turnover progress [44,45]. Where multiple values were available in the cited sources, a representative average percentage was used to define the scenario. Applying these adoption shares to the MASS-GT carrier–tour trajectories yields scenario-specific samples of electrified freight tours.

Table 2
Scenario-year planning limits used in the optimization model.

Year	Adoption rate	Grid limit per station (kW)	Min stations	Min chargers (350 kW / 1 MW)	Max chargers per station (350 kW / 1 MW)
2025	0.75%	2000	15	20 / 19	5 / 2
2030	7.5%	5000	20	100 / 50	14 / 5
2040	27.5%	10,000	35	350 / 150	28 / 10

In addition to demand growth, several infrastructure design limits are adapted across scenarios to reflect the expected evolution of charging technology and grid availability. These include the grid connection capacity per station, minimum infrastructure deployment levels, and charger power configurations used in the planning model. Table 2 summarizes the scenario-dependent parameters adopted in this study. The values reported in Table 2 were determined through several rounds of preliminary experiments and calibration runs to ensure that the optimization model remains computationally tractable while producing realistic infrastructure layouts. These values therefore represent baseline assumptions for the planning analysis. To evaluate the robustness of the results with respect to these assumptions, a series of sensitivity analyses is conducted in Section 5.

4.4. Cost parameters

The objective function (6) includes three infrastructure cost components: the fixed station construction cost C^s, the charger installation cost C^c, and the grid connection cost C^g. In the model, the fixed cost of opening a charging station is set to $C^s = 2000$ ke per site, while the grid connection cost is set to $C^g = 0.350$ ke/kW of installed grid capacity. Charger installation costs are expressed per unit of rated charger power, with $C^c = 500\ e/\text{kW}_{\text{rated}}$, corresponding to 175 k$e$ for a 350 kW charger and 500 ke for a 1 MW charger [46,47].

4.5. Key modelling parameters and assumptions

The main numerical parameters used in the charging demand generation, candidate identification, and planning model are summarized below.

- Charging trigger rule: Charging events are initiated when the simulated battery SoC falls below the charging trigger $\text{SoC}^{\text{trig}} = 20\%$ and the remaining trip distance cannot be completed while maintaining a reserve SoC of $\text{SoC}^{\text{res}} = 15\%$. The charging target is capped at $\text{SoC}^{\text{max}} = 80\%$, which avoids operating in the high-SoC region where charging power tapering leads to disproportionately longer charging times.
- Driving-speed assumption, routing resolution, and time discretization: For travel-time estimation, trucks are assumed to drive at an average speed of 80 km/h. Driven routes are resampled to a spatial resolution of $\Delta s = 2$ km to obtain a uniform trajectory representation. All time-dependent scheduling constraints are expressed using a discrete time step of $\Delta t = 15$ minutes.
- Driving time and rest break parameters: To approximate EU driving time rules in the event generation model, the maximum continuous driving duration is set to $H^{\text{drive}} = 4.5$ h and the break duration is set to $H^{\text{break}} = 45$ min, consistent with the requirement that a driver

must take a break of at least 45min after a driving period of 4.5 hours [48].

- Pre- and post-network trip buffers: For tours entering or exiting the modelled highway network, additional buffer distances are introduced to represent travel segments occurring outside the national study area. These segments account for cross-border freight movements that start or end in neighboring countries. To approximate charging opportunities outside the modelled region, it is assumed that a high-power charging station may be located at most $s_{r,k}^{\text{out}} = 60\,\text{km}$ from the border along the external corridor. This assumption is aligned with the AFIR-based spacing logic described in Section 3.3.
- Planning and scheduling parameters: The infrastructure planning model enforces a minimum service level of $\alpha = 0.9$, meaning that at least 90% of events with feasible candidate stations must be served. Tour-level waiting budgets are limited to a base value of $W_r^{\max} = 15$ minutes per charging event. This value is used as the reference case in the main analysis, while additional sensitivity experiments are conducted for alternative waiting time limits.
- Vehicle side charging power assumption: In the base case, the maximum charging C-rate of the truck battery is set to $2C$. This parameter defines the vehicle side charging power limit used in the effective charging power calculation.

5. Results and discussion

This section presents the outcomes of the optimization framework for the Netherlands under both demand-driven and AFIR-compliant planning approaches across the considered scenario years and charging powers. We report and discuss the resulting charging network configurations, service performance metrics, total system costs, and the trade-offs between demand-driven infrastructure planning and AFIR deployment requirements. In addition, sensitivity analyses are conducted for several key model parameters to examine their impact on infrastructure planning and operational scheduling outcomes.

5.1. Demand characterization on the modelled corridor network

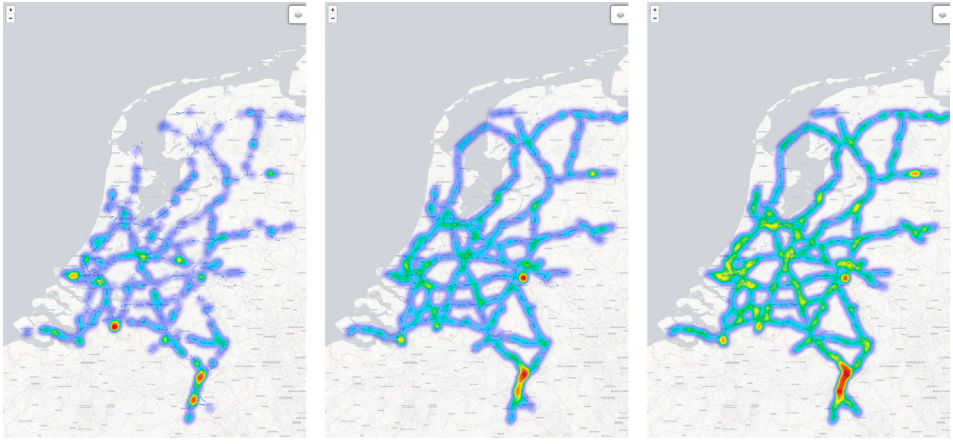
Before presenting the optimization outcomes, the spatial distribution of the synthesized charging demand is first visualized and the preprocessing used to restrict demand to the modelled road network. Since the scope of this article is highway en-route charging for HDEVs, the infrastructure planning domain is intentionally limited to motorways and national roads that represent the primary long-haul freight corridor network. It follows that charging events reconstructed from the tour- and SoC-based simulation are not guaranteed to fall on the modelled planning network. A network-consistency filtering step is therefore applied to retain only those events that can be associated with this corridor network under a predefined proximity criterion. Events classified as out of scope are excluded from subsequent corridor-based analyses, including candidate generation and event–station feasibility evaluation. A diagnostic visualization of this network-association step is provided in Appendix B.1.

Charging demand heatmaps are subsequently constructed using the retained, network-consistent event set, thereby visualizing demand intensity on the studied corridor network (Fig. 4). Demand concentrations are observed along major freight corridors and near cross-border gateways, with increasingly pronounced hotspots around high activity urban–industrial regions as adoption rises. The network-association step is used solely to anchor demand to the modelled network for corridor-based analysis and optimization; it does not imply that charging must occur exactly at the mapped event location. In the planning model, charging events may be served by stations within the vicinity of the trigger location, since detours are permitted through the event–station feasibility construction.

5.2. Demand-driven versus AFIR planning

This subsection compares optimized charging network designs obtained under the demand-driven candidate set and the AFIR candidates. Results are reported for the baseline operational setting ($W_r^{\max} = 15$ min and C-rate = 2C) and for two charger-power assumptions (350 kW and 1 MW). Table 3 shows that both approaches meet the imposed service target when performance is expressed relative to their own feasible event sets: the percentage of service level remains close to 90% across all scenario years and charger power assumptions. However, the two candidate generation approaches do not yield identical feasible event sets. In particular, the demand-driven candidate set achieves full event feasibility on the modelled corridor network (100%), whereas the AFIR benchmark consistently exhibits a small feasibility gap (approximately 98.5%). This feasibility gap is attributable to the fact that AFIR candidates are positioned primarily to satisfy corridor spacing requirements and are drawn from a more geographically constrained set of candidate locations, whereas demand-driven candidates are derived from demand hotspots and thus cover local demand with higher granularity.

Because the two candidate generation approaches yield different feasible event sets, the service level computed as scheduled/feasible events within each approach can mask differences in the absolute served demand volume. This effect is illustrated in Fig. 5, where the



(a) 2025 (b) 2030 (c) 2040

Fig. 4. Heatmap of simulated charging demand events for the three scenario years on the modelled road network in a day.

Table 3
Base-case planning results for 350 kW and 1 MW chargers.

(a) 350 kW chargers								
Year	Approach	Grid limit per station (MW)	No. of stations	Avg chargers per station	Avg grid connection per station (MW)	Avg waiting time (min)	Service level (%)	Infra cost (M€)
2025	AFIR	2	17	2.18	0.76	4.49	90.4	45.01
	Demand-driven		17	2.24	0.78	4.57	90.5	45.30
2030	AFIR	5	22	7.64	2.67	3.46	90.1	93.98
	Demand-driven		24	7.46	2.61	4.44	90.4	101.25
2040	AFIR	10	39	17.51	6.13	4.87	91.5	281.19
	Demand-driven		40	17.38	6.08	5.95	91.1	286.76
(b) 1 MW chargers								
Year	Approach	Grid limit per station (MW)	No. of stations	Avg chargers per station	Avg grid connection per station (MW)	Avg waiting time (min)	Service level (%)	Infra cost (M€)
2025	AFIR	2	17	1.12	1.12	7.66	90.1	50.15
	Demand-driven		17	1.24	1.24	7.74	90.2	51.85
2030	AFIR	5	23	3.43	3.43	4.78	90.0	113.15
	Demand-driven		26	3.15	3.15	6.64	90.2	121.70
2040	AFIR	10	36	7.92	7.92	6.58	90.0	314.25
	Demand-driven		39	7.51	7.51	7.01	90.8	327.05

scheduled share is additionally normalized by a fixed reference denominator: the feasible event set obtained under the demand-driven candidate set. Under this normalization, the AFIR designs satisfy only about 88.5–90.1% of the demand-driven feasible demand volume (depending on the year and charger power), even though they achieve approximately 90% scheduled/feasible within their own feasible sets. In contrast, the demand-driven designs remain close to the intended service target under the same reference (approximately 90.2–91.1%), reflecting both full feasibility and comparable scheduling performance. These results highlight that, for AFIR candidate sets with reduced feasibility coverage, reporting service level (scheduled/feasible) alone may overstate the effective demand satisfaction when compared against a common demand baseline.

Additionally, the base case results reveal a clear distinction between spatial deployment and capacity densification. Across scenario years, the number of stations increases only moderately, whereas the average number of chargers per station rises much more strongly (Table 3). This indicates that demand growth is absorbed primarily by reinforcing capacity at a relatively stable set of locations rather than by proportional expansion in the number of sites. The same pattern is visible in Fig. 6: the dominant charging station locations remain similar across years, but marker sizes increase substantially in the most demand-intensive corridor segments and gateway regions, reflecting progressive reinforcement at persistent hotspot locations. At the same time, stations are still deployed in lower-demand regions such as northern corridors, despite their weaker heatmap intensity in Fig. 4. This follows from the global service level requirement, which prevents feasible demand in lower-intensity regions from being systematically neglected. The generally smaller marker

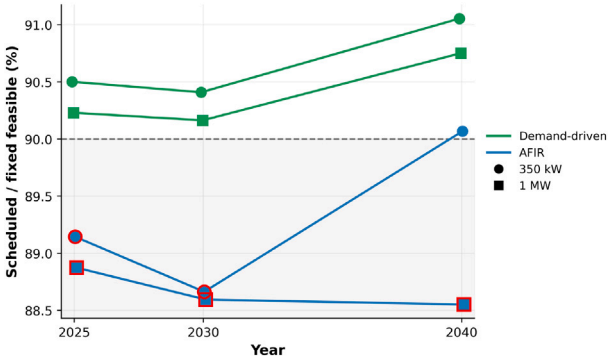
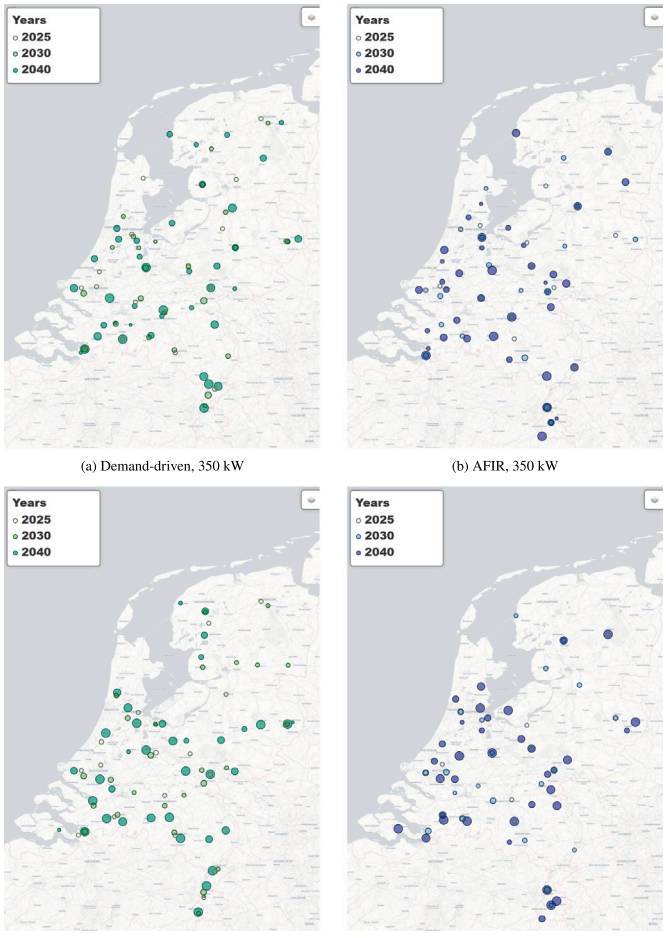


Fig. 5. Scheduled demand normalized by demand-driven feasible events.



(a) Demand-driven, 350 kW (b) AFIR, 350 kW

(c) Demand-driven, 1 MW (d) AFIR, 1 MW

Fig. 6. Optimized charging station layouts across scenario years. Marker size indicates installed grid connection capacity per station, marker colour indicates scenario year.

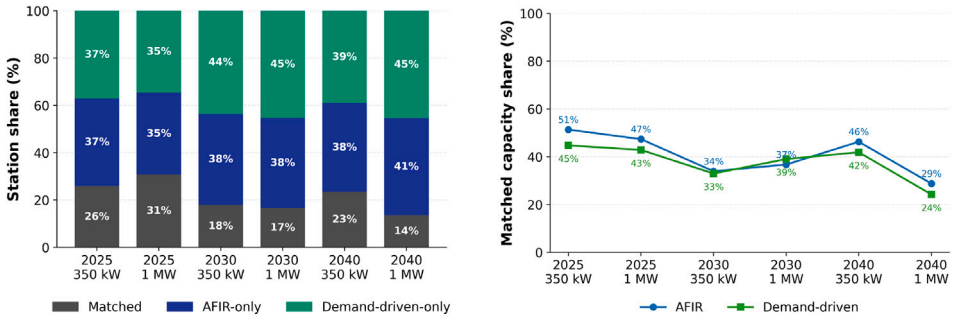


Fig. 7. Spatial similarity between AFIR and demand-driven base case layouts after spatial matching. (a) Overlap of the selected station sets, reported as shares of the combined station set. (b) Share of total installed capacity located at matched stations in each approach.

sizes in these areas indicate that only limited grid connection capacity is needed there. In contrast, high demand regions often evolve toward clusters of nearby large stations rather than a single dominant site, suggesting that local grid limits and charger-count constraints eventually bind and force additional nearby capacity to absorb further demand growth. The main difference between the 350 kW and 1 MW cases lies in station level parallelism rather than network geography. Because each 1 MW charger consumes a larger share of the available station connection capacity, fewer charging points can be installed within the same grid limit. This reduces flexibility in serving temporally overlapping demand and leads to consistently higher waiting times, especially in the higher-demand years.

To further compare the spatial similarity of the optimized layouts, Fig. 7 reports both the station set overlap after spatial matching and the share of total installed capacity located at matched stations. Stations are classified as matched when the selected AFIR and demand-driven sites fall within 5 km of each other; shares are computed relative to the union of selected stations. Fig. 7(a) shows that the matched share remains limited across all cases, ranging from 14% to 31% of the combined station set. The remaining stations are approach-specific, with AFIR-only stations accounting for 35–41% and demand-driven-only stations accounting for 35–45%. Thus, similar service levels and aggregate network performance are achieved through different spatial realizations of the charging network.

Although only a limited percentage of stations are spatially matched, Fig. 7(b) shows that these matched sites often account for a much larger share of installed capacity. Across cases, matched stations represent 24–51% of AFIR capacity and 24–45% of demand-driven capacity, indicating that both approaches converge on a small core of operationally important locations while differing more in secondary placement decisions. This contrast is strongest in 2040: under 350 kW charging, matched stations account for only 23% of the combined station set but already represent 46% of AFIR capacity and 42% of demand-driven capacity, whereas under 1 MW charging the corresponding shares fall to 14%, 29%, and 24%, respectively. This indicates greater spatial divergence when local capacity pressure is higher.

5.3. Operational characteristics of the optimized charging networks

Aggregate charging profiles. Fig. 8 compares the aggregate daily charging load of the optimized networks obtained with the AFIR and demand-driven candidate-generation approaches. The corresponding station-level variability is shown in Appendix B.4. Since the selected station sets are not generally co-located across the two approaches, aggregate network profiles provide the most consistent basis for operational comparison. Across all cases, both approaches produce very similar temporal profiles, characterized by lower loading during the early morning hours followed by a broad increase toward midday and the afternoon. This similarity holds even though the two approaches do not yield identical feasible event sets, scheduled event counts, total delivered energy, or unmet demand. The main difference is therefore not the temporal pattern itself, but the overall scale of the network load as freight electrification increases from 2025 to 2040. The total daily energy served grows strongly with year. In 2025, all cases remain close to 0.10 GWh/day, increasing to about 0.99–1.01 GWh/day in 2030 and 3.65–3.76 GWh/day in 2040. The aggregate peak load follows the same trend. For 350 kW charging, the peak rises from 8.10 to 9.06 MW in 2025 to 54.52–58.14 MW in 2030 and 218.35–223.04 MW in 2040. For 1 MW charging, the corresponding peaks are 10.20–10.87 MW, 61.75–64.43 MW, and 222.62–229.78 MW. These results indicate that the dominant operational challenge at the system-level is driven by the growth in charging demand over time, while the candidate generation approach has only a secondary influence on the aggregate temporal pattern.

The demand-driven solutions generally serve slightly more charging events and slightly more total energy than the AFIR solutions, which is reflected in marginally higher aggregate peaks in most cases. However, this difference remains small relative to the year-to-year growth. For example, in 2040 the total served energy differs by about 58 MWh between the 350 kW AFIR and demand-driven cases and by about 86 MWh between the corresponding 1 MW cases, while the profile shapes remain closely aligned. This suggests that the choice between AFIR and demand-driven candidates affects primarily the spatial allocation of infrastructure and the station-level distribution of load, rather than fundamentally changing the aggregate daily charging pattern.

A further observation is that, within the same year, the 1 MW cases tend to exhibit somewhat higher aggregate peaks than the 350 kW cases, even though the total served energy is similar. This is consistent with shorter charging durations and stronger temporal concentration of delivered energy when higher charger power is available. The effect is most visible in 2030, where the peak increases from 54.52 to 61.75 MW for AFIR and from 58.14 to 64.43 MW for the demand-driven case. By 2040, the difference remains present but becomes smaller in relative terms, as the overall scale of network demand dominates the profile.

Inter-charge driving durations and regulatory rest times. In order to relate the charging patterns to operational driving behaviour, Fig. 9 shows the distribution of driving hours since the previous charging event for all the charging events/trip start times in the 2040 case. Most charging events occur after approximately 3–4.5 h of driving, with the highest concentration below the 4.5 h threshold, which is broadly consistent with the EU rule requiring a break of at least 45 min after a driving period of 4.5 h. At the same time, the distribution exhibits a limited right tail beyond 4.5 h, indicating that not every modelled charging event occurs exactly at that threshold. Overall, the histogram supports the interpretation that a large

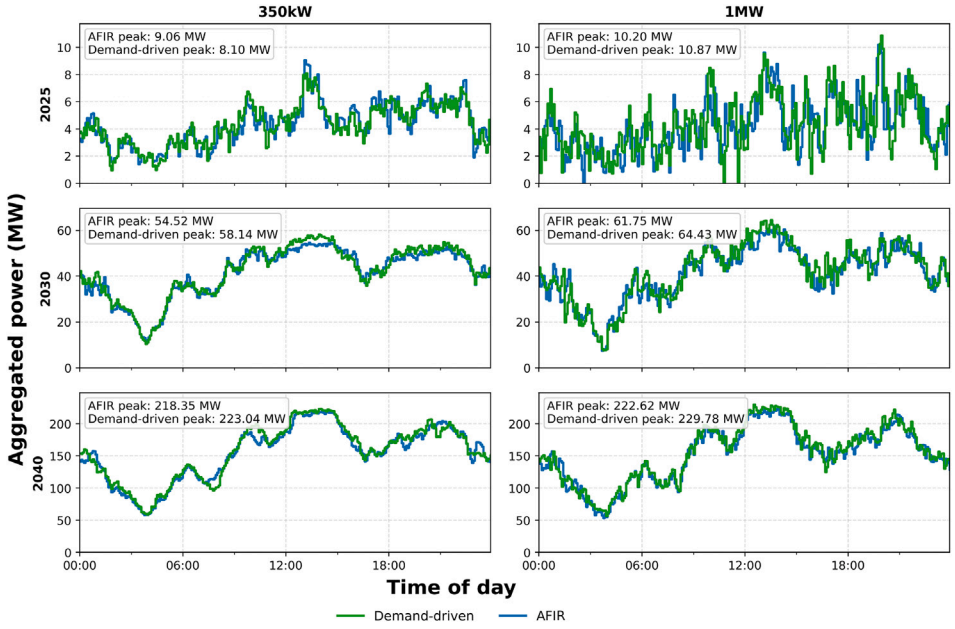


Fig. 8. Aggregate daily charging power profiles of the optimized networks for the AFIR and demand-driven candidate-generation approaches, for 2025, 2030, and 2040 under 350 kW and 1 MW charging assumptions.

share of charging demand is concentrated around practically relevant break intervals rather than being distributed arbitrarily along the route.

Effect of CC–CV charging on charging profiles. Building on the CC–CV charging model introduced earlier, this paragraph examines its effect on charging session duration relative to the constant power approximation. In the simulated charging strategy, vehicles typically begin charging at around 20% SOC and are not charged beyond 80% SOC. As a result, most charging events remain within the mid-SOC range, where power tapering is present but the strongest CV tail near full charge is not yet reached. The effect of the CC–CV formulation is therefore systematic but moderate.

Fig. 10 shows that the CC–CV representation increases the average session duration for all representative high-utilization stations under both the AFIR and demand-driven solutions. For 350 kW charging, the increase ranges from approximately 17% to 21%, while for 1 MW charging it ranges from approximately 20% to 24%. The consistently larger increase in the 1 MW cases is consistent with the C-rate sensitivity results: when charger power is higher, the vehicle-side charging trajectory becomes more influential, whereas under 350 kW charging the charger-side limit remains dominant over a larger portion of the session.

The absolute session durations show the same pattern. For the representative 350 kW stations, the average session time increases from about 25–29 min under the constant power approximation to about 30–36 min with CC–CV charging. For the representative 1 MW stations, the corresponding increase is from about 8.6–10.4 min to about 10.4–12.9 min. This has a direct implication for the assumed waiting time budgets. In the 1 MW cases, combining CC–CV charging with the base case waiting allowance of 15 min yields a total stop duration of roughly 25–28 min, while the 30 min waiting time sensitivity case implies about 40–43 min.

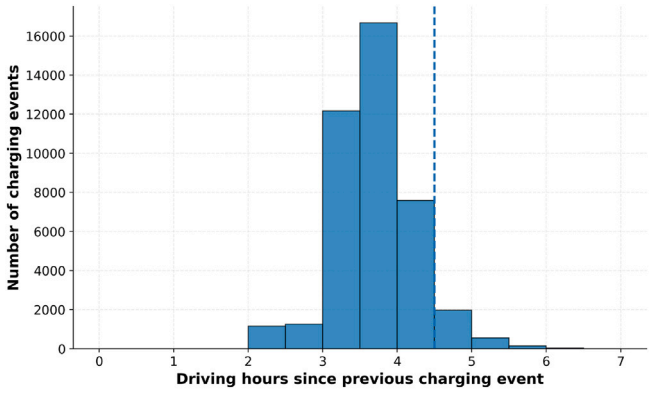


Fig. 9. Distribution of driving hours since the previous charging event in the 2040 case.

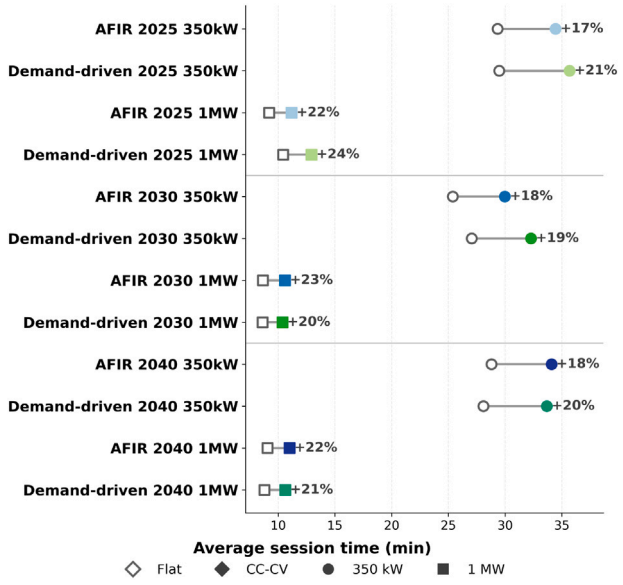


Fig. 10. Increase in average session duration between the constant power charging approximation and the CC–CV charging for representative high-utilization stations.

These values remain broadly compatible with the 45 min break horizon used in the driving-time model. By contrast, for 350 kW charging, the base case waiting allowance of 15 min already implies a total stop duration of roughly 45–51 min, so shorter dwell assumptions such as 30 min would generally not be sufficient. No incomplete CC–CV sessions were observed in the evaluated cases, indicating that the fitted charging curve fully covers the SOC range encountered in the simulations.

Overall, incorporating a CC–CV charging profile mainly elongates individual charging sessions rather than fundamentally changing the charging regime. This is operationally relevant because longer sessions increase charger occupation time and therefore reinforce, to some extent, the sensitivity of the infrastructure design to vehicle-side charging power assumptions, especially in the 1 MW cases. The results also show that lower power public charging is more strongly constrained by stop duration compatibility, whereas megawatt charging remains better aligned with practical en-route break windows.

5.4. Sensitivity analysis

The base case solutions are obtained under a set of operational, vehicle, and infrastructure assumptions that are uncertain in practice and may change as HDEV technology, charging infrastructure, and grid-connection conditions evolve. To evaluate the robustness of the planning results, a sensitivity analysis is performed for four key parameters. The first parameter is the vehicle energy consumption rate, which captures uncertainty in real world truck operation arising from payload, driving speed, ambient temperature, route characteristics, and auxiliary loads. The second parameter is the maximum allowable waiting time, which reflects the trade-off between service quality and operational flexibility. The third parameter is the maximum grid connection capacity per station, which represents different levels of local grid availability and possible grid-connection constraints. The fourth parameter is the effective vehicle charging power, expressed through the assumed battery charging C-rate, which determines charging duration and therefore charger occupancy. In all cases, the scheduled share remains close to the imposed 90% service target. Consequently, the sensitivity analysis mainly quantifies how infrastructure size, grid connection capacity, cost, and waiting time change in order to maintain the required service level.

Sensitivity to vehicle energy consumption. The energy demand of HDEVs is uncertain because real world consumption depends on operational and environmental factors such as payload, average speed, ambient temperature, road gradient, road conditions, and auxiliary loads. Recent real world analyses of HDEVs show that these factors can lead to substantial variation in observed energy consumption and charging behaviour [49,50]. Therefore, the effect of uncertainty in vehicle energy consumption is assessed as part of the sensitivity analysis.

To define realistic sensitivity ranges, the analysis draws on the real world HDEV study of Öko Institut [49]. In that study, the average energy consumption $\eta_{v,\mathrm{avg}}$ in kWh/km is represented as a function of average speed, ambient temperature, vehicle weight, and altitude difference:

$$\eta_{v,\mathrm{avg}} = m_1 \exp(-k_1 V_{\mathrm{avg}}) + m_2 T + m_3 W + m_4 A + m_5 \tag{22}$$

where V_{avg} is the average speed in km/h, T is the ambient temperature in °C, W is the vehicle weight in tonnes, and A is the altitude difference in metres. The reported coefficients are

$$m_1 = 5.1304, \quad k_1 \simeq 0.17, \quad m_2 = -0.0132, \quad m_3 = 0.0183, \quad m_4 = 0.0015, \quad m_5 = 0.7091 \tag{23}$$

Eq. (22) shows that lower temperatures, higher vehicle weights, and positive altitude differences increase the expected energy consumption. The speed effect is nonlinear because average speed enters the model through the exponential term; increasing V_{avg} reduces this term, with the strongest marginal effect at low speeds. Fig. 11 illustrates the corresponding change in predicted energy consumption relative to a base case

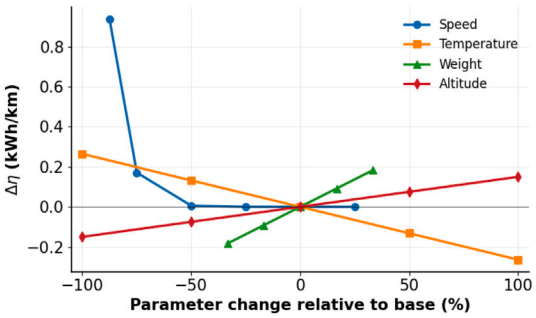


Fig. 11. Change in predicted energy consumption relative to the base case under representative operational and environmental variations.

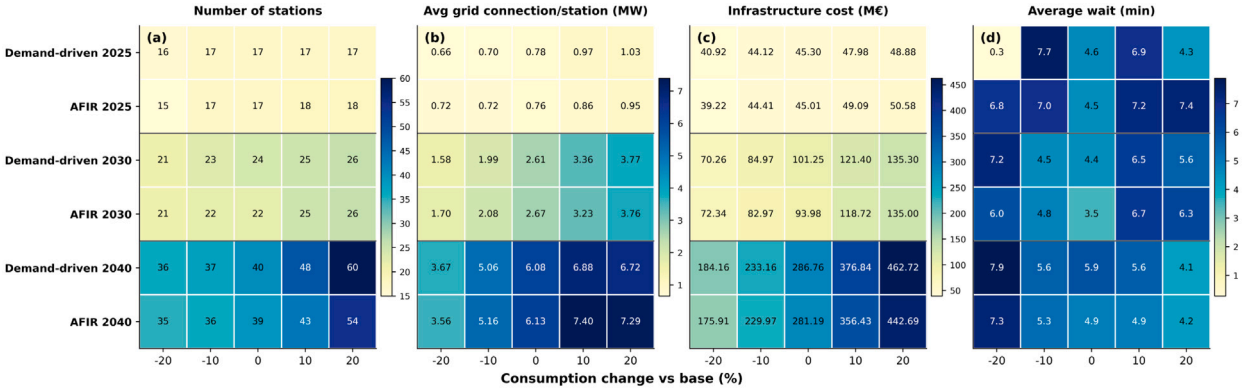
with $V_{\mathrm{avg}} = 80$ km/h, $T = 20$°C, $W = 30$ t, and $A = 0$ m. The figure shows that representative variations in temperature, weight, and altitude can each shift the predicted consumption by the order of 0.1–0.2 kWh/km, while the speed effect is comparatively small around higher cruising speeds but increases strongly in the lower speed range with much stop-and-go. These trends indicate that realistic operational and environmental conditions can have a non-negligible impact on HDEV energy demand and therefore justify the use of a bounded sensitivity range around the base case consumption values.

The base case simulations use fixed energy consumption rates, denoted by η_v for each vehicle type v. Since the trip level data do not consistently include all explanatory variables in Eq. (22), the empirical regression is used here to inform the sensitivity range rather than to replace the baseline formulation directly. Accordingly, the consumption sensitivity is implemented by varying the base-case vehicle-type consumption rates from −20% to + 20% around their nominal values.

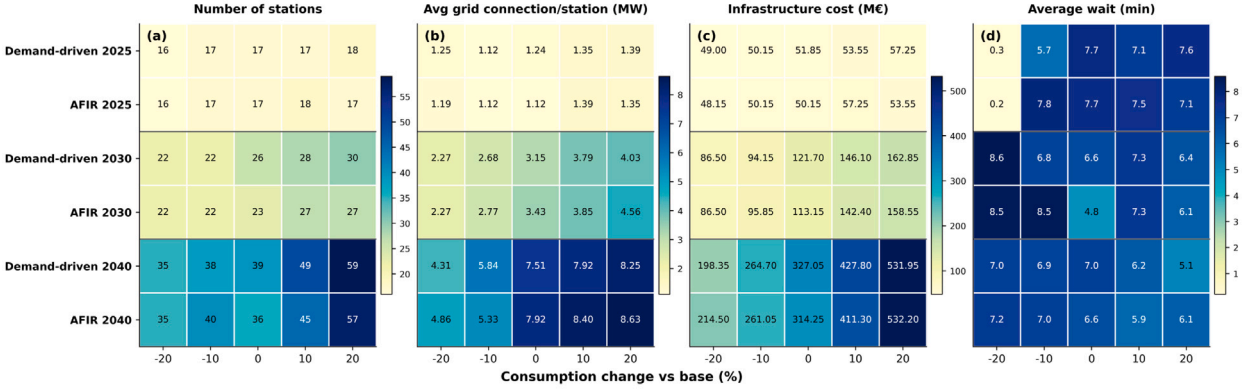
Fig. 12 shows the sensitivity of the planning outcomes to changes in the HDEV energy consumption rate for both 350 kW and 1 MW charger power assumptions. For both charging power assumptions, the sensitivity is limited in 2025. The number of selected stations remains almost unchanged across the consumption range, indicating that at low electrification levels the network design is mainly driven by spatial coverage and route feasibility rather than by the total energy volume. This means that reducing or increasing consumption does not substantially reduce the required number of stations, because a minimum spatial distribution is still needed to serve the freight routes.

The effect of consumption uncertainty becomes more visible in 2030 and is strongest in 2040. Higher consumption generally increases the number of selected stations, required grid connections, and total infrastructure costs. In the 1 MW demand-driven case in 2040, the number of stations increases from 39 in the base case to 59 under +20% consumption. For AFIR, the corresponding value increases from 36 to 57 stations. The 350 kW case shows the same qualitative trend: in the demand-driven 2040 case, the number of stations increases from 40 in the base case to 60 under +20% consumption, while the AFIR case increases from 39 to 54 stations. These results confirm that consumption uncertainty has a substantial effect on infrastructure scale in the high demand scenarios, regardless of charger power level.

The average grid connection per station also tends to increase with energy consumption, although the response is not strictly monotonic. This is because higher consumption can be accommodated either by reinforcing existing stations or by opening additional locations that distribute the demand more widely. For example, in the 1 MW 2040 demand-driven case, the average grid connection increases from 7.51 MW in the base case to 8.25 MW under +20% consumption. These results indicate that additional energy demand is not absorbed only by enlarging existing stations; the model also opens additional locations to distribute charging demand across the network.



(a) 350 kW chargers.



(b) 1 MW chargers.

Fig. 12. Sensitivity to changes in energy consumption.

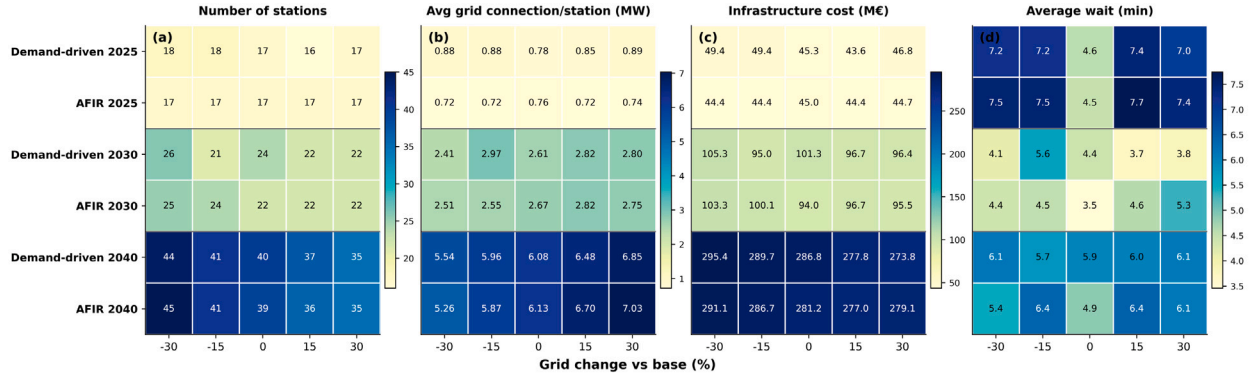
The waiting time results are less monotonic and this is expected, since waiting time is not minimized in the objective. In some cases, higher consumption increases waiting time due to stronger local capacity pressure. In other cases, waiting time decreases because the optimization responds to higher demand by opening more stations and distributing charging events over a wider network. For example, in the 1 MW 2040 demand-driven case, the average waiting time decreases from 7.0 min in the base case to 5.1 min under +20% consumption, despite the increase in total demand. Waiting time should therefore be interpreted as the combined result of demand intensity and the infrastructure expansion triggered by that demand.

Sensitivity to grid connection limits. Fig. 13(a) and (b) show the response to varying the maximum grid limit per station from −30% to +30% relative to the base case. Across both charger power assumptions, tighter grid limits generally force a more spatially distributed network: more stations are opened, the average grid connection per station decreases, and infrastructure costs increase. Relaxing the grid limit tends to have the opposite effect, allowing capacity to be concentrated at fewer stations with larger connections. However, this benefit saturates once the grid constraint is no longer strongly binding.

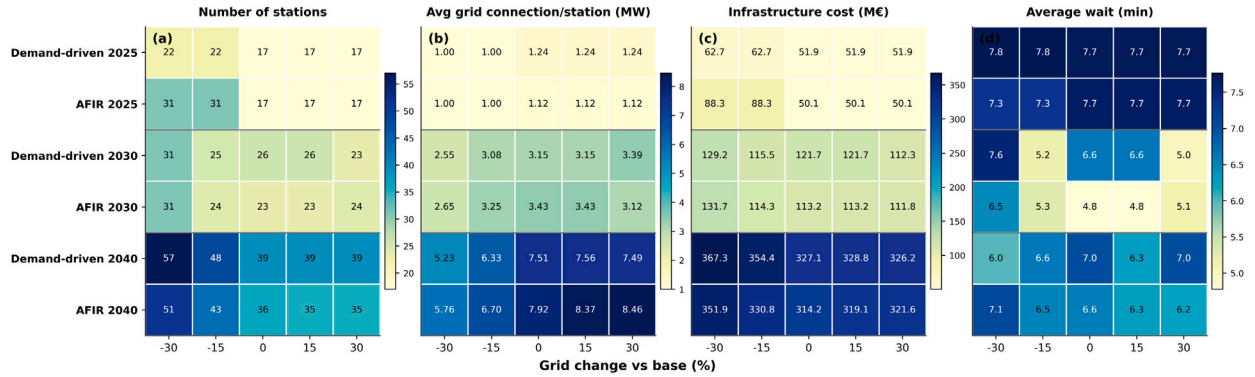
For 350 kW charging, the grid limit sensitivity is weak in 2025 and becomes more pronounced in 2030–2040. In 2025, the solutions are nearly insensitive, particularly for AFIR, where the station count remains fixed at 17 and the infrastructure cost stays close to 45 M€. By 2030, relaxing the grid limit allows capacity to be concentrated at fewer sites, reducing both station count and cost; this effect is clearer for the demand-driven candidate set than for AFIR. The strongest response appears in 2040, where tighter grid limits force a more spatially distributed network, while relaxed limits allow stronger consolidation. In the demand-driven case, the number of stations decreases from 44 to 35 as the grid limit increases from −30% to +30%, accompanied by an increase in average grid connection per station from 5.54 to 6.85 MW and a reduction in infrastructure cost from 295.4 to 273.8 M€. AFIR shows the same pattern, with station count falling from 45 to 35 and cost decreasing from 291.1 to 279.1 M€.

For 1 MW charging, the same mechanism is much stronger because each charger consumes a larger share of the available station connection. As a result, tighter grid limits translate more directly into additional station openings. This is already visible in 2025: at −30%, the station count rises from 17 to 22 in the demand-driven case and from 17 to 31 in the AFIR case, with corresponding cost increases from 51.9 to 62.7 M€ and from 50.1 to 88.3 M€. Above the base grid limit, however, the 2025 designs remain unchanged, indicating that the grid constraint no longer binds. The same pattern persists in 2030 and 2040: tighter limits increase both station count and cost, whereas grid relaxation allows some consolidation. However, the gains from relaxation saturate once the binding locations are relieved. This is most evident in the 2040 demand-driven case, where the station count remains 39 from the base case to +30%, and the infrastructure cost changes only marginally from 327.1 to 326.2 M€.

A further observation is that average realized waiting time is only weakly affected by the grid sweep and does not vary monotonically. This is expected, since waiting time is not minimized in the objective; once feasibility is achieved, additional grid headroom is primarily used to reduce or reconfigure infrastructure. The station-level charging point distributions reported in B.3 are consistent with this interpretation: tighter grid limits compress the distribution toward smaller sites, whereas grid relaxation allows a larger share of high capacity stations, especially in the 1 MW cases.



(a) 350 kW chargers.



(b) 1 MW chargers.

Fig. 13. Sensitivity to per station grid connection limits.

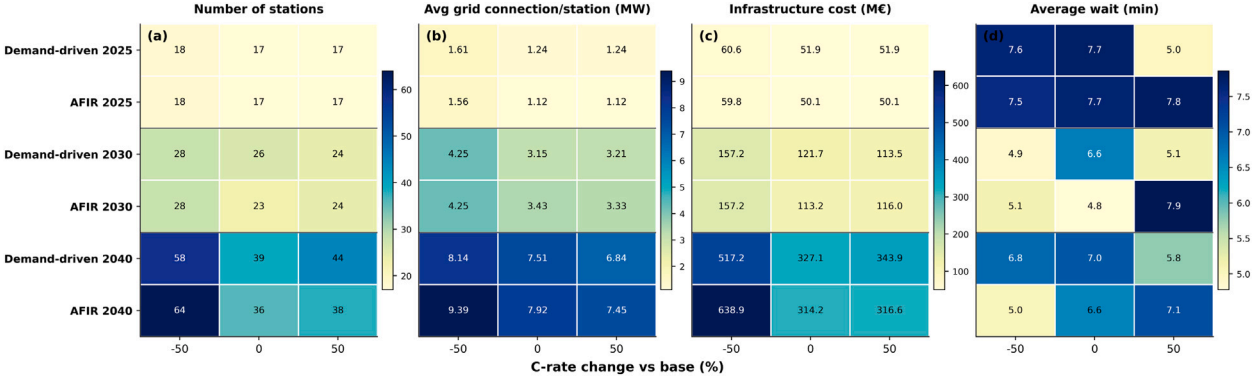


Fig. 14. Sensitivity to vehicle charging C-rate for 1 MW chargers.

Sensitivity to vehicle charging C-rate. The sensitivity analysis considers uncertainty in the vehicle-side charging limit by scaling the base C-rate of 2C by $\{-50\%, 0, +50\%\}$, i.e., 1C, 2C, and 3C. Recall from 12 that the assumed C-rate changes the charging duration and therefore the degree of temporal overlap among sessions.

The 350 kW solutions are effectively invariant across the C-rate sweep. Even under the lowest tested C-rate (1C), the implied vehicle maximum power remains at or above 350 kW for the battery capacities observed at charging events (Appendix B.2). Consequently, P_e^{eff} is almost always capped by the charger rating $P^{\text{ch}} = 350$ kW rather than by $P_v^{\max}$, so the optimized layouts remain essentially unchanged.

By contrast, the 1 MW cases are highly sensitive to the assumed C-rate, as shown in Fig. 14. Reducing the C-rate to 1C makes the vehicle-side limit binding, lengthens charging sessions, and increases concurrency, which in turn requires substantially more infrastructure. The effect strengthens with demand. In 2025, the response is still modest, with the station count increasing from 17 to 18 for both candidate sets and the infrastructure cost rising from 51.9 to 60.6 M€ in the demand-driven case and from 50.1 to 59.8 M€ in the AFIR case. In 2030, the increase is pronounced: both candidate sets require 28 stations at 1C, compared with 26 and 23 at 2C, while the infrastructure cost increases to 157.2 M€ from 121.7 M€ and 113.2 M€, respectively. In 2040, the low-C-rate case produces the strongest response of the entire sensitivity analysis: the demand-driven solution expands from 39 to 58 stations and from 327.1 to 517.2 M€, while the AFIR solution expands from 36 to 64 stations and from 314.2 to 638.9 M€.

Increasing the C-rate from 2C to 3C has a much smaller and less systematic impact. The battery-capacity distribution at charging events

(Appendix B.2) implies that for a large fraction of events $P_v^{\max}$ is already at or above 1 MW at 2C; therefore, P_e^{eff} remains capped by the charger rating and additional C-rate headroom yields limited further reductions in charging duration. In 2025, the designs remain unchanged. In 2030, the demand-driven case becomes slightly smaller and cheaper, whereas the AFIR case remains close to the base design. In 2040, the solutions reconfigure, but not in a uniformly improving direction: for example, the demand-driven case shifts from 39 to 44 stations and from 327.1 to 343.9 M€, while the AFIR case changes only slightly from 314.2 to 316.6 M€. Thus, the high C-rate case mainly alters the selection among nearby integer solutions, whereas the low C-rate case fundamentally changes the required network scale.

To complement the aggregate indicators in Figs. 14 and 15 the station level distribution of installed charging points for the 1 MW cases. The figure confirms that the 1C case is structurally different from the 2C and 3C cases. Under 1C, the distributions shift toward larger station sizes and, in the 2040 cases, a considerable share of stations clusters near the upper bound on installed charging points, indicating that longer charging sessions are absorbed not only through opening more stations but also through stronger capacity expansion at active sites. By contrast, the 2C and 3C distributions overlap substantially, consistent with the much smaller changes observed in station count and infrastructure cost.

Sensitivity to allowable waiting time. Fig. 16(a) and (b) show the effect of increasing the maximum allowable waiting time per charging event from 15 to 30, 45, and 60 min. Relaxing this limit enlarges the feasible scheduling region and allows the optimizer to trade temporal

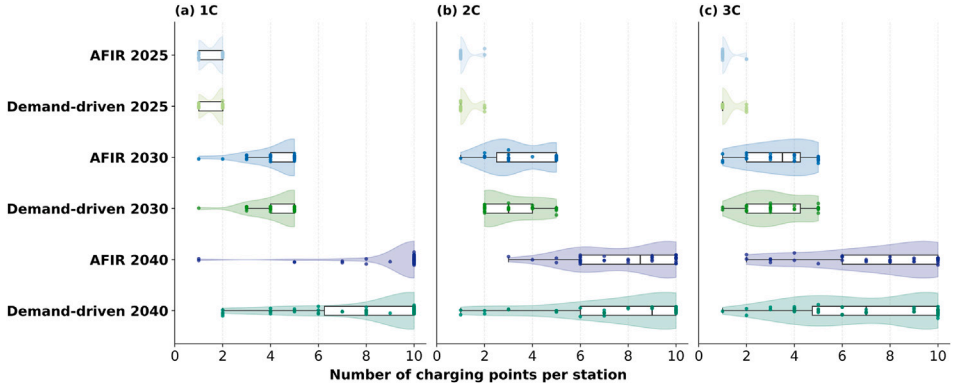
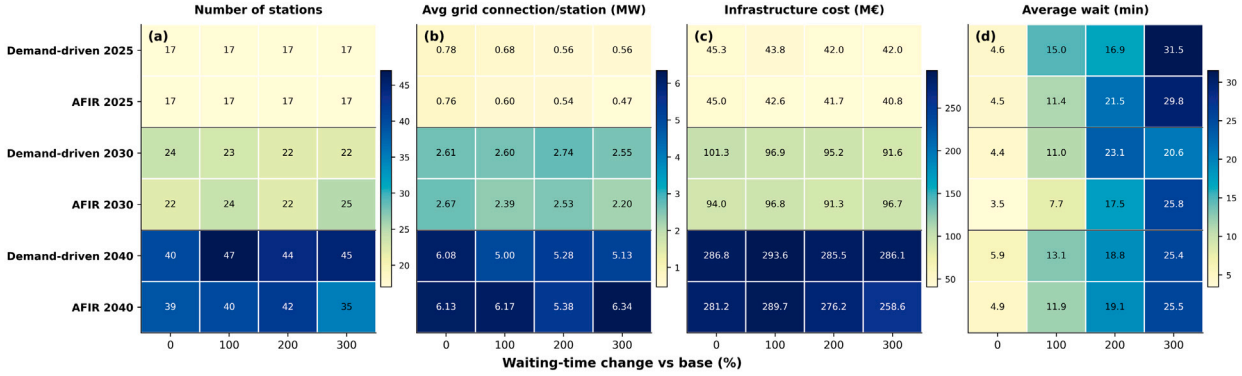
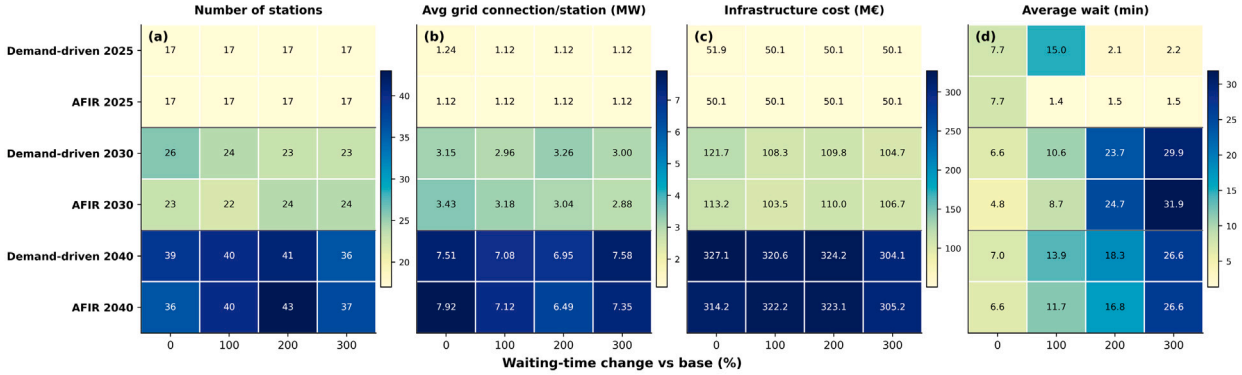


Fig. 15. Distribution of installed charging points per station under 1 MW charging for three C-rate assumptions.



(a) 350 kW chargers.



(b) 1 MW chargers.

Fig. 16. Sensitivity to allowable waiting time of vehicles.

flexibility against infrastructure. However, the response is not a simple monotonic reduction in station count: the additional flexibility is utilized through a combination of station resizing, station reallocation, and schedule reshaping, so infrastructure cost often falls even when the number of stations does not.

For 350 kW charging, the effect is limited in 2025 and mixed in 2030–2040. In 2025, station count remains fixed at 17 for both candidate sets, while infrastructure cost decreases from 45.3 to 42.0 M€ in the demand-driven case and from 45.0 to 40.8 M€ in the AFIR case. In 2030, the demand-driven case becomes somewhat smaller and cheaper as waiting flexibility increases, whereas the AFIR response is non-monotonic. In 2040, the demand-driven case shows little cost benefit despite changes in station count, while AFIR responds more favorably, with cost decreasing from 281.2 to 258.6 M€ and station count falling from 39 to 35 at the highest waiting allowance.

For 1 MW charging, the value of waiting flexibility is more pronounced, especially in 2030 and 2040. In 2025, the designs remain essentially unchanged at 17 stations and about 50 M€, indicating that the base infrastructure is already sufficient. In 2030, increasing the waiting allowance reduces infrastructure costs from 121.7 to 104.7 M€ in the demand-driven case and from 113.2 to 106.7 M€ in the AFIR case. In 2040, the same trend continues, with costs decreasing from 327.1 to 304.1 M€ and from 314.2 to 305.2 M€, respectively. Even so, station count remains non-monotonic, indicating that added waiting flexibility changes the cost-optimal network configuration rather than simply shrinking the network.

A key implication is that realized average waiting time is endogenous and not minimized directly. In medium- and high demand cases, it generally increases as the waiting limit is relaxed, reaching about 26–32 min in several 300% cases. In low-demand cases, it can also decrease because the optimizer selects a different feasible schedule or near-optimal design. Waiting time flexibility should therefore be interpreted mainly as an operational planning buffer that reduces or reshapes infrastructure requirements, rather than as a mechanism for improving driver experience. The charging point distributions in Appendix B.3 support the same interpretation.

5.5. Computational performance of the LBBD algorithm

Finally, to assess the tractability of the proposed decomposition approach, this subsection evaluates the computational performance of the LBBD algorithm across the main planning scenarios. All simulations were performed on a single machine with an Intel Xeon Platinum 8358P CPU running at 2.60 GHz and 16 GB RAM. The optimization models were solved using Gurobi version 12.0.2. Fig. 17 summarizes the distribution of the main computational times by demand year. For each year, the results are aggregated over 12 simulations, corresponding to two candidate generation approaches and two charger power assumptions. The figure presents smoothed probability density curves for the total LBBD runtime, the master problem solution time, the subproblem solution time, and the final integrality check / rolling-horizon scheduling time. These curves were constructed by applying kernel density estimation to the logarithm of the recorded runtimes and plotting the resulting densities against a logarithmic time axis, which is appropriate given the strong right skewness of the runtime distributions.

A clear shift in computational burden is observed as demand increases. In 2025, the total runtime remains concentrated at relatively low values, and the master problem accounts for the largest share of the computational effort. This is consistent with a median total LBBD runtime of about 503 s, a median of only 3 LBBD iterations, and 21 generated cuts. In these smaller and sparser instances, a substantial part of the effort is spent in the master problem to identify a cost effective combination of station locations and charger counts that satisfies the service requirement while avoiding unnecessary investment. In 2030, the total runtime remains moderate, with a median of about 226 s, but the distribution becomes broader and the contribution of the scheduling related stages becomes more visible. The median number of LBBD iterations increases to 5, with a median of 50 generated cuts.

The strongest change appears in 2040, where the total runtime distribution shifts towards much larger values and becomes substantially more dispersed. The median total LBBD runtime rises to approximately 4838 s, while the median number of LBBD iterations increases to 21 and the median number of generated cuts to about 127. In these high demand cases, the computational burden is no longer dominated by the master problem. Instead, it is driven mainly by the scheduling related stages, particularly the final rolling-horizon scheduling step and, to a lesser extent, the subproblem solution time. This reflects the fact that, as the number of charging events increases, it becomes significantly more difficult to assign and schedule each event under the coupled charger capacity, grid capacity, and waiting time constraints.

6. Conclusion

This article presents an integrated, operations-oriented framework for planning high power charging infrastructure for HDEVs. It combines tour-level freight demand simulation, route-level energy estimation, candidate site generation, and optimization of station placement and sizing under event-level scheduling constraints. As a result, charging infrastructure is evaluated not only in terms of spatial coverage, but also in terms of its ability to serve temporally-distributed charging demand under realistic operational constraints. Further, a CC–CV-based post-processing step was used to reconstruct charging durations, departure

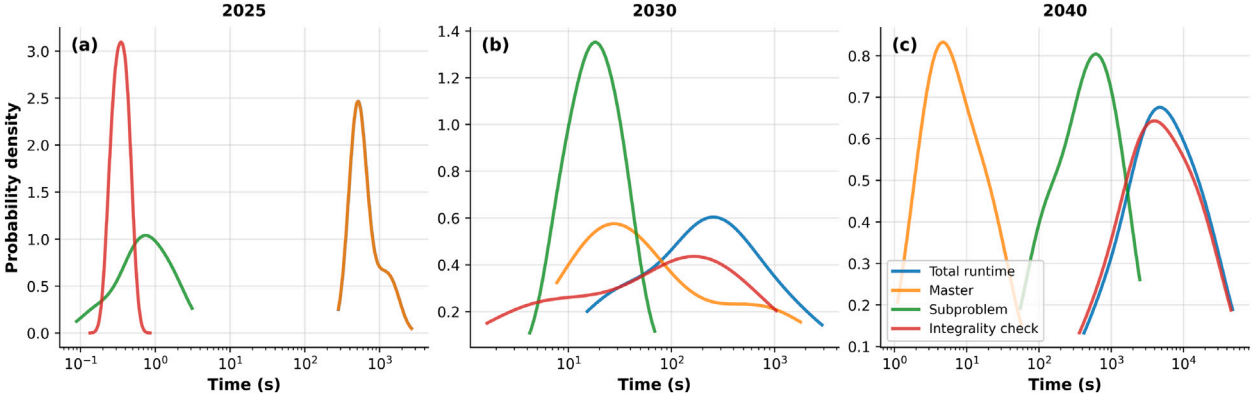


Fig. 17. Distribution of the main computational times of the LBBD algorithm for different demand years.

times, and aggregate station load profiles. In addition, sensitivity analyses on vehicle energy consumption, charging behaviour, grid limits, and waiting time flexibility were conducted to assess the robustness of the planning outcomes. The framework was applied to the Netherlands case for 2025, 2030, and 2040, considering both demand-driven planning and an AFIR-compliant regulatory benchmark, and for two charger power assumptions: 350 kW and 1 MW.

The results show that both planning paradigms can achieve high scheduled shares when evaluated relative to their own feasible event sets, but they differ in effective demand coverage and spatial flexibility. The demand-driven candidate set achieves full event feasibility on the modelled corridor network, whereas the AFIR-based candidate set exhibits a small but persistent feasibility gap because it is restricted to a more geographically-constrained set of corridor compliant locations. Consequently, when demand satisfaction is measured against a common reference denominator, demand-driven layouts consistently serve a larger share of the underlying charging demand. At the same time, AFIR-based solutions can still provide a workable corridor-oriented backbone and, in several base case scenarios, achieve comparable internal service levels with slightly lower infrastructure cost. The key distinction is therefore not whether AFIR-compliant layouts can support a national charging network, but how effectively they capture local freight demand heterogeneity and how much of the modelled charging demand is reachable from the available candidate locations.

A spatial comparison confirms this distinction between the two planning paradigms. Depending on the year and charger power, only about 14–31% of the combined selected station set is spatially matched between the AFIR- and demand-driven solutions. However, the matched subset often represents a much larger share of installed capacity, indicating that both approaches tend to agree on a limited number of strategically important high demand locations while diverging in the placement of complementary sites. This suggests that regulation-driven planning can provide a useful initial backbone, but that demand-driven refinement is important for improving feasibility coverage, capturing local hotspots, and supporting more efficient long-term infrastructure deployment.

Charger power strongly affects the trade-off between charging duration, parallel service capacity, and grid requirements. Relative to 350 kW charging, 1 MW charging substantially reduces the time needed to deliver the same amount of energy, improving compatibility with practical en-route stop and driver break durations. However, under the same per-station grid connection limit, 1 MW configurations allow fewer parallel charging points. This can lead to sharper aggregate power peaks and, in some cases, higher waiting times, despite similar total served energy. Megawatt class charging therefore improves vehicle side charging duration, but it also makes the network more sensitive to local grid constraints, charger availability, and vehicle side charging assumptions. The C-rate sensitivity confirms this effect: 1 MW networks are more affected by changes in the assumed vehicle charging capability, whereas the 350 kW solutions are largely insensitive to the C-rate sweep because the charger side power limit remains dominant.

Our operational analyses further show that AFIR and demand-driven layouts produce similar aggregate daily charging profile shapes, with differences driven mainly by the total amount of served demand rather than by fundamentally different temporal charging patterns. Charging events are concentrated around practically relevant driving intervals, with most events occurring after roughly 3–4.5 h of driving, consistent with the regulatory time horizon at which a driver break is required. Incorporating a CC–CV charging representation increases average charging session duration relative to a constant power approximation by about 17–21% for 350 kW charging and 20–24% for 1 MW charging. This increase remains bounded because the modelled sessions typically start near 20% SOC and do not extend beyond 80% SOC, but it nevertheless shows that vehicle side charging behaviour should not be neglected. After accounting for CC–CV charging, typical average charging times in the 1 MW cases remain around 10–13 min, so even with a 15 min waiting allowance, the total stop duration remains broadly compatible with practical en-route break windows. By contrast, for 350 kW charging, average charging times increase to about 30–36 min, meaning that a 15 min waiting allowance already results in total stop durations of roughly 45–51 min. This indicates that lower power public charging is more constrained by stop duration compatibility, whereas megawatt charging is better aligned with en-route rest opportunities.

Across the scenario years, the results indicate a two stage infrastructure development pattern that also explains why sensitivity effects become stronger over time. In the early adoption stage, especially in 2025, the model primarily establishes spatial coverage to ensure route feasibility, even when charging demand is still limited. Therefore, changes in vehicle energy consumption, grid capacity, C-rate, or allowable waiting time have only a limited effect on the overall network structure, because station selection is governed mainly by coverage needs rather than by capacity pressure. In later scenarios, once this basic coverage exists and charging demand increases, additional demand is served mainly by reinforcing already selected high demand locations. This is reflected in the stronger growth of the average number of chargers and grid connection capacity per station compared with the growth in station count. Consequently, the 2030 and especially 2040 scenarios become more sensitive to operational and technical assumptions, as local capacity, charger throughput, and grid connection limits become increasingly binding.

The different sensitivities affect the planning outcomes through distinct mechanisms. Vehicle energy consumption changes the total amount of charging demand that must be served, and therefore directly affects station count, grid connection requirements, and total infrastructure cost. Grid capacity reductions limit how much demand can be concentrated at each station, causing the model to distribute charging demand over more locations or to reinforce existing sites where possible. C-rate assumptions mainly affect the effective charging duration and charger throughput, with the strongest impact in the 1 MW cases, where vehicle side charging capability can become a limiting factor. By contrast, the 350 kW cases are less affected by C-rate variations because the charger side power limit remains dominant. The allowable waiting time affects the system through a different form of operational flexibility: relaxing this constraint can reduce infrastructure cost, particularly in 1 MW and high demand scenarios, because temporal flexibility can substitute for part of the required physical capacity. However, this effect is not monotonic in the number of selected stations and does not necessarily improve service quality, since waiting time is treated as a feasibility allowance rather than minimized in the objective function. Overall, the results show that energy demand, local grid capacity, vehicle side charging capability, and waiting time assumptions are all important for long-term HDEV charging infrastructure planning, but their relative importance depends on the considered planning outcome, charger power level, and demand year.

From a computational perspective, the LBBD algorithm remains tractable for the completed national scale planning instances, but the results also show that scalability becomes more challenging in high demand scenarios. The computational burden shifts from the investment-oriented master problem in smaller instances toward scheduling-related stages in the 2040 cases, where the number of charging events, generated cuts, and final rolling-horizon scheduling effort increase substantially. This indicates that the proposed decomposition approach is suitable for national scale planning experiments, while further acceleration of the scheduling layer may be needed for larger stochastic, multi-day, or higher-resolution extensions.

Overall, the findings indicate that AFIR-compliant planning can provide a robust corridor-oriented backbone for HDEV charging

infrastructure, while demand-driven planning improves effective demand coverage by adapting more closely to local freight demand patterns. The proposed framework therefore provides a practical decision support tool for policy-makers, infrastructure providers, and grid operators. It can quantify not only how much infrastructure is required, but also where it should be located, how sensitive the design is to uncertain operational assumptions, and how charging-network behaviour evolves as heavy-duty transport electrification scales up.

Future research can extend this framework in several directions. First, waiting time and driver-oriented service quality should be internalized directly in the objective function, so that infrastructure designs reflect both investment efficiency and user experience. Second, future work should validate the charging duration estimates against measured HDEV power–SOC curves, especially for MCS charging, once such data become publicly available. Third, next generation megawatt charging stations with shared power electronics and multiple dispensers should be represented explicitly, since they may alleviate the parallel service limitations observed in the 1 MW cases. Fourth, the planning model can be coupled more deeply with grid side flexibility, including battery storage, renewable generation, staged grid reinforcement, spatially varying grid connection costs, operational expenditure, discounting, and technology cost evolution. Finally, the framework could be expanded in scope to jointly consider public en-route and depot charging, stochastic demand and arrival uncertainty, multi-day demand variation, or cross border corridor planning under evolving European regulations.

CRediT authorship contribution statement

Leila Shams Ashkezari: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Bart Lagae:** Writing – original draft, Software, Methodology, Formal analysis, Data curation. **Neil Yorke-Smith:** Writing – review & editing, Validation, Supervision, Investigation, Formal analysis. **Mahnam Saeednia:** Writing – review & editing, Validation, Supervision, Investigation, Formal analysis. **Pavol Bauer:** Supervision, Resources, Project administration, Funding acquisition. **Gautham Ram Chandra Mouli:** Writing – review & editing, Validation, Supervision, Project administration, Investigation, Formal analysis.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered potential competing interests:

Pavol Bauer reports that financial support was provided by the Ministry of Economic Affairs and Climate Policy. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Generation of the AFIR-based candidate set

The AFIR-based candidate set is generated by solving a placement problem on the road graph to obtain a minimal set of candidate sites satisfying the adopted corridor-spacing requirement. Let $\mathcal{N}$ denote the set of road-network nodes and let $x_n \in \{0,1\}$ indicate whether node $n \in \mathcal{N}$ is selected as an AFIR-based candidate. The objective is to minimize the total number of selected nodes:

$$\min \sum_{n\in\mathcal{N}} x_n. \tag{24}$$

Coverage is enforced using two complementary constraints. First, a node-based coverage constraint requires every node to be within $D^{\max}$ of at least one selected node:

$$\sum_{i\in R_n} x_i \geq 1, \quad \forall n \in \mathcal{N}, \tag{25}$$

where

$$R_n \triangleq \left\{ i \in \mathcal{N} \mid d_{\mathcal{G}}(n,i) \leq D^{\max} \right\} \tag{26}$$

is the set of nodes within $D^{\max}$ of node n under the shortest-path distance $d_{\mathcal{G}}(\cdot,\cdot)$ on the road graph $\mathcal{G}$.

Second, a path-based coverage constraint is imposed over route segments extracted from simulated truck trajectories. Let $\mathcal{K}$ denote the set of trajectory-derived segments with length at most $D^{\max}$, and let $Y_k \subseteq \mathcal{N}$ denote the set of nodes visited along segment $k \in \mathcal{K}$. The corresponding path coverage constraint is

$$\sum_{i\in Y_k} x_i \geq 1, \quad \forall k \in \mathcal{K}. \tag{27}$$

In the present study, a uniform spacing threshold of $D^{\max} = 60$ km is used to generate the AFIR-based candidate set.

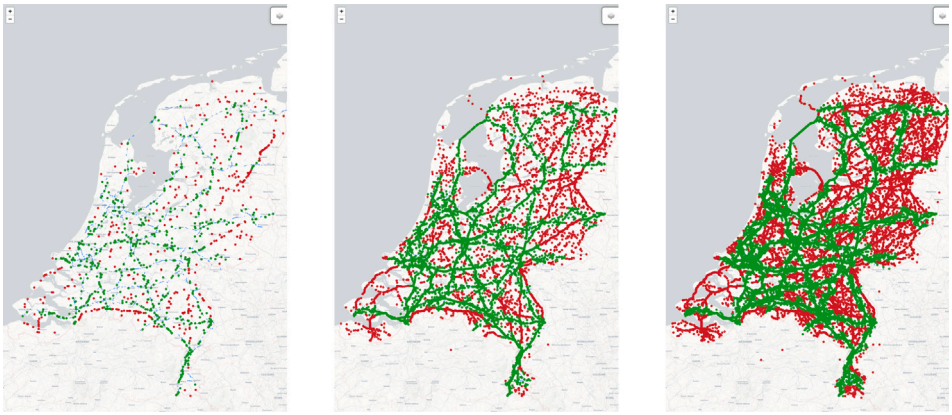
Appendix B. Supplementary results

B.1. Network association diagnostic for ideal charging events

This appendix provides a diagnostic visualization of the network association step used to restrict charging demand events to the modelled planning domain. Charging events that occur on road classes outside the studied corridor network are excluded from subsequent analyses. For transparency, Fig. 18 shows (i) charging events retained as network-consistent after association (green points), and (ii) events classified as out-of-scope (red points). Across the scenario years, a consistent share of reconstructed events is retained on the modelled network. In 2025, 2030, and 2040, 737/1115 (66.1%), 7330/11,381 (64.4%), and 26,773/41,529 (64.5%) events are retained, respectively. The stability of these proportions indicates that the corridor restriction removes a comparable share of out-of-scope demand across adoption levels while preserving the dominant spatiotemporal structure of on-corridor charging demand.

B.2. Battery capacity distribution at charging events

Fig. 19 reports the distribution of battery capacities associated with charging events (2040). This diagnostic supports the interpretation of the C-rate sensitivity results by indicating the typical battery sizes at which the vehicle limited charging power cap $P_v^{\max}$ is evaluated.



(a) 2025 (b) 2030 (c) 2040

Fig. 18. Visualization of the network association step used to filter charging events to the modelled road network.

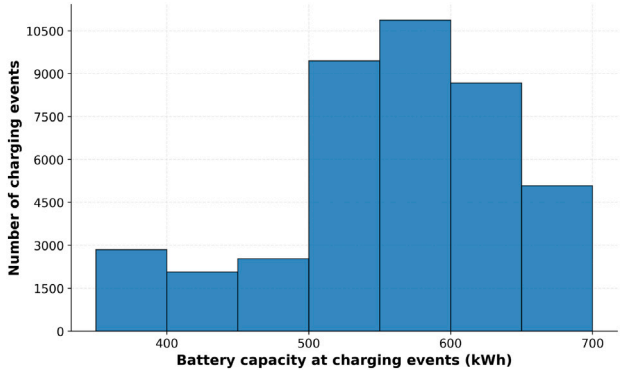


Fig. 19. Distribution of battery capacities associated with charging events in 2040.

B.3. Station-level charging point distributions

Charging point distributions for grid limit sensitivity. Fig. 20 complements the grid sensitivity heatmaps by showing how capacity is distributed across stations. For both charger power levels, tighter grid limits shift the distributions toward smaller station sizes, while relaxed grid limits allow a larger share of medium- and high-capacity stations. This effect is more pronounced for 1 MW charging, where the grid constraint directly limits the number of chargers that can be co-located at a site. The substantial overlap across neighboring cases also confirms that part of the response occurs through discrete re-selection among near-optimal integer designs rather than through a smooth scaling of all stations.

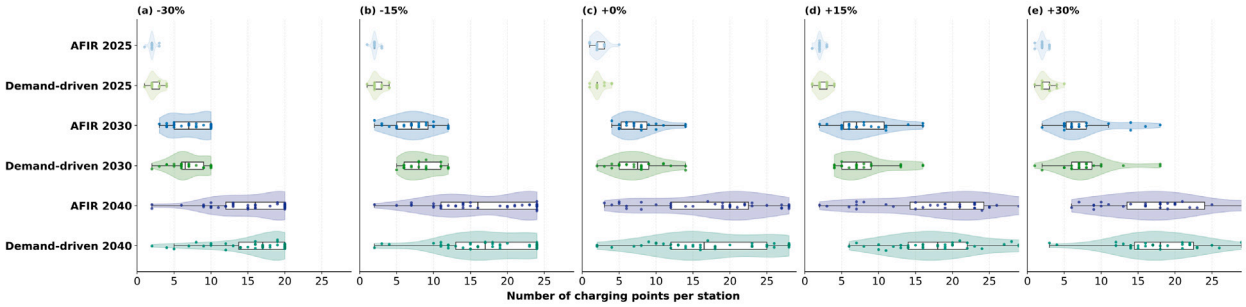
Charging point distribution for waiting time sensitivity. Fig. 21 provides the station-level view of the waiting time sensitivity. Increasing the allowable waiting time generally reduces the need for highly concentrated station sizing, particularly in the 1 MW 2030 and 2040 cases, where the upper tails become less dominant in several scenarios. At the same time, the distributions remain strongly overlapping and do not evolve monotonically across all cases, which is consistent with the heatmap results: additional temporal flexibility changes the cost-optimal network configuration, but not necessarily through a simple reduction in station count or uniform downsizing of every site.

B.4. Station-level charging profile envelopes

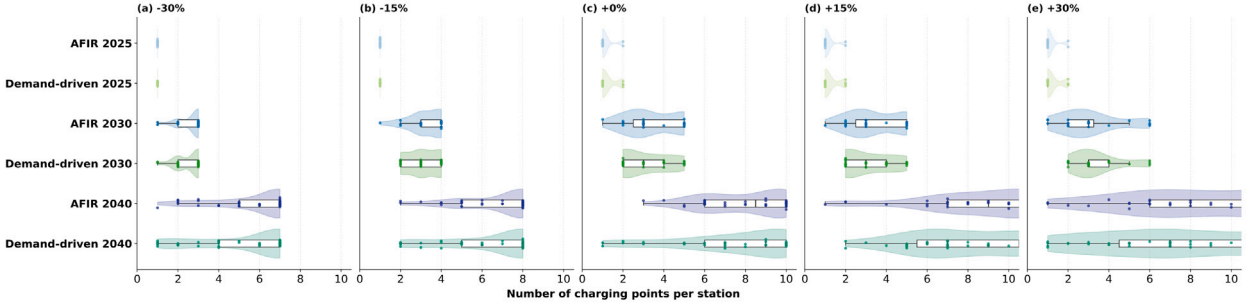
Fig. 22 complements the aggregate network profiles in the main text by showing the distribution of station-level loading patterns within each optimized network. For each case, the median profile summarizes the typical temporal behaviour across selected stations, while the 10–90% envelope quantifies the spread induced by differences in station size, local demand intensity, and temporal overlap of charging sessions.

Across all cases, the spread is widest in the low adoption year 2025, where station utilization is more uneven and a relatively small number of charging events can produce substantial variability in instantaneous power. In 2030 and especially in 2040, the station profiles become more structured and persistent, reflecting the emergence of a clearer hierarchy between lightly used stations and recurrent hotspot locations. The envelopes also indicate that the aggregate smoothness observed in the main text results from averaging over highly heterogeneous station-level behaviour.

The comparison between charger powers is also informative. The 1 MW cases generally exhibit broader short-term fluctuations at the station level, consistent with fewer parallel charging points and stronger sensitivity to the timing of individual sessions. By contrast, the 350 kW cases tend to produce smoother median station profiles, particularly in 2030 and 2040, because demand is distributed over a larger number of parallel charging points. Finally, the AFIR and demand-driven envelopes are broadly similar within each case, which is consistent with the finding in the main text that both approaches produce comparable aggregate temporal patterns even though their selected station sets differ spatially.

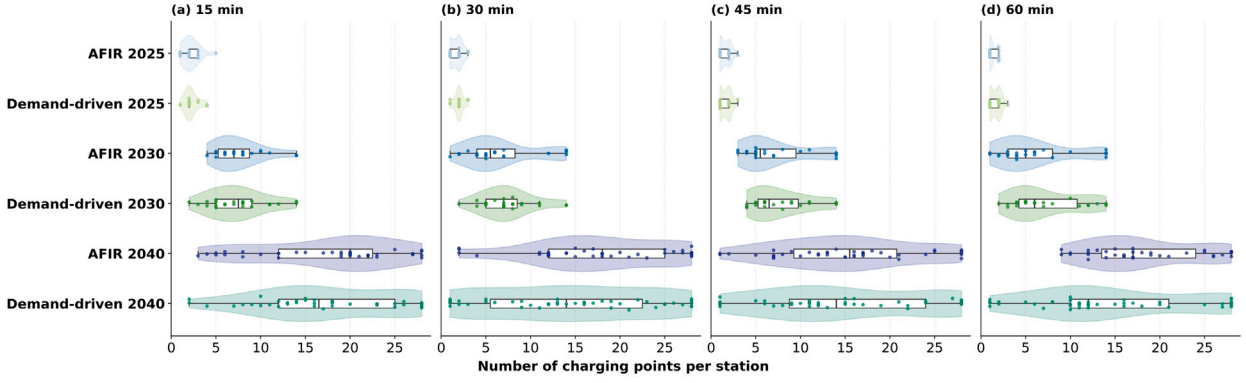


(a) 350 kW chargers.

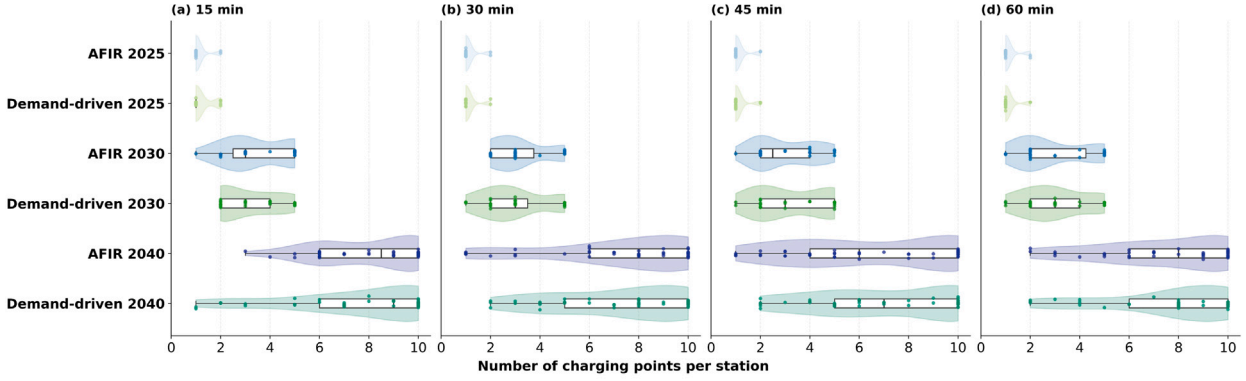


(b) 1 MW chargers.

Fig. 20. Station-level distribution of installed charging points under the grid limit sensitivity.



(a) 350 kW chargers.



(b) 1 MW chargers.

Fig. 21. Station-level distribution of installed charging points under the waiting time sensitivity.

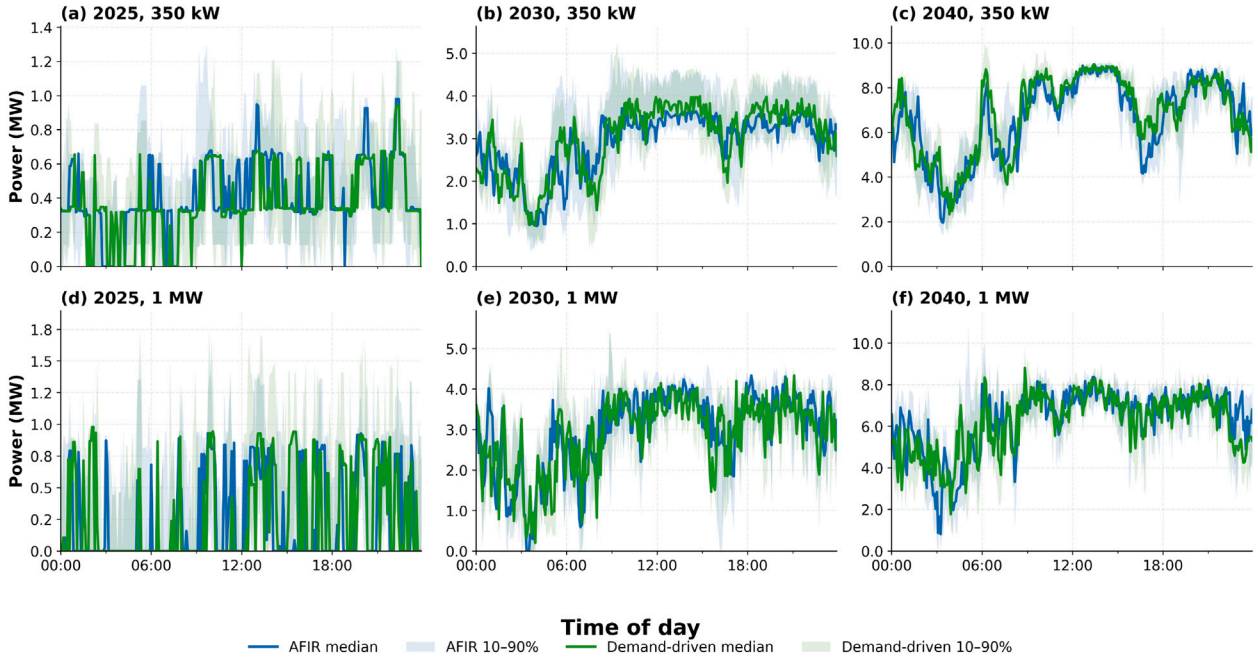


Fig. 22. Station-level charging profile envelopes for representative high-utilization stations from the optimized charging networks.

Data availability

Data will be made available upon request.

References

[1] The European Commission. A European green deal. 2020. https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal_en [Accessed 9 June 2024].

[2] IEA. Global CO_2 emissions from transport by subsector, 2000-2030. 2021. IEA, Paris, https://www.iea.org/data-and-statistics/charts/global-co2-emissions-from-transport-by-subsector-2000-2030 [Accessed 9 June 2023]. Licence: CC BY 4.0.

[3] The European Commission. Green deal: greening freight for more economic gain with less environmental impact. 2023. https://ec.europa.eu/commission/presscorner/detail/en/IP_23_3767 [Accessed 24 July 2024].

[4] Chen Y, Wu G, Sun R, Dubey A, et al. A review and outlook on energy consumption estimation models for electric vehicles. SAE Int. J. Sust. Trans. Energy, Env. Policy 2021;2(1):79–96. https://doi.org/10.4271/13-02-01-0005

[5] Diks C, Knapen L, van den Hoed R. Modeling electric vehicle charging infrastructure deployment: a review and research agenda. In: IOP conference series: earth and environmental science 1234, vol. 1; 2023. https://doi.org/10.1088/2634-4505/ac186a 012034.

[6] Xie F, Li Z, Zhang L, Wang Y. Spatiotemporal estimation of charging demand for electric heavy-duty trucks using large-scale GPS data. Appl Energy 2023;335:120713. https://doi.org/10.1016/j.apenergy.2022.120713

[7] Basma H, Schmidt J. Charging infrastructure needs for battery electric trucks in the European Union by 2030. Tech. rep., International Council on Clean Transportation; Oct 2025: https://theicct.org/publication/charging-needs-battery-electric-trucks-europe-oct25/ [Accessed 11 March 2026].

[8] Tong F, Wolfson D, Jenn A, Scown CD, Auffhammer M. Energy consumption and charging load profiles from long-haul truck electrification in the United States. Environ Res Infrastruct Sustain 2021;1(2):025007. https://doi.org/10.1088/2634-4505/ac186a

[9] Liimatainen H, van Vliet O, Aplyn D. The potential of electric trucks – an international commodity-level analysis. Appl Energy 2019;236:804–14. https://doi.org/10.1016/j.apenergy.2018.12.017

[10] Shoman W, Yeh S, Sprei F, Plötz P, Speth D. Battery electric long-haul trucks in Europe: public charging, energy, and power requirements. Transp. Res. Part D Transp. Environ. 2023;121:103825. https://doi.org/10.1016/j.trd.2023.103825

[11] Sauter V, Speth D, Plötz P, Signer T. A charging infrastructure network for battery electric trucks in Europe, working paper sustainability and innovation, no. s02/2021, fraunhofer institute for systems and innovation research ISI, provided in cooperation with fraunhofer institute for systems and innovation research ISI. 2021, https://www.isi.fraunhofer.de/en/publications/working-papers.html [Accessed 19 April 2025].

[12] Bernard ER, Lutsey N. Charging solutions for battery-electric trucks: infrastructure requirements and costs for early deployment. 2022. https://theicct.org/wp-content/uploads/2022/12/charging-infrastructure-trucks-zeva-dec22.pdf [Accessed 18 August 2025].

[13] Borlaug B, Moniot M, Birky A, Alexander M, Muratori M. Charging needs for electric semi-trailer trucks. Renew Sustain Energy Transit 2022;2:100038. https://doi.org/10.1016/j.rset.2022.100038

[14] Giuliano G, Dessouky M, Dexter S, Fang J, Hu S, Miller M. Heavy-duty trucks: the challenge of getting to zero, Transp. Res. Part D Transp. Environ. 2021;93:102742. https://doi.org/10.1016/j.trd.2021.102742

[15] Walz K, Rudion K. Charging profile modeling of electric trucks at logistics centers. Energies 2024;17(22):5613. https://doi.org/10.3390/en17225613

[16] Eriksson A, Pernestål A. Power demand and queuing at heavy truck semi-public charging points: forestry transport as a case study. Eur. Transp. Res. Rev. 2025;17(47). https://doi.org/10.1186/s12544-025-00744-7

[17] Karlsson J, Grauers A. Agent-based investigation of competing charge point operators for battery electric trucks. Energies 2024;17(12):2901. https://doi.org/10.3390/en17122901

[18] Speth D, Sauter V, Plötz P. Where to charge electric trucks in Europe—modelling a charging infrastructure network. World Electr Veh J 2022;13(9):162. https://doi.org/10.3390/wevj13090162

[19] Speth D, Plötz P, Funke S, Vallarella E. Public fast charging infrastructure for battery electric trucks—a model-based network for Germany. Environ Res Infrastruct Sustain 2022;2(2):025004. https://doi.org/10.1088/2634-4505/ac6442

[20] Hall D, Lutsey N. Estimating the infrastructure needs and costs for the launch of zero-emission trucks. Tech. rep., International Council on Clean Transportation (ICCT); 2019: https://theicct.org/wp-content/uploads/2021/06/ICCT_EV_HDVs_Infrastructure_20190809.pdf [Accessed 18 August 2025].

[21] Sujan V, Sun R, Jatana G, Fan J. Hydrogen and electric charging infrastructure for heavy-duty trucks: a nationally scalable megaregion assessment. SAE Tech Pap 2025. https://doi.org/10.4271/2025-01-8608

[22] Siekmann A, Sujan V, Uddin M, Liu Y, Xie F. Optimizing hydrogen fueling infrastructure plans on freight corridors for heavy-duty fuel cell electric vehicles. SAE Int J Sustain Transp Energy Environ Policy 2024;5(1):89–111. https://doi.org/10.4271/13-05-01-0008

[23] Hurtado-Beltran A, Rilett LR, Nam Y. Driving coverage of charging stations for battery electric trucks located at truck stop facilities. J. Transp. Res. Board 2021;2675(12):850–66. https://doi.org/10.1177/03611981211031542

[24] Speth D, Plötz P, Wietschel M. An optimal capacity-constrained fast charging network for battery electric trucks in Germany. Transp. Res. Part A: Policy Pract. 2025;193:104383. https://doi.org/10.1016/j.tra.2025.104383

[25] Dimitriou P, Nikolopoulou A, Gkiotsalitis K. An exact approach for the charging station location selection problem in urban freight transport. Oper. Res. 2025;25(39). https://doi.org/10.1007/s12351-025-00900-w

[26] Zahedi R, Sheinberg RL, Gowda SN, Chadha JR, SedghiSigarchi K, Gadh R. Phased planning of heavy-duty electric vehicle charging stations: an optimization frame-

work under grid capacity constraints. IEEE Access 2026;14:18316–31. https://doi.org/10.1109/ACCESS.2026.3659818
[27] Then J, Agalgaonkar AP, Muttaqi KM. Planning for medium- and heavy-duty electric vehicle charging infrastructure in distribution networks to support long-range electric trucks. Energies 2025;18(4):785. https://doi.org/10.3390/en18040785
[28] Borlaug B, Muratori M, Gilleran M, Woody D, Muston W, Canada T, Ingram A, Gresham H, McQueen C. Heavy-duty truck electrification and the impacts of depot charging on electricity distribution systems. Nat Energy 2021;6:673–82. https://doi.org/10.1038/s41560-021-00855-0
[29] Zhu X, Mather B, Mishra P. Grid impact analysis of heavy-duty electric vehicle charging stations. In: 2020 IEEE power & energy society innovative smart grid technologies conference (ISGT), 2020. p. 1–5, nREL/CP-5D00-74838. https://doi.org/10.1109/ISGT45199.2020.9087651. https://www.nrel.gov/docs/fy20osti/74838.pdf [Accessed 18 August 2025].
[30] Ye Z, Yu N, Wei R. Joint planning of charging stations and power systems for heavy-duty drayage trucks. Transp. Res. Part D Transp. Environ. 2024;134:104320. https://doi.org/10.1016/j.trd.2024.104320
[31] de Bok M, Tavasszy LA, Thoen S, Eggers L, Kourouniotti I. MASS-GT: an empirical model for the simulation of freight policies. Simul Model Pract Theory 2025;142:103140. https://doi.org/10.1016/j.simpat.2025.103140
[32] GraphHopper GmbH and contributors. Graphhopper directions API documentation: routing, 2025 https://docs.graphhopper.com/openapi/routing [Accessed 4 December 2025].
[33] Council of the European Union. Fit for 55, 05 Dec 2025. https://www.consilium.europa.eu/en/policies/fit-for-55/.; Accessed.
[34] TNO. Duiding van het afir-voorstel op de benodigde opbouw van tank- en laadinfrastructuur in Nederland. Tech. rep., TNO; 2021: https://www.rijksoverheid.nl/documenten/rapporten/2021/11/16/duiding-van-het-afir-voorstel-tno [Accessed 19 April 2025].
[35] Fastned. Mercedes-benz eqs suv: 107.8 kWh battery, CCS connector, 2025 https://www.fastnedcharging.com/en/brands-overview/mercedes-benz?model=EQS%20SUV [Accessed 20 August 2025].
[36] Mercedes-Benz Trucks. eActros 600: battery technology and charging specifications, 2026 https://www.mercedes-benz-trucks.com/int/en/trucks/eactros.html [Accessed 2 July 2026].
[37] Truck D. Megawatt charging in long-haul endurance test: Mercedes-benz trucks sends eActros 600 from Germany to Sweden, 2026 https://www.daimlertruck.com/en/newsroom/pressrelease/megawatt-charging-in-long-haul-endurance-test-mercedes-benz-trucks-sends-eactros-600-from-germany-to-sweden-53327018 [Accessed 2 July 2026].
[38] Trucks V. Volvo electric truck specifications, 2025 https://www.volvotrucks.com/en-en/trucks/electric.html [Accessed 1 March 2025].
[39] Mercedes-Benz Trucks. eActros 600 configurator, 2025 https://www.mercedes-benz-trucks.com/int/en/trucks/eactros-600.html [Accessed 19 April 2025].
[40] National Renewable Energy Laboratory. Electric vehicle Fleet deployment guide, 2025 https://www.nrel.gov [Accessed 19 April 2025].
[41] European Cyclists' Federation. LEVV-LOGIC: light electric freight vehicles in urban logistics, 2025 https://ecf.com [Accessed 19 April 2025].
[42] International Transport Forum. Micromobility: policy insights and data, 2025 https://www.itf-oecd.org [Accessed 19 April 2025].
[43] Rijksdienst voor Ondernemend Nederland (RVO). Duurzame mobiliteit databank: heavy duty vehicles, 2026 https://duurzamemobiliteit.databank.nl/mosaic/en-us/elektrisch-vervoer/heavy-duty-vehicles [Accessed 2 July 2026].
[44] Global Drive to Zero. CALSTART, global roadmap for reaching 100% zero-emission medium- and heavy-duty vehicles by 2040, 2022 https://globaldrivetozero.org/site/wp-content/uploads/2022/06/Global-Roadmap-for-Reaching-100-Zero-Emission-Medium-and-Heavy-Duty-Vehicles-by-2040_63022.pdf [Accessed 18 August 2025].
[45] ElaadNL. Outlook logistiek – update 2025, 2025 https://elaad.nl/wp-content/uploads/downloads/Outlook_Logistiek_2025Q1_def.pdf [Accessed 2 July 2026].
[46] Stedin Netbeheer BV. Elektriciteit tarieven 2024 grootverbruik, 2024 https://www.stedin.net/-/media/project/online/files/tarieven/tarieven-2024/tarieven-2024-elektriciteit-grootverbruik.pdf [Accessed 11 March 2026].
[47] Fröde P, Lee M, Sahdev S. Can public EV fast-charging stations be profitable in the United States? Oct 2023 https://www.mckinsey.com/features/mckinsey-center-for-future-mobility/our-insights/can-public-ev-fast-charging-stations-be-profitable-in-the-united-states [Accessed 11 March 2026].
[48] European Parliament and Council of the European Union. Regulation (EC) no 561/2006 of the european parliament and of the council of 15 March 2006 on the harmonisation of certain social legislation relating to road transport, consolidated version consulted on EUR-lex, Article 7; 2006 https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:02006R0561-20240522 [Accessed 11 March 2026].
[49] Corguillé JL, Hacker F, Dolinga T. Real-world data analysis of battery electric trucks operating in Germany. Tech. rep., eVS38 International Battery, Hybrid and Fuel Cell Electric Vehicle Symposium; Öko-Institut; 2025: https://www.oeko.de/fileadmin/oekodoc/Real-world-data-analysis-of-battery-electric-trucks-operating-in-Germany.pdf [Accessed 2 July 2026].
[50] Mao S. Real-world performance of battery electric heavy-duty vehicles in China: energy consumption, range, and charging patterns. Tech. rep., International Council on Clean Transportation; 2023: https://theicct.org/publication/hdv-china-real-world-performance-apr23/ [Accessed 2 July 2026].