TU Delft

Delft University of Technology

Document Version
Final published version

Licence
CC BY

Citation (APA)
Ramos, V. M. J., & Atasoy, B. (2026). Multi-level optimization for the deployment of battery swapping in inland waterways. *Transportation Research Part E: Logistics and Transportation Review*, *216*, Article 105258. https://doi.org/10.1016/j.tre.2026.105258

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.

Transportation Research Part E 216 (2026) 105258

Contents lists available at ScienceDirect

Transportation Research Part E

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

Multi-level optimization for the deployment of battery swapping in inland waterways

V.M.J. Ramos *, Bilge Atasoy

Department of Maritime and Transport Technology, Faculty of Mechanical Engineering, Delft University of Technology, Delft, 2628CD, The Netherlands

ARTICLE INFO

Keywords:
Inland water transport
Battery swapping
Multi-level optimization
Subsidy design
Infrastructure deployment
Zero-emission shipping

ABSTRACT

Containerized battery swapping offers a promising pathway to zero-emission inland waterway transport. However, its deployment requires coordinated decisions among actors with conflicting objectives. We introduce a multi-level optimization model with three actors: a policymaker who allocates subsidies to maximize adoption, an infrastructure developer who selects station locations to maximize profit, and fleet operators who minimize operating costs. We formulate this as a multi-level optimization problem, reduce it to a bilevel program via Karush–Kuhn–Tucker conditions, and solve it using an iterative bilevel decomposition and a problem-specific large neighborhood search. Experiments on the Dutch–Belgian inland waterway corridor show that a market equilibrium cannot exist without public intervention, the optimal subsidy mix depends on the available budget, and that within an intermediate price range the minimum budget for full adoption remains stable. Sensitivity analyses on costs, battery range, and network size are reported together with an alternative emissions-reduction objective.

1. Introduction

As global commitments toward climate-neutral mobility intensify, decarbonizing the transport sector has become a critical priority. Despite substantial progress in reducing emissions across other sectors, transport remains one of the most challenging sectors to decarbonize. In the European Union, transport accounted for approximately 31% of total greenhouse gas (GHG) emissions (European Environment Agency, 2025). Despite its advantages in energy efficiency and lower emissions intensity, inland water transport (IWT) remains underutilized, accounting for only a 5.1% modal share in Europe (Central Commission for the Navigation of the Rhine, 2024). Recognizing its potential, policy initiatives such as the European Green Deal and the Smart and Sustainable Mobility Strategy emphasize both accelerating transport decarbonization and increasing the role of inland shipping (European Commission, 2019, 2020). Consequently, expanding the IWT fleet while simultaneously transitioning it to zero-emission (ZE) propulsion has become a strategic priority for achieving long-term climate objectives. Recent advances in battery and hydrogen technologies have significantly improved the feasibility of ZE propulsion for IWT. Among the emerging solutions, containerized energy carrier swapping offers a particularly promising approach, enabling both battery-electric and hydrogen fuel cell vessels to overcome operational challenges associated with conventional charging and refueling.

In a containerized battery swapping system, a vessel does not charge onboard. Instead, batteries are charged with renewable electricity at swapping stations and are swapped with depleted ones when the vessel calls at the terminal. Unlike conventional

* Corresponding author.
E-mail address: v.e.ramos@tudelft.nl (V.M.J. Ramos).

https://doi.org/10.1016/j.tre.2026.105258
Received 24 March 2026; Received in revised form 11 September 2026; Accepted 12 September 2026
Available online 23 September 2026
1366-5545/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

onboard charging, this approach reduces vessel downtime to minutes rather than hours, and enables a service-based or pay-per-use business model that lowers upfront capital requirements for fleet operators (Zero Emission Services, 2021a). Early deployments in containerized battery-electric propulsion such as the inland vessels *Alphenaar* and *MS Den Bosch Max Groen* in the Netherlands (Port of Rotterdam Authority, 2021; Zero Emission Services, 2025) illustrate the technical feasibility of electric and battery-swapping concepts for IWT.

Despite these developments, a battery swapping system for IWT cannot scale on its own. Unlike road-based electric vehicles, which can rely on dense and flexible charging networks, electric vessels are constrained to navigable waterways and limited terminals. Swapping stations must be strategically positioned along key corridors and backed by sufficient battery inventory. Success relies on large, upfront, and coordinated investments. Without a functioning network, fleet operators have little incentive to adopt ZE vessels. This makes it financially unattractive without sufficient vessel adoption. This classic "chicken-and-egg" coordination failure means that market forces alone are unlikely to drive the transition. The mutual dependence between adoption and infrastructure is an indirect network effect well documented in two-sided markets for infrastructure-dependent transport (Li et al., 2017; Springel, 2021), and because such network externalities lead private actors to under-provide infrastructure, they form the established rationale for public intervention (Leibowicz, 2018; Rapson and Muehlegger, 2023).

Well-designed public subsidies can play a pivotal role in overcoming these coordination failures. Evidence from the electrification of road transport indicates that such market frictions rarely resolve without policy intervention. Governments have deployed purchase subsidies, tax incentives, and infrastructure support to offset high upfront costs and limited early network coverage. Empirical studies also show that these instruments can significantly accelerate adoption (Pi et al., 2024; Li et al., 2023; Srivastava et al., 2022). Comparable challenges are now emerging in IWT. Key policy questions concern the optimal allocation of support between infrastructure developer and fleet operators, the appropriate targeting of instruments, and the subsidy levels required to trigger self-sustaining adoption dynamics. However, for IWT and battery swapping systems, systematic analysis of subsidy design and its interaction with infrastructure investment decisions remains limited.

We address this gap by developing a multi-level optimization framework for the coordinated deployment of containerized battery swapping infrastructure in IWT. This paper makes the following contributions:

- We develop an integrated optimization framework that jointly captures subsidy design, infrastructure deployment, and fleet adoption for zero-emission IWT. In contrast to existing inland-waterway studies, our framework models the coordination between two distinct public instruments (a per-swap infrastructure subsidy to the developer and a per-swap fleet subsidy to operators) and the strategic responses they induce across all three actors.
- We formulate the integrated optimization problem as a multi-level program reflecting the sequential leader–follower structure among policymaker, developer, and operators, and reduce it to a tractable bilevel program through an exact KKT reformulation of the operator level. To solve this at realistic network sizes, we use an iterative bilevel decomposition with warm-start column generation and a tailored large neighborhood search. These methods achieve near-exhaustive solution quality at much lower computational cost.
- We incorporate operational features specific to inland waterway transport that road-based electrification models do not capture. Because vessels can swap only at a sparse set of fixed network ports, a route becomes feasible only when its entire required station set is in place. This coverage couples infrastructure deployment to route-level adoption.

Applying the framework to the Dutch–Belgian Rhine–Scheldt corridor, we further derive policy insights on the composition of public support across budget levels and on the conditions under which subsidies are necessary to enable deployment.

The remainder of this paper is structured as follows. Section 2 reviews relevant literature on multi-level optimization, subsidy design for clean technologies, and prior work on ZE IWT and battery swapping systems. Section 3 presents the mathematical formulation of our multi-level framework, including the decision variable, objectives, and constraints at each level, along with the solution approaches. Section 4 reports the results of numerical experiments and sensitivity analyses. Section 5 concludes by summarizing the main findings and Section 6 outlines the limitations and directions for future research.

2. Review of related literature

The deployment of containerized battery swapping system in IWT is at the intersection of several research areas: facility location and infrastructure design, subsidy and incentive policy for clean technology adoption, and strategic and operational planning for electrified IWT. While each of these areas is well developed in its own right, they are rarely studied together. This section reviews the relevant work across three areas and identifies the gaps that motivate our contribution.

Bilevel optimization has become a widely used framework for modeling strategic interactions, where planners, infrastructure providers, and operators make interdependent decisions regarding infrastructure deployment, pricing, and technology adoption. In road transportation, applications include battery-swapping station deployment (Zhang et al., 2023), infrastructure planning with endogenous user adoption (Lamontagne et al., 2023; Piedra de la Cuadra and Ortega, 2024), and freight decarbonization policy design (Villamar and Aussel, 2024). Similar approaches have also been adopted in electric aviation for infrastructure sizing and charging logistics planning (Mitici et al., 2022; van Oosterom and Mitici, 2023). Methodologically, these bilevel models are commonly reformulated through KKT conditions, duality theory, or MPEC approaches to capture the hierarchical relationship between infrastructure decisions and user responses (Rahiminia et al., 2023; Li et al., 2023; Ghanbari Motlagh et al., 2025). Because these reformulations couple strategic planning decisions with large-scale operational optimization, they often become computationally challenging. Consequently, decomposition-based methods such as column generation have been widely adopted to exploit problem

structure in transportation planning and routing problems (Capelle et al., 2019; Lienkamp and Schiffer, 2024; Huang et al., 2023). In parallel, large neighborhood search (LNS) and its adaptive variants have demonstrated strong performance for complex electrified transport and vehicle-routing applications (Limmer, 2023; Li et al., 2024; Amiri et al., 2023). Despite these methodological advances, existing studies primarily address infrastructure deployment or policy design in isolation and rarely consider the integrated design of subsidies, infrastructure investment, and technology adoption in inland waterway transport.

Beyond infrastructure planning, the successful deployment of low-emission transport technologies depends critically on effective incentives. A growing body of research has examined how governments can design subsidies and other policy instruments to accelerate technology diffusion. Existing literature employ optimization, game-theoretic, and simulation-based approaches to analyze the interactions among governments, infrastructure providers, and users (Brozynski and Leibowicz, 2022; Chen et al., 2023; Shi et al., 2022; Shao et al., 2017; Luo et al., 2023; Wang et al., 2025). Across a variety of transportation settings, these studies consistently show that the effectiveness of subsidy schemes depends on market structure, stakeholder behavior, and the balance between infrastructure-side and user-side incentives.

More specific to our setting, Cui et al. (2025) design an optimal subsidy scheme for electric ship adoption in inland shipping. They analyzed three different subsidy schemes related to charging station investment, electric ship purchases, and charging fees. Wang et al. (2022) model and optimize a governmental subsidy plan for LNG-fueled maritime transportation, considering the decisions at the government, port, and ship levels. Han et al. (2026) analyze the coordination deadlock among four stakeholders in the adoption of hydrogen-powered inland vessels through an evolutionary game framework. Together, these studies span the leading alternative-fuel pathways for waterborne transport and consistently show that subsidy effectiveness depends on alignment of multiple actors. However, none of these jointly optimizes infrastructure siting alongside subsidy design, which is one of the central features of our framework.

Studies on electrified IWT have examined a range of infrastructure planning, voyage planning, and energy management problems. Early strategic work focused on charging infrastructure deployment. For example, Villa et al. (2020) formulate the electric riverboat charging station location problem to determine the location and capacity of shore-side charging facilities that support vessel operations on river networks. Zhang et al. (2024) integrate battery-swapping station deployment with voyage planning through a bilevel optimization framework, demonstrating how infrastructure availability and waterway conditions jointly influence routing and station investment decisions.

Alongside infrastructure planning, a number of work has focused on operational planning and energy management for electric vessels. The studies by Gu et al. (2021), Guo and Dai (2023), Zhang et al. (2023) develop voyage planning models for vessels equipped with swappable battery systems, optimizing routing and swapping decisions under predefined infrastructure configurations. Sun et al. (2023) investigate speed and energy management under dynamic electricity prices and navigation conditions, while Yang et al. (2024) extend these ideas to fleet-level voyage scheduling. Similarly, Zhang et al. (2022) jointly optimize battery capacity and sailing speed to balance operational cost, travel time, and vessel range.

More recently, Guo et al. (2025) develop a mixed-integer optimization framework that jointly determines replenishment locations, technology choices, and replenishment quantities under partial replenishment strategies. For day-ahead decisions, Meherishi et al. (2026) propose a two-phase optimization framework for battery replenishment and repositioning in electric barge networks, minimizing battery investment and logistics costs. Focusing on fleet operations, Liu et al. (2026) develop a mixed-integer programming model for coordinating charging and battery-swapping activities across multiple vessels and ports under charging capacity and battery inventory constraints.

Taken together, the existing literature reveals a clear separation between two research streams. Studies on subsidy design primarily examine policy effectiveness and stakeholder behavior but generally abstract from network-level infrastructure planning and operational characteristics of inland shipping. In contrast, IWT electrification studies develop increasingly sophisticated infrastructure and operational optimization models but typically treat policy incentives and technology adoption as exogenous. As a result, there remains no integrated framework that simultaneously determines subsidy policies, battery-swapping infrastructure deployment, and fleet adoption decisions while accounting for the operational characteristics of inland waterway transport.

To address this gap, we develop a multi-level optimization model for containerized battery swapping in IWT that captures the strategic interaction among a policymaker, an infrastructure developer, and a fleet of heterogeneous operators. The framework jointly determines infrastructure-side and operator-side subsidies, battery-swapping station deployment decisions, and fleet adoption patterns. Through this integrated perspective, we analyze how subsidy policies, infrastructure investment, and operator participation interact to shape the emergence of zero-emission inland shipping networks under different budget and deployment conditions.

3. Multi-level optimization model

Decarbonizing inland waterway transport with containerized battery swapping requires coordinated decisions by three distinct but interdependent actors: the policymaker, the infrastructure developer, and the fleet operator. Each actor pursues its own objective while responding to the decisions made at the other levels. We formalize these interactions as a multi-level optimization problem.

3.1. Fleet operator (L3)

We start at the operator level, where we model technology adoption as a route-level trip allocation decision. We consider multiple fleet operators with different cost structures. For each operator $k \in \mathcal{K}$, a given annual demand τ_{rk} on route r can be served either by conventional diesel vessels or by zero-emission vessels using containerized battery swapping. The operator decides what portion of

trips on each route will be served by each technology. We assume that these allocation decisions are primarily driven by generalized operating cost, while acknowledging that behavioral and organizational factors influence adoption in practice.

L3 sets and parameters:

$\mathcal{K}$	set of fleet operators $k \in \mathcal{K}$
$\mathcal{I}$	set of candidate station nodes $i \in \mathcal{I}$
$\mathcal{R}$	set of routes $r \in \mathcal{R}$, with endpoints $o_r, d_r \in \mathcal{I}$
$\mathcal{S}_r \subseteq \mathcal{I}$	required station set for route r
ℓ_r	longest sub-leg between consecutive stations in $\mathcal{S}_r$ [km]
n_r	number of battery packs carried on board on route r
s_r	number of swap events per one-way trip on route r
τ_{rk}	annual trips on route r done by operator k
t_r	sailing-hours per trip on r [hr]
R_b	usable battery range per pack [km]
C_k^{DS}	generalized operating cost on diesel for operator k [€/hour]
p	swap price set by the infrastructure developer [€/swap]
x_i	station-build decisions (from L2)
π_F	per-swap fleet subsidy to the operators [€/swap] (from L1)

A ZE vessel operating on route r carries containerized battery packs of limited range and must exchange them at swap stations along its path. We capture this through a required station set $\mathcal{S}_r \subseteq \mathcal{I}$: the set of ports that must be equipped with a swap station for a zero-emission voyage on route r to be feasible. This set always contains the route endpoints and, on longer routes, one or more intermediate ports where pack range would otherwise be exceeded. Its cardinality $|\mathcal{S}_r|$ gives the number of stations at which those exchanges occur, and it also determines how far the vessel sails between them. The on-board pack count n_r is dimensioned by the longest such sub-leg ℓ_r, and since the full inventory is exchanged at each visited station, the number of swap events per one-way trip is $s_r = |\mathcal{S}_r|\, n_r$ (Appendix C). We treat $\mathcal{S}_r$ and s_r as exogenous inputs, computed from the network topology and battery range prior to optimization. Endogenizing the choice of intermediate swap ports, by jointly optimizing station location with routing and replenishment decisions, couples strategic infrastructure planning with operational decisions, and is left as an extension. Operator heterogeneity is concentrated on the diesel cost side, where differences across firms are most pronounced in practice. On the zero-emission side, a common swap price p and subsidy π_F apply uniformly across all operators.

L3 decision variables:

$q_{rk}^{G} \geq 0$	number of zero-emission trips on route r by operator k
$q_{rk}^{DS} \geq 0$	number of diesel trips on route r by operator k

Objective function

$$\min_{q^G, q^{DS}} \sum_{k \in \mathcal{K}} \sum_{r \in \mathcal{R}} \left(C_k^{\mathrm{DS}}\, t_r\, q_{rk}^{DS} + (p - \pi_F)\, s_r\, q_{rk}^{G} \right). \tag{1}$$

Subject to

$$q_{rk}^{DS} + q_{rk}^{G} = \tau_{rk}, \qquad \forall r \in \mathcal{R},\ k \in \mathcal{K} \tag{2}$$

$$q_{rk}^{G} \leq \tau_{rk}\, x_i, \qquad \forall r \in \mathcal{R},\ k \in \mathcal{K},\ \forall i \in \mathcal{S}_r \tag{3}$$

$$q_{rk}^{G} \geq 0, \quad q_{rk}^{DS} \geq 0, \qquad \forall r \in \mathcal{R},\ k \in \mathcal{K}. \tag{4}$$

The objective (1) minimizes total generalized operating cost across all routes and operators, combining the diesel cost $C_k^{\mathrm{DS}} t_r$ per trip with the subsidized cost of zero-emission trips $(p - \pi_F) s_r$. Constraint (2) requires that, for each operator k, the trips allocated on route r equal that operator's annual demand τ_{rk}. Constraint (3) ensures that zero-emission trips on route r are only feasible when every port in $\mathcal{S}_r$ is equipped with a swap station, directly embedding the range-feasibility logic into the adoption decision. Non-negativity of all trip allocations is enforced by constraint (4).

The outcome of this level is the route-level adoption q_{rk}^{G} and swap demand volume $s_r q_{rk}^{G}$ as functions of the swap price p, the per-swap subsidy π_F, and the station location vector x. These outcomes feed into the infrastructure developer model L2 (Section 3.2), where swap demand determines station throughput and revenue. At L1 (Section 3.3), total adoption $\sum_k \sum_r q_{rk}^{G}$ enters the policymaker's objective, while the fleet subsidy π_F contributes to the budget outlay.

Operators differ in their cost of running diesel vessels, reflecting differences in fuel consumption, fleet age, and operating conditions. We represent these differences through C_k^{DS}, the generalized per-hour diesel operating cost. On the zero-emission side, the

common swap price p and subsidy π_F can likewise be interpreted as covering both the direct battery-swapping charge and any amortized costs passed through into the per-swap fee. For a given fleet subsidy π_F, each operator–route pair (r, k) has a distinct adoption threshold

$$\pi_F^{\min}(r,k) = p - \frac{C_k^{\mathrm{DS}}\, t_r}{s_r}, \tag{5}$$

Because the L3 problem is linear, each operator's optimal choice on a route is all-or-nothing. It serves the entire demand with ZE vessels once $\pi_F \geq \pi_F^{\min}(r,k)$, and otherwise remains fully on diesel. Route-level adoption therefore follows a staircase pattern, as π_F rises, operators with higher diesel costs cross their thresholds first and ZE share of demand increases. Whether this adoption materializes at all still depends on the station vector x, which the developer decides in the higher level. We represent adoption purely through cost minimization under infrastructure availability, abstracting from behavioral factors such as risk perception, technology familiarity, contractual commitments, and environmental preferences. This keeps the lower level tractable and focuses the analysis on how subsidy design and infrastructure deployment interact to shape zero-emission uptake. Richer behavioral models are left for future work.

3.2. Infrastructure developer (L2)

At the second level, the infrastructure developer chooses the physical network that determines whether the operator adoption modeled in L3 can materialize. We assume a single representative infrastructure developer, or a coalition acting as a coordinated platform, that operates all stations and sets a common swap price p to all operators and routes. The developer decides where to build battery-swapping stations, x_i, among the candidate locations $i \in \mathcal{I}$, and how many battery packs to allocate at each location, b_i, in response to the subsidy instruments set by the policymaker at L1 and in anticipation of fleet operator behavior at L3.

The developer thus occupies a central position in the framework, translating public support into physical network infrastructure and shaping adoption opportunities through the coverage and capacity of the swapping service. This central position also gives the developer a leader role in the developer-operator relationship, as infrastructure investment decisions must precede operators' adoption decisions. Battery swap stations are capital-intensive and long-lived assets that make zero-emission operation physically infeasible without sufficient station coverage. Infrastructure availability is therefore a precondition for fleet-side adoption rather than a simultaneous co-decision, and the developer must commit to its network before observing the operators' response. The same leader-follower structure is established in the electric-vehicle charging deployment literature, where infrastructure provision likewise precedes and enables adoption (Brozynski and Leibowicz, 2022; Lamontagne et al., 2023).

L2 parameters

$SW_i^{\max}$	maximum number of swaps per year that station i can serve
$BSS^{\min}$, $BSS^{\max}$	minimum and maximum number of stations in the network
C_i^{BSS}	annualized cost (capex + opex) to build and operate a swap station at i [€/year]
C^{bat}	annualized cost (capex + opex) of one battery pack [€/year]
δ	maximum number of full charge–discharge cycles one battery pack can deliver per year
s_r	swap events generated per zero-emission trip on route r
$\mathcal{S}_r$	required station set for route r
p	swap price charged to fleet operators [€/swap]
π_I	per-swap infrastructure subsidy to the developer (from L1)
q_{rk}^G	zero-emission trips on route r by operator k (outcome of L3)

L2 decision variables

$x_i \in \{0,1\}$	1 if a station is built at node i, 0 otherwise
$b_i \in \mathbb{Z}$	effective number of battery packs installed at station i

Objective function

$$\max_{x,b}\ (p+\pi_I)\sum_{r\in\mathcal{R}}\sum_{k\in\mathcal{K}} s_r\, q_{rk}^G \;-\; \sum_{i\in\mathcal{I}}\left(C_i^{\mathrm{BSS}}\, x_i + C^{\mathrm{bat}}\, b_i\right). \tag{6}$$

Subject to

$$\sum_{k\in\mathcal{K}}\sum_{r:\, i\in\mathcal{S}_r} \frac{s_r}{|\mathcal{S}_r|}\, q_{rk}^G \leq \delta\, b_i, \qquad \forall i \in \mathcal{I} \tag{7}$$

$$\sum_{k\in\mathcal{K}}\sum_{r:\, i\in S_r}\frac{s_r}{|S_r|}\,q^G_{rk}\le SW_i^{\max}\,x_i, \qquad \forall i\in\mathcal{I} \tag{8}$$

$$BSS^{\min}\le\sum_{i\in\mathcal{I}}x_i\le BSS^{\max} \tag{9}$$

$$x_i\in\{0,1\},\quad b_i\ge 0, \qquad \forall i\in\mathcal{I}. \tag{10}$$

The objective (6) maximizes the developer's annual profit. The first term represents revenue from the battery-swapping service: the developer earns the swap price p per swap from fleet operators, plus the per-swap infrastructure subsidy π_I from the policymaker, with s_r swaps generated per zero-emission trip on route r. The second term captures the annualized cost of building and operating stations and holding battery inventory.

Constraints (7) and (8) govern the capacity of each station. Each zero-emission trip on route r generates s_r swap events, distributed across the $|S_r|$ stations of its required set, so a station $i\in S_r$ handles $s_r/|S_r|$ swaps per trip; summed over the stations of the route this recovers the total s_r swaps that enter the revenue term. Constraint (7) ensures the annual swap load at station i does not exceed the battery-cycle capacity δb_i, while constraint (8) imposes an operational throughput ceiling $SW_i^{\max}$, capturing physical limits such as the number of swap bays, chargers, grid connection, and staffing capacity. Constraint (9) restricts the total number of stations to lie between $BSS^{\min}$ and $BSS^{\max}$, reflecting corridor-level planning requirements or practical limits on deployable sites. Finally, (10) enforces the binary nature of the station-siting decisions; battery inventory b_i is treated as a continuous effective annual quantity.

In this formulation, we do not track individual batteries as they move through the network. Instead, b_i represents the effective number of battery packs assigned to station i on an annual basis. The cycle parameter δ denotes the maximum number of full charge–discharge cycles that each pack can deliver per year, and already accounts for capacity losses due to battery degradation and charging inefficiency. It represents a net usable throughput figure. In practice, when a vessel swaps at a station, it arrives with a depleted pack and departs with a charged one, keeping the number of packs physically present approximately constant over time. Short-term imbalances due to asymmetric routing or repositioning are handled operationally by the developer and are not modeled explicitly here. Constraint (7) therefore enforces a long-run capacity condition consistent with the annual time horizon of the model, allowing us to capture the cost and capacity implications of battery inventory without modeling detailed battery logistics. A more granular representation with time-dependent battery flows and repositioning decisions would significantly increase model complexity and is beyond the scope of this strategic analysis.

Taken together, the L2 problem is a capacitated facility location and sizing problem with endogenous demand. The infrastructure developer weighs the additional swap revenue from a denser network and larger battery inventory against the fixed and variable costs of building and equipping stations. Through the dependence of q^G_{rk} on x, the L2 solution embeds expectations about how infrastructure coverage and fleet-side subsidies translate into swap demand and network profitability. Both subsidy levels, however, are themselves chosen above the developer by the policymaker, who sets them to steer the whole system toward its own objective.

3.3. Policymaker (L1)

The policymaker sits above the developer and sets the two subsidies that drive the responses modeled in L2 and L3. This public authority seeks to accelerate the adoption of ZE technology in IWT and has two financial instruments available. A per-swap infrastructure subsidy π_I paid to the developer, which improves station profitability and expands network coverage, and a per-swap fleet subsidy π_F paid directly to operators, which lowers the marginal cost of ZE operation on routes that are already served. Together these instruments allow the policymaker to allocate public support between the supply and demand sides of the swapping ecosystem under a limited budget G.

L1 parameters

G	public budget ceiling [€/year]
p	swap price set by the infrastructure developer [€/swap]
s_r	number of swap events per zero-emission trip on route r (from L3)
$\pi_F^{\max}$	maximum allowable per-swap subsidy to fleet operators
$\pi_I^{\max}$	maximum allowable per-swap subsidy to the infrastructure developer
$x_i,\ q^G_{rk}$	equilibrium L2/L3 outcomes induced by (π_F,π_I) at price p

L1 decision variables

$\pi_F\ge 0$	per-swap subsidy paid to fleet operators
$\pi_I\ge 0$	per-swap subsidy paid to the infrastructure developer

Objective function

$$\max_{\pi_F,\pi_I}\ \sum_{k\in\mathcal{K}}\sum_{r\in\mathcal{R}}q^G_{rk} \tag{11}$$

Subject to

$$(\pi_F + \pi_I) \sum_{k \in \mathcal{K}} \sum_{r \in \mathcal{R}} s_r \, q^G_{rk} \leq G \tag{12}$$

$$0 \leq \pi_F \leq \pi_F^{\max} \tag{13}$$

$$0 \leq \pi_I \leq \pi_I^{\max} \tag{14}$$

The budget constraint (12) limits total annual policy outlay. Since both subsidies are disbursed per swap, the constraint bounds the combined per-swap public expenditure $(\pi_F + \pi_I)$ against the total network-wide swap volume $\sum_{k,r} s_r \, q^G_{rk}$, rather than capping each instrument independently. Bounds (13) and (14) impose regulatory or practical limits on subsidy levels: $\pi_F^{\max}$ would typically not exceed p, so that the subsidized swap cost remains non-negative and operators are not paid to swap, while $\pi_I^{\max}$ can be calibrated to reflect the scale of station construction costs.

The subsidy pair (π_F, π_I) affects adoption exclusively through the lower levels: for any given (π_F, π_I), the developer and operators solve their respective problems at L2 and L3, yielding equilibrium outcomes x_i, b_i, q^G_{rk}, and q^{DS}_{rk} that feed back into the objective and budget constraint. The L1 problem therefore optimizes over subsidy design while treating these outcomes as the endogenous response of a profit-seeking developer and cost-minimizing fleet operators.

We use adoption maximization as the primary policy objective, in line with the fleet transition pathway defined by the Central Commission for the Navigation of the Rhine (Central Commission for the Navigation of the Rhine, 2021). Maximizing (11) directly operationalizes this pathway. It incentivizes the policymaker to deploy public funds where they unlock the largest fleet transition, subject to the infrastructure and cost-competitiveness constraints embedded in L2 and L3. We acknowledge that policy objectives are in practice multidimensional. Beyond adoption maximization, a policymaker might instead focus on welfare or in the minimization of emissions for a given adoption target. The current structure accommodates such alternatives through a modification of the L1 objective alone, leaving the L2–L3 equilibrium unchanged.

3.4. Solution methods

The full model forms a multi-level optimization problem in which the three stakeholders interact strategically. Solving this structure as a monolithic mathematical program would be computationally intractable in general. We therefore adopt a two-step strategy. First, we analytically eliminate the fleet operator problem (L3) by replacing it with its optimality conditions, reducing the model to a bilevel problem between policymaker and developer. Second, we solve this bilevel problem with two methods: an iterative bilevel decomposition and a problem-specific Large Neighborhood Search.

We evaluate both methods through a common lexicographic policy objective. Let $Q(\pi_F, \pi_I)$ denote total ZE trips and $\text{Outlay}(\pi_F, \pi_I) = (\pi_F + \pi_I)\, B(\pi_F, \pi_I)$ the associated annual public expenditure, where $B(\pi_F, \pi_I)$ is the total swap volume induced by the equilibrium responses of the developer and operators at the given subsidies. The policymaker maximizes ZE adoption subject to the budget and policy bounds,

$$Q^\star = \max_{(\pi_F, \pi_I)} Q(\pi_F, \pi_I), \tag{15a}$$

$$\text{s.t.} \quad (\pi_F + \pi_I)\, B(\pi_F, \pi_I) \leq G, \tag{15b}$$

$$0 \leq \pi_F \leq \pi_F^{\max}, \quad 0 \leq \pi_I \leq \pi_I^{\max}, \tag{15c}$$

and then, among all subsidy pairs achieving this maximum adoption up to tolerance $\varepsilon > 0$, minimizes public expenditure:

$$\min_{(\pi_F, \pi_I)} \text{Outlay}(\pi_F, \pi_I) \quad \text{s.t.} \quad Q(\pi_F, \pi_I) \geq Q^\star - \varepsilon, \ \text{(15b)-(15c)}.$$

This criterion is enforced directly through the master problem design in the iterative bilevel decomposition, and through the acceptance rule in the problem-specific LNS. We now develop each step of the strategy in turn, beginning with the elimination of L3.

3.4.1. Reformulating the operator level (L3)

The L3 problem is a linear program, so its Karush-Kuhn-Tucker (KKT) conditions are necessary and sufficient for global optimality, a property we exploit to replace the explicit operator problem with its KKT system. For each route-operator pair (r, k), we introduce dual variables associated with the demand-balance and station feasibility constraints. Yielding a system of (i) primal feasibility conditions, (ii) dual feasibility conditions, and (iii) complementary slackness conditions that link binding constraints to positive duals. The full KKT system and resulting reformulation are given in Appendix A.

Embedding this KKT system within the infrastructure developer's problem (L2) collapses the two lower levels into a single-level mathematical program with equilibrium constraints (MPEC). For any given policy (π_F, π_I) and swap price p, the developer's problem internalizes the optimal response directly. We denote this combined problem as [L2 + KKT(L3)]. The original tri-level problem thereby reduces to a bilevel problem in which the policymaker (L1) as the upper level and the infrastructure–fleet equilibrium [L2 + KKT(L3)] is the lower level. The two succeeding methods differ only in how they search the policymaker's decision space over this equilibrium.

3.4.2. Iterative bilevel decomposition

The first method solves the bilevel problem using an iterative decomposition with two phases. An initialization phase constructs a pool of candidate infrastructure-fleet equilibria directly from the network structure, providing a warm start; a master-oracle loop then refines this pool until bilevel consistency is achieved. The full procedure is given in Algorithm 1.

Algorithm 1: Warm-started iterative bilevel decomposition.

Input: Network, demand, and cost data; swap price p; budget G; station bounds $n_{\min}, n_{\max}$; tolerances $\varepsilon_Q, \varepsilon_B$; max iterations $N_{\max}$

Output: Policy $(\pi_F^\star, \pi_I^\star)$ and equilibrium $(x^\star, b^\star, q^{G,\star}, q^{DS,\star})$

Phase 1: Initialization;
$\mathcal{L} \leftarrow \{\ell_\emptyset\}$; // initialise with empty column
for $n = n_{\min}$ **to** $n_{\max}$ **do**
  for *each* $S \subseteq \mathcal{I}$ *with* $|S| = n$ **do**
    Compute $\mathcal{R}(S)$ and thresholds $\{\pi_F^{\min,(r,k)}\}$ via (17);
    for *each unique breakpoint tier* π_F^t **do**
      $\ell \leftarrow$ `BuildTierColumn`(S, π_F^t) via (19);
      if $\ell \notin \mathcal{L}$ ***and*** $B_\ell > 0$ **then** $\mathcal{L} \leftarrow \mathcal{L} \cup \{\ell\}$;
Remove columns violating (20) from $\mathcal{L}$;
Phase 2: Iterative decomposition;
for $t = 1$ **to** $N_{\max}$ **do**
  `// Master step`
  $(\pi_F^{(t)}, \pi_I^{(t)}, \ell_{\text{sel}}) \leftarrow$ `SolveMaster`$(\mathcal{L}, G)$;
  `// Oracle step`
  $(\hat{x}, \hat{b}, \hat{q}^G, \hat{q}^{DS}, \hat{Q}, \hat{B}, \hat{S}, \hat{C}^{\text{infra}}, \hat{\pi}_F^{\min}) \leftarrow$ `SolveL2+KKT(L3)`$(\pi_F^{(t)}, \pi_I^{(t)})$;
  `// Consistency check via (22)`
  if $|\hat{Q} - Q_{\ell_{\text{sel}}}| \le \varepsilon_Q$ ***and*** $|\hat{B} - B_{\ell_{\text{sel}}}| \le \varepsilon_B$ ***and*** $\hat{S} = S_{\ell_{\text{sel}}}$ **then**
    return $(\pi_F^{(t)}, \pi_I^{(t)}, \hat{x}, \hat{b}, \hat{q}^G, \hat{q}^{DS})$;
  `// Enrich pool with oracle column`
  if *oracle equilibrium* $\notin \mathcal{L}$ **then**
    Add ℓ_{new} to $\mathcal{L}$ storing $(\hat{x}, \hat{b}, \hat{q}^G, \hat{q}^{DS}, \hat{Q}, \hat{B}, \hat{S}, \hat{C}^{\text{infra}}, \hat{\pi}_F^{\min})$;
    ; // budget-infeasible columns added as ghost columns: blocked from selection via (21f) but calibrate v
return $(\pi_F^{(t)}, \pi_I^{(t)}, \hat{x}, \hat{b}, \hat{q}^G, \hat{q}^{DS})$;

We index candidate equilibria by $\ell \in \mathcal{L}$. Each column ℓ stores a feasible configuration $(x_i^\ell, b_i^\ell, q_{rk}^{G,\ell}, q_{rk}^{DS,\ell})$ together with the annualized infrastructure cost $C_\ell^{\text{infra}} = \sum_{i\in\mathcal{I}} C_i^{\text{BSS}} x_i^\ell + C^{\text{bat}} b_i^\ell$ and the scalar summaries

$$Q_\ell = \sum_{k\in\mathcal{K}} \sum_{r\in\mathcal{R}} q_{rk}^{G,\ell}, \qquad B_\ell = \sum_{k\in\mathcal{K}} \sum_{r\in\mathcal{R}} s_r\, q_{rk}^{G,\ell}, \qquad S_\ell = \sum_{i\in\mathcal{I}} x_i^\ell,$$

The pool contains two types of columns. *Analytical columns* are constructed directly from the network structure without oracle calls, providing a warm start that covers the policy space. *Oracle columns* are computed by solving [L2 + KKT(L3)] at a specific policy and represent the developer's true equilibrium response there. They are added during the oracle phase to confirm and, where necessary, complete the pool.

For each station subset $S \subseteq \mathcal{I}$ with $|S| \in [n_{\min}, n_{\max}]$, the covered routes are those for which all required swap stations are present:

$$\mathcal{R}(S) = \{r \in \mathcal{R} : S_r \subseteq S\}. \tag{16}$$

For each (r,k) with $r \in \mathcal{R}(S)$, the per-route adoption threshold and its column-level aggregation are:

$$\pi_F^{\min,(r,k)} = \max\left(0,\ p - \frac{C_k^{\text{DS}}\, t_r}{s_r}\right), \tag{17}$$

$$\pi_F^{\min,\ell} = \max_{(r,k):\, q_{rk}^{G,\ell} > 0} \pi_F^{\min,(r,k)}. \tag{18}$$

Let $\{\pi_F^1 < \pi_F^2 < \cdots < \pi_F^T\}$ denote the unique values taken by (17) across all covered route–operator pairs. For each breakpoint tier $t = 1, \ldots, T$ we construct one analytical column with trip allocation

$$q_{rk}^{G,\ell} = \begin{cases} \tau_{rk} & \text{if } r \in \mathcal{R}(S) \text{ and } \pi_F^{\min,(r,k)} \le \pi_F^t, \\ 0 & \text{otherwise,} \end{cases} \qquad q_{rk}^{DS,\ell} = \tau_{rk} - q_{rk}^{G,\ell}. \tag{19}$$

This construction ensures the pool contains one column for every adoption level that can arise as π_F crosses a breakpoint, covering the policy space analytically.

Before the iterative loop begins, analytical columns for which no valid policy can satisfy the budget constraint are removed from $\mathcal{L}$. Specifically, a column ℓ is retained only if

$$\left(\pi_F^{\min,\ell} + \max\left(0,\ \frac{C_\ell^{\text{infra}}}{B_\ell} - p\right)\right) B_\ell \ \le\ G, \tag{20}$$

where the second term in (20) is the minimum per-swap infrastructure subsidy the developer requires to break even, given the swap revenue $p\,B_\ell$ already earned.

Given a pool $\mathcal{L}$ of columns, the master problem approximates the policymaker's problem (L1) by jointly optimizing the subsidy levels (π_F, π_I) and selecting one column from the pool. Let $y_\ell \in \{0,1\}$ equal 1 if column ℓ is selected and 0 otherwise, let $z_\ell \in \{0,1\}$ equal 1 if column ℓ is budget-feasible at the chosen policy and 0 otherwise, and let $\upsilon \in \mathbb{R}$ be an auxiliary variable representing the infrastructure developer's profit under the chosen column and subsidy π_I. The master problem is:

$$\max_{\pi_F,\pi_I,y,z,\upsilon} \quad \sum_{\ell\in\mathcal{L}} Q_\ell\, y_\ell \tag{21a}$$

$$\text{s.t.} \quad \sum_{\ell\in\mathcal{L}} y_\ell = 1, \tag{21b}$$

$$\pi_F \ge \sum_{\ell\in\mathcal{L}} \pi_F^{\min,\ell}\, y_\ell, \tag{21c}$$

$$(\pi_F+\pi_I)\,B_\ell \ \le\ G + M_G\,(1-z_\ell), \quad \forall\,\ell\in\mathcal{L}, \tag{21d}$$

$$(\pi_F+\pi_I)\,B_\ell \ \ge\ (G+\varepsilon) - M_G\,z_\ell, \quad \forall\,\ell\in\mathcal{L}, \tag{21e}$$

$$y_\ell \ \le\ z_\ell, \quad \forall\,\ell\in\mathcal{L}, \tag{21f}$$

$$\upsilon \ \ge\ \Pi_\ell(\pi_I) - M_P\,(1-z_\ell), \quad \forall\,\ell\in\mathcal{L}, \tag{21g}$$

$$\upsilon \ \le\ \Pi_\ell(\pi_I) + M_P\,(1-y_\ell), \quad \forall\,\ell\in\mathcal{L}, \tag{21h}$$

$$0 \le \pi_F \le \pi_F^{\max}, \quad 0 \le \pi_I \le \pi_I^{\max}, \tag{21i}$$

$$y_\ell \in \{0,1\}, \quad z_\ell \in \{0,1\}, \quad \forall\,\ell\in\mathcal{L}, \tag{21j}$$

$$\upsilon \in \mathbb{R}, \tag{21k}$$

where $\Pi_\ell(\pi_I) = (p+\pi_I)\,B_\ell - C_\ell^{\text{infra}}$ is the developer's annual profit under column ℓ and subsidy π_I. The constant M_G is calibrated to be strictly larger than both G and the maximum possible budget expression, ensuring that (21d) and (21e) correctly deactivate when z_ℓ is fixed. The constant M_P bounds the maximum profit any column can earn, deactivating the lower bound on υ in (21g) for budget-infeasible columns. The small tolerance $\varepsilon > 0$ in (21e) approximates the strict inequality $(\pi_F+\pi_I)B_\ell > G$ required when $z_\ell = 0$. The objective (21a) maximizes total ZE trips over the selected column. Constraint (21b) enforces the selection of exactly one column. Constraint (21c) ensures the chosen fleet subsidy is sufficient to sustain the ZE adoption embedded in the selected column. Constraints (21d) and (21e) jointly define z_ℓ as a budget-feasibility indicator: when $z_\ell = 1$, (21d) enforces $(\pi_F+\pi_I)B_\ell \le G$; when $z_\ell = 0$, (21e) enforces $(\pi_F+\pi_I)B_\ell > G$. Constraint (21f) prevents selection of any budget-infeasible column. Constraints (21g) and (21h) enforce the best-response condition: υ equals the maximum developer profit over all budget-feasible columns, and the selected column must attain this maximum. Excluding budget-infeasible columns from the lower bound via M_P in (21g) is essential, without it, a column with large B_ℓ that is unaffordable at the current policy would inflate υ, making it impossible for any selectable column to satisfy (21h). Finally, (21i)–(21k) declare policy bounds and variable domains.

From the master's proposed policy $(\bar\pi_F, \bar\pi_I)$, we solve the oracle subproblem to compute the developer's true best response, yielding an equilibrium $(\hat x_i,\ \hat b_i,\ \hat q_{rk}^{G},\ \hat q_{rk}^{DS})$ with summaries $\hat Q$, $\hat B$, $\hat S$, $\hat C^{\text{infra}}$, and $\hat\pi_F^{\min}$. Bilevel consistency is checked by comparing these oracle summaries against those of the master's selected column ℓ_{sel}:

$$d_Q \ =\ |\hat Q - Q_{\ell_{\text{sel}}}|, \qquad d_B \ =\ |\hat B - B_{\ell_{\text{sel}}}|, \qquad d_S \ =\ |\hat S - S_{\ell_{\text{sel}}}|. \tag{22}$$

The algorithm terminates if $d_Q \le \varepsilon_Q$, $d_B \le \varepsilon_B$, and $d_S = 0$. The exact equality on d_S is warranted because station-siting variables are binary and admit no approximation error. If consistency is not achieved, the oracle equilibrium is added to $\mathcal{L}$ and the master is re-solved over the augmented pool. If the oracle column is budget-infeasible at the current policy, it is still added to $\mathcal{L}$ but cannot be selected; it participates only in calibrating υ via (21g), providing a corrective signal that discourages subsidy allocations that make high-revenue configurations unaffordable.

The algorithm terminates in a finite number of iterations. The station-siting variables $x_i \in \{0,1\}$ take values in a finite set, so the number of distinct feasible infrastructure–fleet equilibria is finite. Each oracle call either confirms bilevel consistency, triggering termination, or returns an equilibrium not already in the pool, which is then added. Since columns are never removed and the pool cannot exceed the finite set of distinct equilibria, termination occurs after finitely many iterations. At termination, the returned solution satisfies the bilevel consistency condition: the developer's best response under the master's proposed policy coincides with the selected column, confirming that no profitable deviation exists for the developer at the returned policy. This constitutes the algorithm's correctness guarantee.

Beyond consistency, the initialization provides structural coverage of the infrastructure decision space. Phase 1 enumerates every feasible station configuration $S \subseteq \mathcal{I}$ with $|S| \in [n_{\min}, n_{\max}]$, and for each configuration constructs one analytical column at every adoption breakpoint that can arise as the fleet subsidy π_F sweeps its range. The pool therefore contains a column for every infrastructure configuration and every adoption level reachable on it, so no feasible configuration is unrepresented. Consequently, the only

approximation in the method concerns the continuous policy variables (π_F, π_I), which the master problem optimizes directly rather than by enumeration; the discrete infrastructure dimension carries no representational gap. The oracle phase then serves to validate consistency and to calibrate the developer's best-response variable v. This calibration uses budget-infeasible columns, which prevent the master from proposing policies that are superficially attractive but practically unachievable. Its role is therefore to confirm and refine, not to repair missing configurations, since no feasible configuration is absent from the pool.

3.4.3. Problem-specific large neighborhood search

The iterative bilevel decomposition relies on the completeness of the column pool to guarantee solution quality. As the network grows, the number of potential equilibria expands combinatorially, making exhaustive enumeration impractical and increasing both memory requirements and master problem solve times. To address this scalability limitation, we introduce a problem-specific large neighborhood search (LNS) as a complementary solution method.

LNS is a destroy–repair metaheuristic that iteratively removes part of the current solution and re-optimizes over the freed variables. The solution has two components: the continuous subsidy vector (π_F, π_I) and the discrete infrastructure configuration x_i. While standard LNS is typically applied to purely combinatorial problems, the framework extends naturally to this mixed continuous-discrete setting by defining tailored destroy and repair operators for each component. In both cases, the repair step consists of fully re-solving the [L2 + KKT(L3)] oracle, which reconstructs the infrastructure–fleet equilibrium in response to a given candidate policy and partial fixing of locations. The dominant computational cost per iteration is therefore the oracle solve, regardless of which neighborhood type is selected.

At each iteration, the algorithm probabilistically alternates between two neighborhood move types:

1. **Policy moves** (probability ρ_{pol}): the destroy step perturbs the incumbent subsidy vector $(\pi_F^{\text{inc}}, \pi_I^{\text{inc}})$ by sampling a displacement uniformly from $[-\Delta_F, \Delta_F] \times [-\Delta_I, \Delta_I]$ and projecting onto the feasible subsidy space $[0, \pi_F^{\max}] \times [0, \pi_I^{\max}]$, yielding a candidate policy (π'_F, π'_I). The repair step re-solves [L2 + KKT(L3)] under (π'_F, π'_I), reconstructing the infrastructure–fleet equilibrium from scratch. These moves explore the continuous subsidy space and allow the algorithm to escape local optima in the policy dimension.
2. **Infrastructure moves** (probability $1 - \rho_{\text{pol}}$): the destroy step releases a random subset $\mathcal{F}_{\text{free}} \subseteq \mathcal{I}$ of $k \in [k_{\min}, k_{\max}]$ station locations from the incumbent configuration, fixing the remainder at their incumbent values $x_i = x_i^{\text{inc}}$ for all $i \notin \mathcal{F}_{\text{free}}$. The repair step re-solves [L2 + KKT(L3)] under the incumbent policy with the freed locations unconstrained, allowing the developer to reconfigure that subset of the network. These moves identify network topologies that are profit-maximizing under the current subsidy allocation.

The two neighborhood types jointly explore the coupled continuous–discrete solution space. Policy moves identify subsidy levels that unlock new infrastructure configurations; infrastructure moves identify network topologies that are profit-maximizing under the current subsidy allocation. The coupling between both dimensions is preserved by always evaluating candidate solutions through the same [L2 + KKT(L3)] oracle. The full algorithm is given in Algorithm 2.

We adopt a greedy lexicographic acceptance rule consistent with the solution criterion (15). A candidate solution (π'_F, π'_I) with outcomes (Q', Outlay') replaces the incumbent if it yields strictly more ZE trips, or if it achieves the same number of ZE trips at strictly lower public expenditure. A probabilistic acceptance scheme such as simulated annealing could be incorporated as a straightforward extension, but preliminary experiments showed no improvement in solution quality over the greedy variant at the instance sizes considered.

The parameter values are grounded in the structure of the search space. The move probability ρ_{pol} balances search effort between the continuous, two-dimensional subsidy space and the discrete infrastructure space with at most $BSS^{\max}$ configurations, and the perturbation bounds (Δ_F, Δ_I) are set small enough to preserve local structure around the incumbent, while the iteration budget and multi-start scheme provide coverage of the policy box. The number of stations released per infrastructure move is drawn uniformly, ensuring moves are neither too local nor too disruptive.

4. Analysis

4.1. Experimental design

We apply the proposed framework to the Dutch–Belgian inland waterway network, drawing on publicly available data, reports, and insights from emerging zero-emission shipping initiatives in the Rhine–Scheldt corridor, including the deployment vision of ZES. The study network covers 10 strategic inland ports connected by 87 shipping service routes, each route r defined by an origin–destination pair with distance d_r and travel time t_r. Fig. 1 illustrates the geographic distribution of the selected ports and the waterway topology.

Route-level trip demand is based on the 2025 inland vessel statistics maintained by the Dutch Ministry of Infrastructure and Water Management (Rijkswaterstaat, 2025). As vessel-level origin–destination flows and trip frequencies are not publicly available, route-specific annual trip volumes are constructed by combining aggregate port-visit statistics with route distances and an assumed average cruising speed of 15 km/hr. Demand is treated as deterministic and time-invariant over the planning horizon, abstracting from seasonal fluctuation and demand uncertainty.

We assume a standardized ZES-type 3MWh battery container giving an operating range of approximately 90 km per pack. Applying the pack requirement of Section 3.1 to the case-study network, most routes are served at their endpoints only ($|S_r| = 2$), where the longest sub-leg is the full route distance. For the corridors longer than 180 km, spanning the route in a single leg would require an impractical number of on-board packs, so an exogenous mid-route station is introduced, giving $|S_r| = 3$ and a longest sub-leg of at

Algorithm 2: Problem-specific large neighborhood search.

Input: Network and cost data; swap price p; budget G; policy bounds $[\pi_F^{\min}, \pi_F^{\max}]$, $[\pi_I^{\min}, \pi_I^{\max}]$; step sizes (Δ_F, Δ_I); infra neighborhood sizes $k_{\min}, k_{\max}$; policy move probability ρ_{pol}; max iterations $N_{\max}$

Output: Policy $(\pi_F^\star, \pi_I^\star)$ and equilibrium $(x^\star, b^\star, q^{G,\star}, q^{DS,\star})$

Initialization;
Choose initial policy $(\pi_F^{(0)}, \pi_I^{(0)})$ (e.g., random seeds or boundary corners);
$(x^{\text{inc}}, b^{\text{inc}}, q^{G,\text{inc}}, q^{DS,\text{inc}}, Q^{\text{inc}}, \text{outlay}^{\text{inc}}) \leftarrow \texttt{SolveL2+KKT(L3)}(\pi_F^{(0)}, \pi_I^{(0)})$;
$(\pi_F^\star, \pi_I^\star) \leftarrow (\pi_F^{(0)}, \pi_I^{(0)})$, $(Q^\star, \text{outlay}^\star) \leftarrow (Q^{\text{inc}}, \text{outlay}^{\text{inc}})$;
Main LNS loop;
for $i = 1$ **to** $N_{\max}$ **do**
  Draw $u \sim \mathcal{U}[0,1]$;
  if $u < \rho_{pol}$ **then**
    `// Policy move`
    $(\pi'_F, \pi'_I) \leftarrow \texttt{PolicyMove}(\pi_F^{\text{inc}}, \pi_I^{\text{inc}};\ \Delta_F, \Delta_I)$;
    $(x^c, b^c, q^{G,c}, q^{DS,c}, Q^c, B^c, S^c, \text{outlay}^c) \leftarrow \texttt{SolveL2+KKT(L3)}(\pi'_F, \pi'_I)$;
  else
    `// Infrastructure move`
    Draw k uniformly from $\{k_{\min}, \dots, k_{\max}\}$;
    $x^{\text{fix}} \leftarrow \texttt{InfraMove}(x^{\text{inc}}, k)$;
    $(x^c, b^c, q^{G,c}, q^{DS,c}, Q^c, B^c, S^c, \text{outlay}^c) \leftarrow \texttt{SolveL2+KKT(L3)}(\pi_F^{\text{inc}}, \pi_I^{\text{inc}}, x^{\text{fix}})$;
    $(\pi'_F, \pi'_I) \leftarrow (\pi_F^{\text{inc}}, \pi_I^{\text{inc}})$;
  if *solution is feasible* **then**
    `// incumbent update`
    if $Q^c > Q^{\text{inc}} + \varepsilon$ ***or*** $|Q^c - Q^{\text{inc}}| \le \varepsilon$ ***and*** $\text{outlay}^c < \text{outlay}^{\text{inc}}$ **then**
      $(\pi_F^{\text{inc}}, \pi_I^{\text{inc}}) \leftarrow (\pi'_F, \pi'_I)$;
      $(x^{\text{inc}}, b^{\text{inc}}, q^{G,\text{inc}}, q^{DS,\text{inc}}, Q^{\text{inc}}, \text{outlay}^{\text{inc}}) \leftarrow (x^c, b^c, q^{G,c}, q^{DS,c}, Q^c, \text{outlay}^c)$;
    `// global-best update`
    if $Q^c > Q^\star + \varepsilon$ ***or*** $|Q^c - Q^\star| \le \varepsilon$ ***and*** $\text{outlay}^c < \text{outlay}^\star$ **then**
      $(\pi_F^\star, \pi_I^\star) \leftarrow (\pi'_F, \pi'_I)$;
      $(Q^\star, \text{outlay}^\star) \leftarrow (Q^c, \text{outlay}^c)$;
return $(\pi_F^\star, \pi_I^\star, x^\star, b^\star, q^{G,\star}, q^{DS,\star})$;
Function `PolicyMove`$(\pi_F, \pi_I;\ \Delta_F, \Delta_I)$**:**
  Draw $\delta_F \sim \mathcal{U}[-\Delta_F, \Delta_F]$,;
    $\delta_I \sim \mathcal{U}[-\Delta_I, \Delta_I]$;
  $\pi'_F \leftarrow \min\{\pi_F^{\max}, \max\{\pi_F^{\min}, \pi_F + \delta_F\}\}$;
  $\pi'_I \leftarrow \min\{\pi_I^{\max}, \max\{\pi_I^{\min}, \pi_I + \delta_I\}\}$;
  return (π'_F, π'_I);
Function `InfraMove`(x^{inc}, k)**:**
  Sample $F_{\text{free}} \subseteq \mathcal{I}$ with $|F_{\text{free}}| = k$ uniformly at random;
  $x^{\text{fix}} \leftarrow \{x_i^{\text{inc}} : i \in \mathcal{I} \setminus F_{\text{free}}\}$;
  return x^{fix};

most 172 km. No sub-leg in the network exceeds 180 km, so every route requires at most two packs on board. The route characteristics, and the five routes with a required middle station, are given in Appendix B.

The infrastructure developer incurs station capital and operational costs that vary by location, reflecting land availability, permitting, and construction requirements. We distinguish two port tiers. The major ports (Amsterdam, Antwerp, Rotterdam, Utrecht) carry a higher annualized cost of 1.5 million euro, while the remaining ports are assigned 1.0 million euros. Battery procurement is standardized, with an annualized cost of €77,500 per pack. Swapping operations, charging durations, and maintenance downtime are captured implicitly through an annual cycle capacity of 200 swaps per pack (EPEVER, 2025). Station throughput capacity $SW_i^{\max}$ is set high enough that the capacity constraint is effectively non-binding, so deployment scale is limited only by the budget G. Infrastructure related cost parameters are expressed as combined, annualized capital and operating costs and are derived from publicly available industry and regulatory sources and documented in Appendix D.

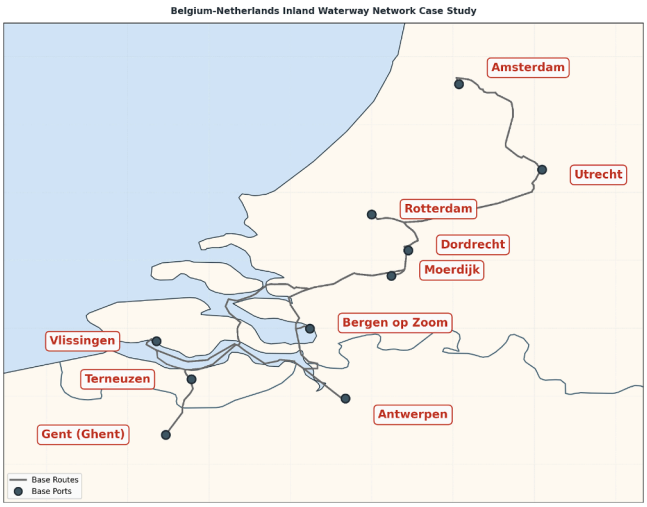


Fig. 1. Dutch–Belgian network case study.

To capture heterogeneous cost structures and adoption tiers, route demand is disaggregated into three operator categories defined by vessel commissioning year, following the age distribution of the European dry-cargo fleet reported in the CCNR annual report (Central Commission for the Navigation of the Rhine , 2022): *standard operators* (commissioned before 2000, 83% of active fleet capacity), *modern operators* (2000–2021, 16%), and *recent entrants* (2022–2025, 1%). The recent entrant class represents the newest vessels, postdating the CCNR reporting window, and is included as a small frontier group to capture the recent builds. We use commissioning year as a proxy for engine age, since fuel efficiency is closely tied to build year. Engine age determines operational efficiency and therefore the diesel sailing cost C_k^{DS}. Older vessels have higher specific fuel consumption and thus higher hourly diesel cost, approximately €76/hour for standard operators, €70/hour for modern operators, and €65/hour for recent entrants, consistent with the estimation basis of TNO (2009). We assume each operator type controls a homogeneous fleet within its category, and that all vessels use standardized, interoperable battery containers.

4.2. Benchmarking

We compare three methods to evaluate the proposed solution approaches: (i) the iterative bilevel decomposition (IBD), (ii) the problem-specific Large Neighborhood Search (LNS), and (iii) a sequential grid search (GS) serving as a reference. All methods are evaluated against the solution criterion (15) and call the same [L2 + KKT(L3)] oracle to evaluate any candidate subsidy pair. This oracle is solved exactly as an indicator-linearized MILP (A.3), so the embedded complementarity is handled. For IBD, a secondary optimization step is applied after convergence to minimize public outlay while maintaining the achieved adoption level. We report for each instance and method, the best subsidy pair found, the resulting public outlay, and average wall-clock solving time. Each method is executed across 15 independent runs per instance, and all reported metrics represent averages over these runs.

Benchmark instances of increasing size are constructed by subsampling the full 10-port Dutch-Belgian network into port subsets of size 5, 8, and 10, retaining only routes whose origin and destination both belong to the subset. For each network size, three operator configurations are defined by progressively including additional operator types: K = 1 includes only standard operators, K = 2 adds modern operators, and K = 3 includes all three types. This yields 12 instances in total, denoted I-R-K in Table 1, where I, R, and K refer to the number of ports, routes, and operator types respectively.

The GS method enumerates a uniform discretization $\Pi_F \times \Pi_I$ over the policy bounds solving [L2 + KKT(L3)] at each of the 2500 grid points. GS evaluates the full extent of the policy box at the chosen resolution, leaving no grid region unsampled, but it is limited to its fixed grid points and does not resolve policies that fall between them. The limitation is the cost, since the number of oracle solves scales as $|\Pi_F| \times |\Pi_I|$. It serves as a valuable reference for assessing whether IBD and LNS systematically miss high-quality policy regions at the considered resolution. Refer to Algorithm 3 for the formulation.

For the LNS, the policy-move probability is set to $\rho_{\mathrm{pol}} = 0.5$, the perturbation step sizes to $\Delta_F = 30$ and $\Delta_I = 50$, and the maximum number of iterations to $N_{\max} = 300$. These values were selected through a grid evaluation over ρ_{pol} and the perturbation step sizes on the smallest and largest benchmark instances, in which full adoption was attained in every configuration and outlay was insensitive to the settings provided sufficient search effort is allocated to the policy dimension; the best-performing setting was then held fixed across all reported experiments. All experiments are conducted on a standard workstation and implemented in Python 3.12 using the Gurobi 12.0, running Windows 11 with an Intel Core 13th generation processor and 16GB of RAM.

Table 1 reveals consistent performance differences across the three methods. All three methods achieve full ZE adoption in every instance, confirming that the policy space contains multiple solutions that attain full adoption at different cost levels. Across all 15 runs per instance, IBD and GS are deterministic and return identical policies and outlays on every run, so no run-to-run variation is

Algorithm 3: Sequential Grid Search over (π_F, π_I).

Input: Instance and cost data; grids Π_F (for π_F) and Π_I (for π_I); unit price p; budget G
Output: Best feasible policy $(\pi_F^\star, \pi_I^\star, Q^\star, B^\star, S^\star, \text{outlay}^\star, \text{profit}^\star)$, or $\varnothing$ if none

best $\leftarrow \varnothing$;
foreach $\pi_F \in \Pi_F$ **do**
  foreach $\pi_I \in \Pi_I$ **do**
    $(Q, B, S, \text{outlay}, \text{profit}) \leftarrow$ `SolveL2+L3`(π_F, π_I);
    if *feasible* **then**
      if $Q > Q^\star$ *OR* $(|Q - Q^\star| \le \varepsilon$ *AND* $\text{outlay} < \text{outlay}^\star)$ **then**
        $(\pi_F^\star, \pi_I^\star) \leftarrow (\pi_F, \pi_I)$;
        $(Q^\star, \text{outlay}^\star, \ldots) \leftarrow (Q, \text{outlay}, \ldots)$;
        best $\leftarrow (\pi_F^\star, \pi_I^\star, Q^\star, \text{outlay}^\star, \ldots)$;
  return $(\pi_F^\star, \pi_I^\star, Q^\star, B^\star, S^\star, \textit{outlay}^\star, \textit{profit}^\star)$;

Table 1
Benchmarking IBD, LNS, and GS across instances (averages over 15 independent runs). All instances achieve 100% ZE adoption.

		IBD				LNS				GS (2,500 policies)			
Instance I-R-K	Q	π_F [€/swap]	π_I [€/swap]	Outlay [€M]	Time [s]	π_F [€/swap]	π_I [€/swap]	Outlay [€M]	Time [s]	π_F [€/swap]	π_I [€/swap]	Outlay [€M]	Time [s]
5-10-1	79,851	137.5	41.1	56.8	0.03	140.8	44.4	59.0	3.32	137.8	61.2	63.4	29.77
5-10-2	490,584	137.5	15.1	298.4	0.06	141.4	18.7	313.2	5.29	137.8	20.4	309.4	36.58
5-10-3	494,576	175.0	15.0	374.7	0.09	181.3	20.7	398.3	8.70	178.6	20.4	392.4	36.80
8-30-1	94,820	196.4	42.9	78.2	0.88	202.8	53.6	83.8	15.60	199.0	61.2	85.0	75.61
8-30-2	582,555	196.4	15.3	425.1	1.63	206.2	25.7	465.6	29.23	199.0	20.4	440.5	92.30
8-30-3	587,296	214.3	15.3	464.7	4.16	222.7	26.3	504.1	14.56	214.3	20.4	475.1	237.37
10-60-1	125,562	196.4	48.3	102.9	8.01	206.0	56.7	110.5	7.78	199.0	61.2	109.5	284.14
10-60-2	771,430	196.4	16.2	549.5	25.27	202.8	22.8	583.1	14.54	199.0	20.4	567.0	333.47
10-60-3	777,710	214.2	16.2	600.4	98.38	219.0	19.7	621.9	23.32	214.3	20.4	611.5	199.83
10-80-1	176,235	196.4	47.7	150.2	13.8	205.3	55.6	160.5	18.23	199.0	61.2	160.1	105.14
10-80-2	1,082,752	196.4	16.1	803.4	40.33	202.8	22.8	853.0	34.17	199.0	20.4	829.4	165.40
10-80-3	1,091,565	214.2	16.1	877.9	263.2	219.2	25.2	931.4	49.34	214.3	20.4	894.5	348.28

reported for them. The stochastic LNS attains full adoption on every run of every instance, with only modest run-to-run variation in outlay (standard deviations of under 10% of the mean). The methods therefore compete purely on expenditure and solving time.

IBD attains the lowest public outlay in every instance, and its solutions align closely with those of GS. This indicates that the decomposition does not miss high-quality regions of the policy space at the resolution tested. IBD reaches these solutions without evaluating thousands of policy points. GS solving time grows steeply with instance size and exceeds 240 s on the largest case, while IBD remains an order of magnitude faster. GS is thus a useful verification tool on small instances but becomes impractical as the network grows.

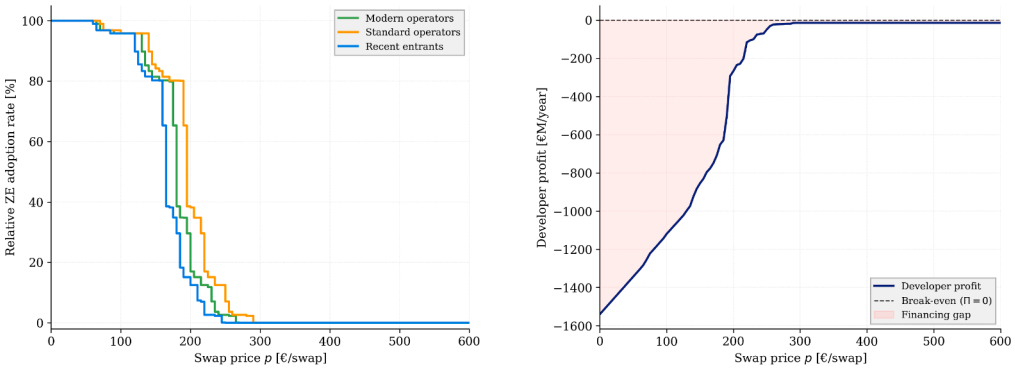
LNS is the fastest method on the large instances and completes every run, but it returns higher outlays than both IBD and GS, and the gap widens with instance size. Two features of its design explain this behavior. First, the greedy lexicographic acceptance rule stops improving outlay once full adoption is reached, leaving the search on whichever policy first attained full adoption rather than the cheapest one. Second, the infrastructure moves reconfigure only a small subset of stations per iteration, which makes it increasingly difficult to escape station configurations that are locally but not globally optimal as the discrete decision space grows. IBD avoids both limitations, the warm-start enumeration provides broad infrastructure coverage independent of the search trajectory, and the post-convergence minimization step explicitly identifies the cheapest policy sustaining full adoption. Because IBD delivers the most budget-efficient solutions at competitive runtimes, we adopt it as the solution method for the remainder of the analysis.

4.3. Subsidy analysis

We now turn to the interaction between subsidy design and pricing. We begin with the market outcome under no policy intervention, which establishes the baseline that motivates the use of subsidies (Section 4.3.1). We then activate both instruments and vary the swap price to examine how the optimal subsidy mix and the required budget respond to pricing (Section 4.3.2). Finally, we fix the price and vary the public budget to study how adoption and the two instruments evolve as support increases (Section 4.3.3).

4.3.1. Zero subsidy

We first characterize the market equilibrium in the absence of policy intervention. This baseline identifies the conditions under which ZE adoption and infrastructure deployment can emerge through pricing alone, and quantifies the coordination failure that



(a) ZE adoption by operator type vs. swap price

(b) Developer profit vs. swap price

Fig. 2. Price sensitivity analysis under zero subsidy.

motivates the subsidy analysis that follows. For any covered route r, the adoption condition derived in (5) reduces under zero subsidies ($\pi_F = 0$) to:

$$p \leq \frac{C_k^{\mathrm{DS}} t_r}{s_r}, \tag{23}$$

On the other hand, the developer will only deploy stations if the revenues cover the costs. From the L2 objective (6), the break-even condition under zero infrastructure subsidy is:

$$p\,B \geq \sum_i C_i^{\mathrm{BSS}} x_i + \sum_i C^{\mathrm{bat}} b_i, \tag{24}$$

where B is the total annual swap volume, and C_i^{BSS} and C^{bat} now denote the *total annualized* station and battery costs, each combining capital recovery with the corresponding operating expense.

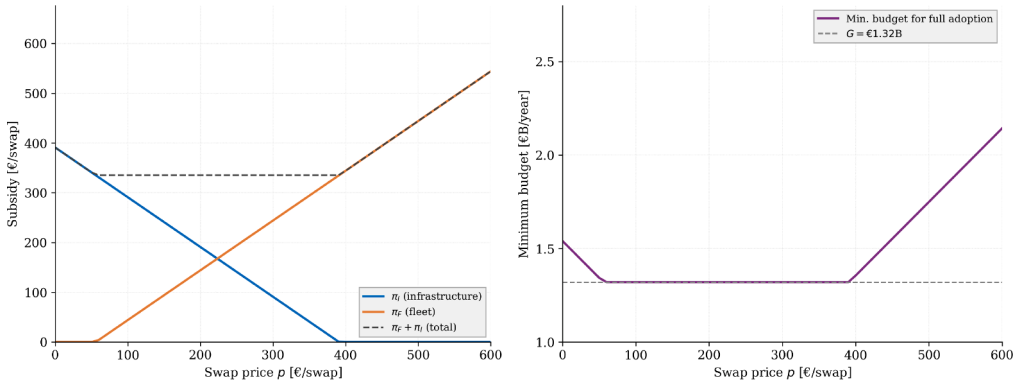
Fig. 2 illustrates the resulting price sensitivities on both sides of the market. Fig. 2a shows ZE adoption by operator type as the swap price rises. Adoption declines in a staircase pattern, and the order in which the operator types drop out follows their diesel costs. Standard operators, with the highest diesel costs, retain ZE operation longest, while recent entrants, with the lowest diesel costs, abandon it first. The spread of these thresholds is the demand-side expression of the cost heterogeneity built into the model. Fig. 2b shows the infrastructure developer side. The developer profit remains strictly negative across the entire price range. At prices low enough to induce operator adoption, swap revenues are insufficient to recover infrastructure and battery costs; at prices high enough to cover those costs, operator adoption collapses. The two viability conditions (23) and (24) therefore do not overlap at any price, confirming that a purely private-market equilibrium with battery-swapping deployment does not exist in the studied network. This coordination failure, where profitable operation and infrastructure cannot coexist at any single price, is precisely what the subsidy instruments introduced in the next section are designed to resolve.

4.3.2. Price sweep with subsidies at full adoption

We then activate both policy instruments (π_F, π_I) and solve the resulting bilevel problem using the iterative bilevel decomposition method. We aim to get full deployment and vary the swap price p to identify how the optimal subsidy mix and minimum required budget respond to pricing decisions. For each swap price level, we solve the bilevel problem with a sufficiently large budget to guarantee full adoption. Consistent with the lexicographic criterion (15), we then solve a secondary master problem that minimizes total public outlay, identifying the cheapest policy that sustains full adoption at each price.

Fig. 3a reveals a substitution pattern between the two instruments. As the swap price p increases, the optimal infrastructure subsidy π_I declines while the fleet subsidy π_F increases, and their sum $\pi_F + \pi_I$ remains broadly stable across the intermediate price range. A higher swap price strengthens the developer's revenue position and reduces the need for infrastructure support, but the operator then requires higher fleet-side compensation to preserve cost-competitiveness. The swap price thus redistributes public support between the two stakeholders without substantially altering the total per-swap requirement.

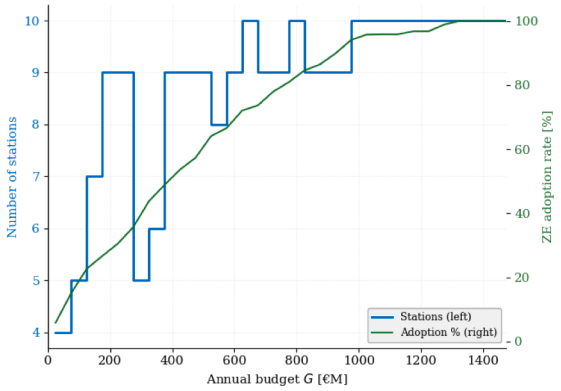
Fig. 3b reveals an asymmetric budget curve. The minimum budget required for full adoption stays broadly flat across the price range from approximately €60 to €390/swap, bottoming at €1,320M, and rises sharply outside it. Within this band both the developer's revenue position and operator cost-competitiveness are satisfied without excess public expenditure. Below it, prices are too low to sustain developer viability, and above it too high to maintain operator competitiveness. The swap price is therefore a key design variable, and the policymaker should target this intermediate band to minimize fiscal exposure while preserving full adoption. The reference price used in the remainder of the analysis, $p =$ €250/swap, lies inside this band, so the budget sweep that follows is conducted at a price that is already budget-efficient.



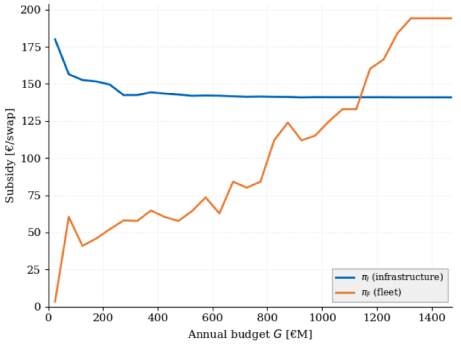
(a) Optimal subsidy instruments vs. swap price

(b) Minimum budget for full adoption vs. swap price

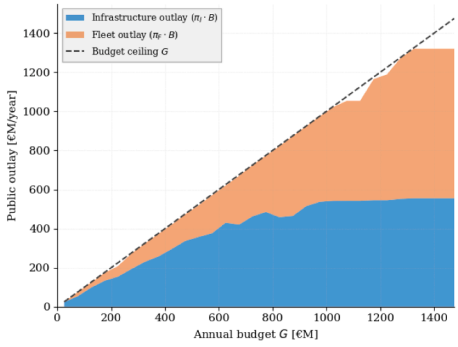
Fig. 3. Swap price sensitivity for full ZE adoption.



(a) ZE adoption and infrastructure scaling vs. budget



(b) Optimal subsidy instruments vs. budget



(c) Budget allocation vs. budget

Fig. 4. Budget sweep results at p =€250/swap. Results obtained via iterative bilevel decomposition.

4.3.3. Budget sweep with subsidies

Knowing the price sweep and the minimum budget to reach full adoption, we next examine how public expenditure can affect adoption. In this analysis, we identify the subsidies that secure participation, the fleet subsidy at the operator adoption floor and the infrastructure subsidy at the developer break-even, so that the available budget is spent expanding adoption. We set the swap price to p = €250/swap and vary the public budget G to determine how much adoption a given budget can reach. Table 2 presents the results of selected budget values.

Fig. 4a shows that ZE adoption increases and reaches full adoption at €1,320M. The curve exhibits several accelerations, reflecting discrete transitions in the infrastructure network as new stations unlock previously uncovered routes. The same panel shows that the number of open stations does not grow monotonically with budget. At several budget levels, the model substitutes a larger

Table 2
Budget sweep results at selected budget levels (p =€250/swap). Full adoption is reached at G =€1,325M.

Budget (G) [€M]	Trips	Adoption [%]	Budget Util. [%]	π_F [€/swap]	π_I [€/swap]	Swap Volume	Stations
50	112,528	9.9	99.9	60	161	225,056	4
100	199,931	17.6	99.5	37	159	508,370	8
200	334,695	29.5	100	50	150	1,002,172	9
300	446,773	39.3	99.8	58	145	1,476,528	8
400	588,775	51.8	99.8	60	143	1,959,044	8
500	697,901	61.4	99.6	63	141	2,438,036	7
700	858,905	75.6	99.5	79	142	3,165,560	9
900	982,466	86.5	98.3	128	141	3,291,002	9
1,000	1,076,333	94.8	99.0	118	141	3,818,994	10
1,325	1,135,905	100	99.6	194	141	3,938,330	10

number of lower-cost port locations with a smaller set of higher-throughput major ports that collectively serve more demand. These infrastructure-reconfiguration events appear as reductions in station count accompanied by increases in adoption rate, underscoring the value of jointly optimizing infrastructure siting and subsidy design. This non-monotonic behavior is consistent with multi-instrument policy models, which likewise exhibit regime shifts where adoption can dip as support reallocates between instruments (Brozynski and Leibowicz, 2022).

The evolution of the two subsidy instruments in Fig. 4c shows a clear sequencing. At low budgets the infrastructure subsidy π_I carries most of the support, since the binding constraint is making the first high-throughput stations commercially viable for the developer. As the network expands, π_I stabilizes and further adoption must be drawn in through the fleet subsidy π_F, which rises to bring progressively more cost-sensitive operators over their adoption thresholds. Fig. 4d confirms this by decomposing the budget allocation. In the early phase, spending is infrastructure-dominated and shifts toward a growing fleet-subsidy share as adoption approaches saturation. This infrastructure-before-fleet sequencing is supported by the broader literature on electric vehicle adoption, where infrastructure and demand subsidies act as substitutes and optimal support shifts toward the demand side as deployment matures, and where station subsidies are most cost-effective early but lose that edge as spending grows (Springel, 2021). It also aligns with the historical pattern in which transport infrastructure has preceded vehicle adoption (Leibowicz, 2018).

4.4. Sensitivity analysis

We now test the robustness of the model to cost and technology parameters that are most uncertain in practice. Each analysis changes one parameter while holding all others at their grounded baseline values. We report the resulting ZE adoption rate as a function of the annual public budget G at the reference swap price p = €250/swap. We examine cost sensitivity (Section 4.4.1) and battery-range sensitivity (Section 4.4.2), followed by an alternative emissions-minimizing objective (Section 4.4.3) and network-size sensitivity (Section 4.4.4).

4.4.1. Cost sensitivity

We test the two cost parameters that most directly govern the adoption decision, the annualized battery-pack cost C^{bat}, shouldered by the developer, and the diesel operating cost C_k^{DS}, which sets each operator's reference cost for ZE adoption. Operating cost scales with capital cost so that each scenario represents a coherent cost level rather than an isolated component shift. Fig. 5 reports both sweeps.

Fig. 5(a) shows that battery cost is a determinant of how expensive adoption is to achieve. Halving the battery cost moves the adoption frontier sharply leftward. Full adoption is reached at roughly €700M, compared with approximately €1,320M in the baseline. The cost-increase scenarios show that at +50% and +100% battery cost, full adoption is not attained anywhere in the tested budget range, achieving near 85% and 65% respectively at €1,500M budget.

Fig. 5(b) shows the effect on the operator side. Higher diesel costs can widen the cost advantage of zero-emission operation and shift the adoption frontier leftward, so that any given adoption target is reached at lower public cost. The effect, however, exhibits diminishing returns. The +50% and +100% operating diesel cost curves are close to one another and converge to full adoption at a similar budget, indicating that once diesel costs are high enough to push most operators past their adoption thresholds, further increases yield little additional uptake at the margin. The low-fuel cost scenario (−30%) moves the frontier rightward, requiring the largest budget to reach full adoption. This confirms that battery cost, not diesel cost, is therefore the dominant lever on the public budget required for full electrification, reinforcing the infrastructure financing gap as the principal barrier to early deployment. Public support may be most effective when it lowers infrastructure cost directly, through grants or battery-leasing guarantees, since the effect of higher diesel costs saturates once most operators pass their adoption thresholds. This saturation is consistent with empirical evidence on electric vehicles, where higher environmental tax revenues have not translated into greater uptake (Neves and Marques, 2025).

4.4.2. Battery range sensitivity

We next examine how battery energy density and technology improvement affect the required public investment. The usable range per pack determines the per-leg pack requirement n_r and hence the swap requirement s_r. A shorter range forces more swaps per trip,

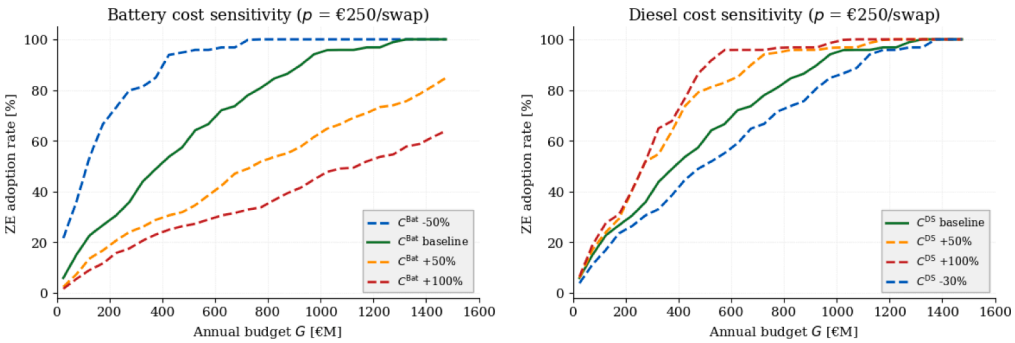


Fig. 5. Cost sensitivity of ZE adoption versus annual budget G. Left: battery-pack cost scenarios; Right: diesel operating-cost scenarios.

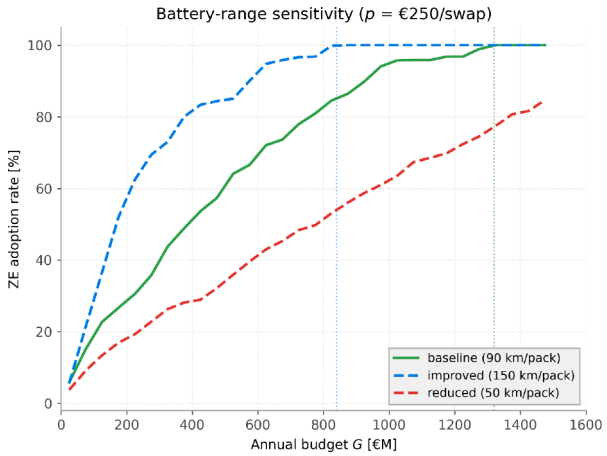


Fig. 6. ZE adoption rate versus annual budget G for three battery-range scenarios. Dashed vertical lines indicate the budget required for full adoption; the reduced-range scenario does not reach full adoption within the tested range.

raising the swap volume the developer must serve and the subsidy needed to make adoption attractive. We compare three scenarios: an improved battery ($R_b = 150$ km/pack), the baseline ($R_b = 90$ km/pack), and a reduced battery ($R_b = 50$ km/pack). Across three scenarios, the station sets S_r are held fixed, and n_r is recomputed from the longest sub-leg ℓ_r. At 150 km/pack it falls to a single pack on most routes; at 50 km/pack it rises to three on most routes, and to four wherever ℓ_r exceeds 150 km. The reduced-range scenario therefore reflects higher swap and inventory requirements rather than infeasibility.

Fig. 6 shows a consistent shift in the adoption frontier with range. The improved battery reaches full adoption at approximately €850M, well below the baseline requirement of €1,320M, because fewer swaps per trip reduce both the developer's inventory burden and the per-trip subsidy required. The reduced-battery scenario moves in the opposite direction and, notably, does not reach full adoption within the tested budget range, attaining only about 84.6% at €1,475M. This is the expected consequence of the energy-balance logic. Under degraded range, n_r and s_r increase, vessels carry additional packs and require more swaps per trip, and the public cost of closing the financing gap rises accordingly. For policymakers, this suggests that supporting battery technology and energy-density improvements, can reduce the public budget required for a given adoption target as effectively as increasing the subsidy itself.

4.4.3. Adoption-emissions trade-off

As an alternative on the policymaker's objective, we consider the case where public funds are directed toward maximizing avoided emissions rather than raw adoption. While adoption maximization treats every zero-emission trip equally, a welfare-oriented planner may instead prioritize trips that displace the most diesel, since these differ substantially across routes and operators. Longer routes and less fuel-efficient vessels avoid more CO_2 per converted trip. To capture this, we weigh each zero-emission trip by its avoided emissions $e_{rk} = c_k d_r \, EF$, where c_k is the operator-specific fuel consumption, d_r the route distance, and EF the diesel emission factor (D). The L1 objective then becomes

$$\max_{\pi_F, \pi_I} \sum_{k \in \mathcal{K}} \sum_{r \in \mathcal{R}} e_{rk} \, q^G_{rk}, \tag{25}$$

This experiment changes only the upper-level selection criterion. The constraints, column pool, the developer break-even and operator participation conditions remain unchanged. The framework therefore accommodates a different policy objective without modification to its lower levels and still produces a stable trade-off across the full budget range.

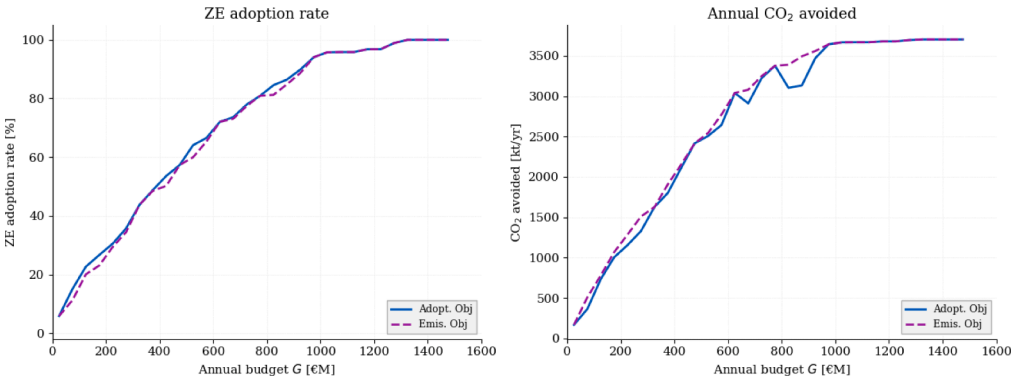


Fig. 7. ZE adoption rate (left) and annual CO_2 avoided (right) versus annual budget G under the adoption and avoided-emissions objectives.

Table 3
Maximizing adoption versus maximizing avoided emissions across the budget sweep (p =€250/swap). At each budget the avoided-emissions objective gives up little adoption while displacing more CO_2, reconfiguring the network toward longer, higher-emission routes.

	Adoption [%]		CO_2 avoided [kt/yr]			Stations	
G [€ M]	Adopt.	Emis.	Adopt.	Emis.	Δ%	Adopt.	Emis.
75	15.1	11.2	363.5	508.9	+40.0	5	7
175	26.7	23.1	1,003.9	1,067.8	+6.4	9	8
275	35.7	34.4	1,329.1	1,509.2	+13.5	5	9
425	53.6	50.2	2,112.1	2,153.3	+2.0	9	8
525	64.1	60.0	2,506.1	2,543.0	+1.5	8	10
675	73.6	73.1	2,910.5	3,079.7	+5.8	9	10
875	86.5	84.8	3,132.9	3,493.2	+11.5	9	10
1325	100.0	100.0	3,702.7	3,702.7	+0.0	10	10

Table 3 presents the comparison across the budget sweep, evaluated at eight representative levels. At every budget the two objectives allocate similar public expenditure to different sets of trips. Adoption maximization favors the cheapest conversions, which are the short routes where the swap requirement is small and the swap volume per converted trip is low. On the other hand, the avoided-emissions objective favors the routes that displace the most diesel. These are the longer routes served by older and less efficient vessels. These trips are more expensive to serve, so fewer of them are converted at a given budget. The adoption gap between the two objectives is small throughout, never exceeding 4.1 percentage points (at G =€525M). The two objectives select an identical network, so the gap closes entirely. The gap stays small because operator cost structures do not differ greatly, so changing the objective affects which trips are served more than how many.

The trade-off is largest when the budget is most constrained. At G =€75M, avoided emissions rise from 363.5 to 508.9 kt/yr, a gain of 40.0%, for a loss of only 3.9 percentage points in adoption. The largest absolute gain occurs at G =€875M, where the emissions objective avoids 3,493.2 kt/yr against 3,132.9 kt/yr under adoption maximization, a difference of about 360 kt/yr, for a loss of just 1.7 percentage points in adoption. As the left panel of Fig. 7 shows, the two adoption curves remain close across the entire budget range. Switching the objective changes the headline adoption rate only marginally. The network beneath it is reconfigured rather than simply enlarged, because the high-emission routes are longer and connect more distant ports, the emissions objective re-sites stations to reach them, which at some budget levels means opening more stations than the adoption objective and at others means reaching a comparable adoption rate with a different, not strictly larger, set of stations. The choice between the two objectives therefore depends on the policy goal rather than on the model. A planner aiming for broad fleet transition and network coverage is served by the adoption objective, while a planner with an explicit climate target can shift deployment toward high-emission routes and avoid substantially more CO_2 at a modest adoption cost, particularly at low-to-moderate budget levels.

4.4.4. Network size

Lastly, we examine how the number of deployed swap stations affects adoption and fiscal requirements under a fixed budget. We force exactly n stations to be open and allow the model to select the best port locations and subsidy mix for each network size. This mimics a phased deployment scenario where the developer commits to a fixed number of stations per period and seeks to maximize adoption within the available budget. Results are reported in Table 4.

Table 4 shows a clear pattern of diminishing marginal returns as the network expands. Early deployment of stations deliver the largest adoption gain, and for each subsequent station, it contributes progressively less. The model selects the highest-throughput port locations first, with each additional station serving a progressively smaller share of uncovered demand. The per-swap subsidy instruments remain stable throughout. This confirms that subsidy design does not need to be redesigned as the network grows, only

Table 4
Network configuration sensitivity results (p =€250/swap).

n	Adoption [%]	ΔAdoption [%]	π_F [€/swap]	π_I [€/swap]	Batteries	Budget used [€B]
2	26.8	—	87	138	6,084	0.27
3	41.3	14.5	91	139	9,384	0.44
4	53.5	12.2	133	140	11,152	0.61
5	63.9	10.4	194	139	13,948	0.93
6	74.4	10.5	194	140	15,813	1.06
7	84.1	9.7	194	139	17,358	1.16
8	91.8	7.7	194	138	18,455	1.23
9	96.7	4.9	194	138	19,249	1.28
10	100.0	3.3	194	138	19,692	1.32

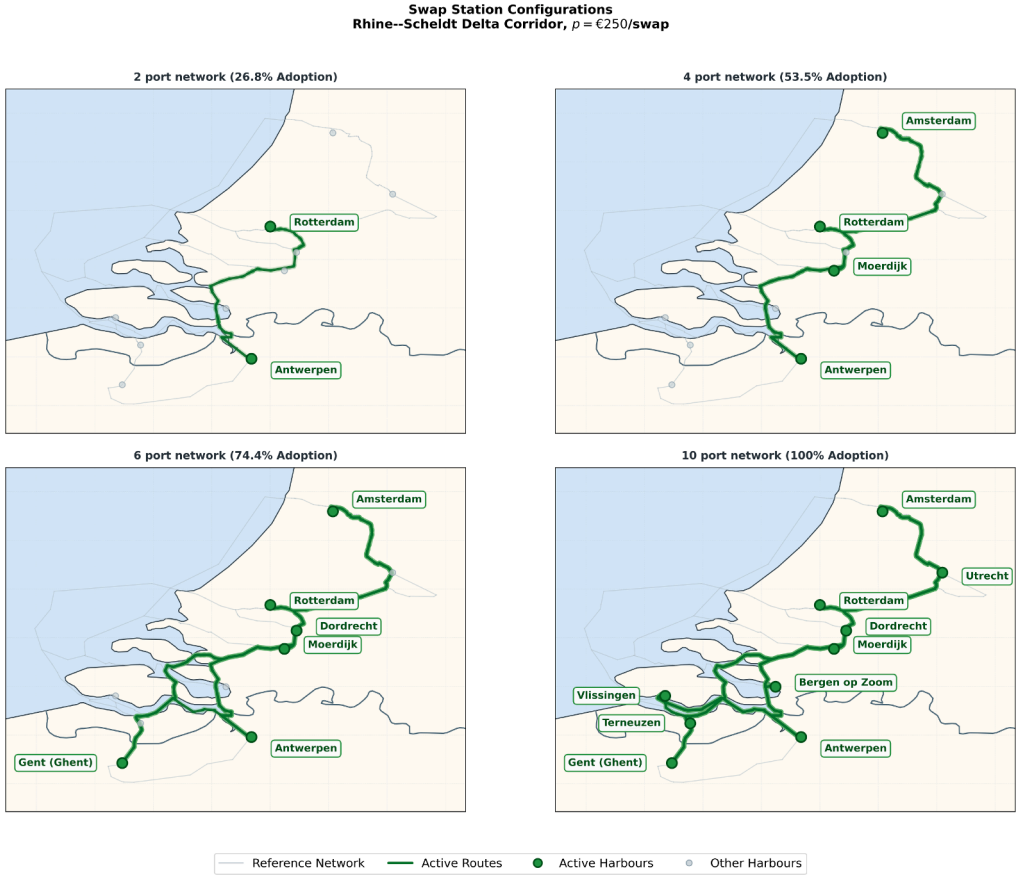


Fig. 8. Network-size sensitivity - maximum achievable adoption and minimum budget required (p =€250/swap).

the public expenditure increases. Results also indicate that in an early deployment phase, meaningful adoption can be achieved well below the committed budget ceiling.

Fig. 8 visualizes the evolution of the swap network across four selected deployment scenarios. Starting with only two stations, the model concentrates coverage on the high-volume Rotterdam–Antwerpen corridor, which alone accounts for 26.8% of total fleet trips. Expanding to four stations adds Amsterdam and Moerdijk, extending coverage and raising adoption to 53.5%. At six stations, Dordrecht and Gent are incorporated, unlocking the Belgian inland waterway connections and pushing adoption to 74.4%. Full adoption requires all ten stations that cover the remaining routes. The sequential port selection follows a clear geographic logic and corresponds directly to the diminishing marginal adoption gains visible in Table 4.

5. Conclusion

This paper presents a multi-level optimization framework for the deployment of containerized battery swapping in zero-emission inland waterway transport. The model captures the strategic interaction among a policymaker who designs subsidy instruments

under a budget constraint, an infrastructure developer who sets station locations and battery inventory, and fleet operators who decide whether to adopt. Applied to a Dutch–Belgian case study, the framework yields the following findings. First, no pure market equilibrium exists, so public intervention is required to deploy the system at all. Second, the swap price governs how support is split between the two instruments rather than how much is needed, with the minimum budget for full adoption stable across an intermediate price band. Third, the optimal subsidy mix depends on the available budget. Infrastructure-dominated when budgets are limited, shifting toward fleet support as adoption approaches saturation. Lastly, the sensitivity analysis identifies battery cost and range as the dominant levers on the required budget.

The results suggest a few practical directions, though the model provides a snapshot of the optimal network and subsidy mix at a given budget rather than a time-phased plan. First, since the optimal split between instruments depends on the budget, support need not be fixed at a single split. It is best concentrated on infrastructure when budgets are limited and shifted toward the fleet subsidy as adoption grows. Second, the swap price is a useful design lever, because the minimum budget is stable only within an intermediate band, contractual arrangements such as long-term agreements between the developer and operators that keep the price within this band could reduce the public support required. Third, because battery cost and range are the dominant budget levers, support for battery cost reduction and energy-density improvement can complement direct subsidy. Since the model treats unit costs as fixed, scaling up deployment may itself lower these costs through volume and learning effects.

6. Limitations and recommendations

Several limitations of the model suggest directions for improvement. The adoption mechanism models operator decisions as a binary cost comparison. A more realistic treatment would model operators as utility maximizers with heterogeneous preferences considering risk aversion and capital constraints. The framework also models only a single representative infrastructure developer. Extending it to competition between multiple infrastructure providers, with strategic pricing and coverage, is a natural direction as the market matures. The single-period structure also cannot capture how infrastructure availability and adoption rates evolve together over time. A multi-period extension with dynamic demand and budget would allow evaluation of a phased plan and transition pathways.

Furthermore, the developer is modeled as committing to its network before operators respond. In some cases, infrastructure and adoption are coordinated through long-term agreements or joint ventures between different actors. We leave such partnership-based structures for future work. In addition, both subsidy instruments are modeled in per-swap form. Comparing them against lump-sum construction subsidies can also be a further extension. The framework also does not consider the institutional, regulatory, and funding arrangements that sit above this optimization problem, which are interesting future work directions.

In practice, directional flow imbalances and the choice of where to swap en route jointly affect how many batteries are needed and where stations should be placed. Developing tactical and voyage-level decisions such as scheduling and routing-based swap-location optimization would help in determining ideal mid-route station location that can be used as inputs for the model. Finally, the cost and technical parameters, while drawn from industry and regulatory sources, remain estimates rather than figures observed from operational deployments, which are not yet available at scale in IWT. Calibrating against realized swap prices, trip records, and adoption behavior as such systems mature would sharpen the quantitative findings.

CRediT authorship contribution statement

V.M.J. Ramos: Writing – review & editing, Visualization, Validation, Methodology, Formal analysis, Data curation, Conceptualization; **Bilge Atasoy:** Writing – review & editing, Validation, Supervision, Methodology, Funding acquisition, Conceptualization.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the author used Claude (Anthropic) to improve the language and readability of the manuscript. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This research was conducted within the project PATH2ZERO, which is financed by the NWA L2-Thema 2020 Zero Emission Shipping (ZES) program of the Netherlands Organization for Scientific Research (NWO) with Grant NWA.1439.20.001. Partly funded by the European Union (ERC, ADAPT-OR, No 101117675). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

Data availability

Data will be made available on request.

Appendix A. KKT reformulation of the fleet operator problem

This appendix derives the KKT conditions for the fleet operator problem (L3) and shows how they are embedded into the infrastructure developer problem (L2) to form a single-level mathematical program with equilibrium constraints (MPEC). The reformulation treats each route-operator pair (r,k) independently, reflecting operator heterogeneity in diesel costs, and indexes the station feasibility constraints over the route-specific required station set S_r.

Since L3 is a linear program with a convex feasible region, the Karush–Kuhn–Tucker conditions are both necessary and sufficient for global optimality. Replacing the explicit L3 problem with its KKT system is therefore an exact equivalence: any primal-dual solution satisfying the KKT conditions below is an optimal solution to L3, and conversely. This guarantees that the resulting MPEC faithfully represents the joint behavior of the infrastructure developer and the fleet operators.

A.1. Fleet operator problem (L3)

For given infrastructure configuration (x,b) and subsidies (π_F,π_I), each operator type $k\in\mathcal{K}$ independently minimizes transport costs on each route $r\in\mathcal{R}$:

$$\min_{q^G_{r,k},\,q^{DS}_{r,k}} \quad C^{\mathrm{DS}}_k\, t_r\, q^{DS}_{r,k} + (p-\pi_F)\, s_r\, q^G_{r,k} \tag{A.1a}$$

$$\text{s.t.} \quad q^{DS}_{r,k} + q^G_{r,k} = \tau_{r,k} \tag{A.1b}$$

$$q^G_{r,k} \le \tau_{r,k}\, x_i, \quad \forall i\in S_r \tag{A.1c}$$

$$q^G_{r,k}\ge 0, \quad q^{DS}_{r,k}\ge 0 \tag{A.1d}$$

A.2. Lagrangian and KKT conditions

Assign dual multipliers to each constraint: an unrestricted multiplier $\lambda_{r,k}\in\mathbb{R}$ for the demand-balance constraint (A.1b); a non-negative multiplier $\alpha_{r,k,i}\ge 0$ for each station feasibility constraint (A.1c), one per $i\in S_r$; and non-negative multipliers $\eta_{r,k}\ge 0$ and $\mu_{r,k}\ge 0$ for the non-negativity constraints on $q^G_{r,k}$ and $q^{DS}_{r,k}$ respectively:

$$\begin{aligned} q^{DS}_{r,k} + q^G_{r,k} - \tau_{r,k} &= 0 \quad (\lambda_{r,k}\in\mathbb{R}),\\ q^G_{r,k} - \tau_{r,k}\, x_i \le 0, \quad \forall i\in S_r &\quad (\alpha_{r,k,i}\ge 0),\\ -q^G_{r,k} &\le 0 \quad (\eta_{r,k}\ge 0),\\ -q^{DS}_{r,k} &\le 0 \quad (\mu_{r,k}\ge 0). \end{aligned}$$

The Lagrangian for route-operator pair (r,k) is:

$$\begin{aligned} \mathcal{L}_{r,k} = {} & C^{\mathrm{DS}}_k t_r\, q^{DS}_{r,k} + (p-\pi_F) s_r\, q^G_{r,k}\\ & + \lambda_{r,k}(\tau_{r,k} - q^{DS}_{r,k} - q^G_{r,k})\\ & + \sum_{i\in S_r} \alpha_{r,k,i}(q^G_{r,k} - \tau_{r,k}\, x_i)\\ & - \eta_{r,k}\, q^G_{r,k} - \mu_{r,k}\, q^{DS}_{r,k}. \end{aligned}$$

The KKT conditions for each (r,k) are:

Stationarity:.

$$\frac{\partial \mathcal{L}_{r,k}}{\partial q^{DS}_{r,k}} : \quad C^{\mathrm{DS}}_k t_r - \lambda_{r,k} - \mu_{r,k} = 0 \tag{A.2}$$

$$\frac{\partial \mathcal{L}_{r,k}}{\partial q^G_{r,k}} : \quad (p-\pi_F)s_r - \lambda_{r,k} + \sum_{i\in S_r}\alpha_{r,k,i} - \eta_{r,k} = 0 \tag{A.3}$$

Primal and dual feasibility:. Constraints (A.1b)–(A.1d) and $\alpha_{r,k,i}\ge 0$ for all $i\in S_r$, $\eta_{r,k},\mu_{r,k}\ge 0$.

Complementary slackness:.

$$\alpha_{r,k,i} \cdot (\tau_{r,k}\, x_i - q^G_{r,k}) = 0, \quad \forall i \in S_r \tag{A.4}$$

$$\eta_{r,k} \cdot q^G_{r,k} = 0 \tag{A.5}$$

$$\mu_{r,k} \cdot q^{DS}_{r,k} = 0 \tag{A.6}$$

A.3. MPEC reformulation: embedding L3 into L2

The infrastructure developer problem (L2) maximizes profit subject to capacity constraints and the requirement that operator choices (q^G, q^{DS}) satisfy the KKT conditions of L3. Since these conditions are necessary and sufficient for L3 optimality, the resulting MPEC is equivalent to the original bilevel L2–L3 problem:

$$\max_{x,b,q^G,q^{DS},\lambda,\alpha,\eta,\mu} \quad \sum_{r\in\mathcal{R}}\sum_{k\in\mathcal{K}} (p+\pi_I)\, s_r\, q^G_{r,k} - \sum_{i\in\mathcal{I}} \left[C^{\text{BSS}}_i\, x_i + C^{\text{bat}}\, b_i\right] \tag{A.7a}$$

$$\text{s.t.} \quad \sum_{k\in\mathcal{K}} \sum_{r:\, i\in S_r} \frac{s_r}{|S_r|}\, q^G_{r,k} \le \delta\, b_i, \quad \forall i \in \mathcal{I} \tag{A.7b}$$

$$\sum_{k\in\mathcal{K}} \sum_{r:\, i\in S_r} \frac{s_r}{|S_r|}\, q^G_{r,k} \le SW^{\max}_i\, x_i, \quad \forall i \in \mathcal{I} \tag{A.7c}$$

$$BSS^{\min} \le \sum_{i\in\mathcal{I}} x_i \le BSS^{\max} \tag{A.7d}$$

$$x_i \in \{0,1\}, \quad b_i \ge 0, \quad \forall i \in \mathcal{I} \tag{A.7e}$$

$$\text{KKT conditions (A.1b), (A.1c), (A.1d), (A.2) – (A.6)} \quad \forall (r,k) \in \mathcal{R}\times\mathcal{K} \tag{A.7f}$$

A.4. Linearization of complementarity constraints

The bilinear complementarity conditions (A.4)–(A.6) are nonlinear and must be linearized before the MPEC can be solved by a mixed-integer linear programming solver. For each (r,k), we introduce binary indicator variables $z^{\alpha}_{r,k,i} \in \{0,1\}$ for each $i \in S_r$, and $z^{\eta}_{r,k}, z^{\mu}_{r,k} \in \{0,1\}$, yielding the equivalent linear conditions:

$$z^{\alpha}_{r,k,i} = 1 \;\Rightarrow\; q^G_{r,k} = \tau_{r,k}\, x_i, \quad \forall i \in S_r \tag{A.8}$$

$$z^{\alpha}_{r,k,i} = 0 \;\Rightarrow\; \alpha_{r,k,i} = 0, \quad \forall i \in S_r \tag{A.9}$$

$$z^{\eta}_{r,k} = 1 \;\Rightarrow\; q^G_{r,k} = 0 \tag{A.10}$$

$$z^{\eta}_{r,k} = 0 \;\Rightarrow\; \eta_{r,k} = 0 \tag{A.11}$$

$$z^{\mu}_{r,k} = 1 \;\Rightarrow\; q^{DS}_{r,k} = 0 \tag{A.12}$$

$$z^{\mu}_{r,k} = 0 \;\Rightarrow\; \mu_{r,k} = 0 \tag{A.13}$$

Indicator constraints of this form are natively supported by modern MILP solvers such as Gurobi, which enforce the logical implications internally without requiring an explicit penalty parameter. This avoids the calibration of a finite big-M constant, the conventional alternative for linearizing complementarity, in which each condition $\mu \cdot g(x) = 0$ is replaced by $\mu \le Mz$ and $g(x) \le M(1-z)$. A valid M must bound both the dual and primal terms simultaneously, and since these differ in magnitude in our model, any single constant large enough to be valid is loose for most pairs, weakening the linear relaxation.The indicator formulation removes this risk, at the cost of a specialized branching procedure that is marginally slower on small instances but proved negligible across all instances tested. The resulting problem is a mixed-integer linear program with $|\mathcal{R}| \times |\mathcal{K}|$ sets of KKT conditions, each containing $|S_r| + 2$ complementarity pairs.

Appendix B. Sample Route Inputs (top 10) and Mid-route swap station definition

Table B.1
Route characteristics (origin–destination, total trips, distance, and swaps).

Route	Origin	Destination	Total trips $\sum_k \tau_{rk}$	Distance (km) d_r	Swap requirement s_r
R1	*Rotterdam*	*Antwerp*	154,343	117	4
R2	*Antwerp*	*Rotterdam*	152,615	117	4
R3	*Amsterdam*	*Rotterdam*	87,995	114	4
R4	*Rotterdam*	*Amsterdam*	78,531	114	4
R5	*Antwerp*	*Ghent*	36,738	66	2
R6	*Antwerp*	*Amsterdam*	34,738	198	6
R7	*Vlissingen*	*Antwerp*	34,220	64	2
R8	*Amsterdam*	*Antwerp*	31,266	198	6
R9	*Rotterdam*	*Moerdijk*	29,577	42	2
R10	*Moerdijk*	*Antwerp*	27,141	75	2

Swap requirements follow $s_r = |S_r|\, n_r$ with $n_r = \lceil \ell_r / R_b \rceil$ (C). Routes served only at their endpoints have $|S_r| = 2$ and $\ell_r = d_r$; the corridors with an intermediate swap station (R6 and R8 above) have $|S_r| = 3$ and are listed in Table B.2.

Table B.2
Routes with an exogenous mid-route swap station.

Route	Distance (km)	Mid-route station	Sub-legs (km)	ℓ_r (km)	n_r
Amsterdam ↔ Antwerp	198	Dordrecht	107 / 91	107	2
Amsterdam ↔ Ghent	235	Moerdijk	125 / 110	125	2
Amsterdam ↔ Terneuzen	214	Bergen op Zoom	172 / 42	172	2
Amsterdam ↔ Vlissingen	224	Moerdijk	125 / 99	125	2
Utrecht ↔ Ghent	190	Bergen op Zoom	125 / 65	125	2

The mid-route station applies to both directions, giving $S_r = \{o_r, m_r, d_r\}$ and $n_r = \lceil \ell_r / R_b \rceil = 2$, hence $s_r = |S_r|\, n_r = 6$ (Table B.1). All other routes use $S_r = \{o_r, d_r\}$, with $\ell_r = d_r$.

Appendix C. Battery-swap requirement and station sets

Each zero-emission one-way trip requires a battery inventory determined by the distance the vessel sails between two consecutive exchanges and the single-pack range $R_b = 90$ km. Let S_r denote the set of battery-swap stations visited by route r, ordered along the sailing path, and let

$$\ell_r = \max_j d\left(i_j, i_{j+1}\right), \qquad i_j \in S_r, \tag{C.1}$$

denote the longest sub-leg between two consecutive stations. Since the vessel receives no intermediate replenishment within a sub-leg, the on-board inventory must cover ℓ_r, so the vessel carries

$$n_r = \left\lceil \frac{\ell_r}{R_b} \right\rceil, \tag{C.2}$$

where n_r is the minimum integer number of battery packs carried on board. At each visited station the vessel exchanges its entire onboard inventory. Hence, the total number of battery swaps per one-way trip is

$$s_r = |S_r|\, n_r, \tag{C.3}$$

where $|S_r|$ is the number of swap stations on the route.

For most routes, battery exchanges occur only at the origin and destination, so $S_r = \{o_r, d_r\}$, $\ell_r = d_r$, and $s_r = 2n_r$. For the five longest corridors, spanning the full route in a single leg would require three packs on board, which exceeds the container slots realistically available on a vessel. An exogenous mid-route station m_r is therefore introduced (Table B.2), giving $S_r = \{o_r, m_r, d_r\}$ and $s_r = 3n_r$. The intermediate stop shortens the longest sub-leg to $\ell_r \le 2R_b$, reducing the on-board requirement to $n_r = 2$ packs and giving $s_r = 6$ on these corridors. A route can be operated only if every station in S_r is opened. Consequently, each station handles exactly n_r battery swaps per one-way trip.

We assume a uniform pack range across operator types, swaps as instantaneous discrete events, and swap demand allocated uniformly across a route's required ports. In the baseline ($R_b = 90$ km) no sub-leg in the case-study network exceeds $2R_b = 180$ km, so $n_r \le 2$ throughout: the two-pack limit is a property of the network under the given station sets rather than an imposed constraint. Under the reduced-range scenario ($R_b = 50$ km) the same station sets are retained and n_r rises to three packs on most routes, and to four on those whose longest sub-leg exceeds 150 km, reflecting the additional container slots required at degraded energy density. The pack range and battery costs are varied in the sensitivity analysis (Section 4.4).

Appendix D. Operator diesel cost and avoided-emission coefficients

The operator-specific diesel sailing cost C_k^{DS} and the per-trip avoided-emission coefficient e_{rk} are both derived from a single specific-fuel-consumption (SFC) basis, so that an operator's private fuel cost and the social CO_2 benefit of converting its trips remain mutually consistent.

Let P be the laden service power, v the cruising speed, ρ the diesel density, p_{f} the diesel price, EF the combustion emission factor, and SFC_k the engine-age-specific fuel consumption of operator type k. The distance-based fuel intensity of a type-k vessel is

$$c_k = \frac{\mathrm{SFC}_k\, P}{1000\, \rho\, v} \qquad [\mathrm{L/km}], \tag{D.1}$$

where the factor 1000 converts grams to kilograms. Since SFC_k increases with engine age, older operator classes have a larger c_k, and hence both a higher diesel cost and a larger avoided emission per converted trip. The diesel sailing cost per operating hour and the CO_2 avoided per converted trip on route r then follow directly from c_k:

$$C_k^{\mathrm{DS}} = c_k\, v\, p_{\mathrm{f}} \quad [€/\mathrm{h}], \qquad e_{rk} = c_k\, d_r\, \mathrm{EF} \quad [\mathrm{kg\ CO_2/trip}]. \tag{D.2}$$

The diesel cost of serving route r with operator k is $C_k^{\mathrm{DS}}\, t_r$, entering the operator (L3) problem through the route travel time t_r. The total CO_2 avoided by an infrastructure–fleet configuration (column) ℓ is

$$E_\ell = \sum_{r\in R} \sum_{k\in K} e_{rk}\, q^{\mathrm{G}}_{rk,\ell}, \tag{D.3}$$

the objective coefficient used in the emission-weighted sensitivity analysis.

D.1. Parameter values and sources

Table D.1
Cost, technical, and emission parameters used in the case study.

Symbol	Description	Value	Source
Infrastructure and battery cost parameters			
C_i^{BSS} (major)	Annualized swap station cost, major ports	1.5 €M/yr	Liu et al. (2026)
C_i^{BSS} (other)	Annualized station cost, other ports	1.0 €M/yr	
C^{bat}	Annualized battery pack cost[a]	€77,500 pack/yr	Meherishi et al. (2026)
δ	Battery cycle capacity	200 swaps/pack/yr	EPEVER (2025)
R_b	Usable range per battery pack	90 km	Zero Emission Services (2021b)
Diesel-cost and emission parameters			
P	Laden service power	500 kW	Meherishi et al. (2026)
v	Cruising speed	15 km/h	
ρ	Diesel density	0.85 kg/L	Netherlands Enterprise Agency (2025)
p_{f}	Diesel price	0.55 €/L	Central Commission for the Navigation of the Rhine (CCNR)
EF	Combustion emission factor	2.68 kg CO_2/L	Netherlands Enterprise Agency (2025)
$\mathrm{SFC}_{\mathrm{new}}$	Specific fuel consumption, recent engines	200 g/kWh	TNO (2009)
$\mathrm{SFC}_{\mathrm{modern}}$	SFC, modern engines	215 g/kWh	TNO (2009)
$\mathrm{SFC}_{\mathrm{standard}}$	SFC, standard engines	235 g/kWh	TNO (2009)

Table footnote [a]Annualized by the authors from the per-pack capital cost reported in Meherishi et al. (2026), assuming a 10-year service life and a 5% discount rate, and operating expenditure equal to 5% of capex.

All parameters are held fixed across the analysis. For the cost and battery-range scenarios in Section 4.4, we vary the relevant subset.

References

Amiri, A., Zolfagharinia, H., Hassanzadeh Amin, S., 2023. Routing a mixed fleet of conventional and electric vehicles for urban delivery problems: considering different charging technologies and battery swapping. Int. J. Syst. Sci.: Oper. Logist. 10 (1), 2191804. https://doi.org/10.1080/23302674.2023.2191804

Brozynski, M.T., Leibowicz, B.D., 2022. A multi-level optimization model of infrastructure-dependent technology adoption: overcoming the chicken-and-egg problem. Eur. J. Oper. Res. 300, 755–770. https://doi.org/10.1016/j.ejor.2021.10.026

Capelle, T., Cortés, C.E., Gendreau, M., Rey, P.A., Rousseau, L.M., 2019. A column generation approach for location-routing problems with pickup and delivery. Eur. J. Oper. Res. 272 (1), 121–131. https://doi.org/10.1016/j.ejor.2018.05.055

Central Commission for the Navigation of the Rhine, 2021. CCNR roadmap for reducing inland navigation emissions. https://www.ccr-zkr.org/files/documents/Roadmap/Roadmap_en.pdf.

Central Commission for the Navigation of the Rhine, 2024. 2. freight transport on inland waterways. https://inland-navigation-market.org/chapitre/2-transport-fluvial-de-marchandises/?lang=en. Inland Navigation in Europe - Market Observation.

Central Commission for the Navigation of the Rhine(CCNR), 2025. Inland navigation in Europe — market insight.

Central Commission for the Navigation of the Rhine (CCNR), 2022. 2. CCNR market observation - annual report 2022. https://www.ccr-zkr.org/files/documents/om/om22IIen.pdf.

Chen, R., Meng, Q., Yu, J.J., 2023. Optimal government incentives to improve the new technology adoption: subsidizing infrastructure investment or usage? Omega 114, 102740. https://doi.org/10.1016/j.omega.2022.102740

Cui, Z., Fu, X., Wu, X., Wang, K., 2025. Optimal subsidy scheme for electric ship adoption in inland shipping. Transp. Res. a 199, 104591. https://doi.org/10.1016/j.tra.2025.104591
European Commission, 2019. The European green deal.
European Commission, 2020. Sustainable and smart mobility strategy - putting european transport on track for the future.
European Environment Agency, 2025. Climate - sustainability of Europe's mobility systems. https://www.eea.europa.eu/en/analysis/publications/sustainability-of-europes-mobility-systems-2025/climate.
EPEVER, 2025. How long do lifePO batteries last in solar energy systems? EPEVER Articles. https://www.epever.com/articles/how-long-do-lifepo.
Ghanbari Motlagh, S., Oladigbolu, J., Li, L., 2025. A review on electric vehicle charging station operation considering market dynamics and grid interaction. Appl. Energy 392, 126058. https://doi.org/10.1016/j.apenergy.2025.126058
Gu, M., Wen, S., Zhu, M., Chen, C., Huang, Y., 2021. Voyage optimization for all-electric vessels integrated with swappable containerized battery. In: 2021 International Conference on Power Electronics Systems and Applications. IEEE, pp. 586–591. https://doi.org/10.1109/ICPES53652.2021.9683806
Guo, S., Dai, L., 2023. Voyage optimization for swappable battery-powered inland containerships: considering heterogenous flow velocity and multi-battery swapping. In: 2023 IEEE International Conference on Power Systems Technology (ICPS). IEEE, pp. 1–7. https://doi.org/10.1109/ICPSAsia58343.2023.10294788
Guo, S., Wang, Y., Yue, M., Dai, L., Fang, S., Zhang, S., Hu, H., 2025. Optimization of energy replenishment for inland electric ships considering multi-technology adoption and partial replenishment. J. Mar. Sci. Eng. 13 (11), 2092. https://doi.org/10.3390/jmse13112092
Han, F., Li, J., Luo, Y., Long, F., Ji, Y., Wang, Z., 2026. Breaking the coordination deadlock of hydrogen application in inland shipping: a four-stakeholder evolutionary game analysis. Ocean Eng. 355, 125053. https://doi.org/10.1016/j.oceaneng.2026.125053
Huang, S., Liu, K., Zhang, Z.H., 2023. Column-and-constraint-generation-based approach to a robust reverse logistic network design for bike sharing. Transp. Res. B: Methodol. 173, 90–118. https://doi.org/10.1016/j.trb.2023.04.010
Lamontagne, S., Carvalho, M., Frejinger, E., Gendron, B., Anjos, M.F., Atallah, R., 2023. Optimising electric vehicle charging station placement using advanced discrete choice models. INFORMS J. Comput. 35 (5), 1195–1213. https://doi.org/10.1287/ijoc.2022.0185
Leibowicz, B.D., 2018. Policy recommendations for a transition to sustainable mobility based on historical diffusion dynamics of transport systems. Energy Policy 119, 357–366. https://doi.org/10.1016/j.enpol.2018.04.066
Li, B., Xie, K., Zhong, W., Huang, X., Wu, Y., Xie, S., 2023. Operation management of electric vehicle battery swapping and charging systems: a bilevel optimization approach. IEEE Trans. Intell. Transp. Syst. 24 (1), 528–540. https://doi.org/10.1109/TITS.2022.3211883
Li, S., Tong, L., Xing, J., Zhou, Y., 2017. The market for electric vehicles: indirect network effects and policy design. J. Assoc. Environ. Resour. Econ. 4, 89–133. https://doi.org/10.1086/689702
Li, Z., Wei, Y., Park, J.H., 2024. An improved bilevel algorithm based on ant colony optimization and adaptive large neighborhood search for routing and charging scheduling of electric vehicles. IEEE Trans. Transp. Electrif. https://doi.org/10.1109/TTE.2024.3398113
Lienkamp, B., Schiffer, M., 2024. Column generation for solving large scale multi-commodity flow problems for passenger transportation. Eur. J. Oper. Res. 314, 703–717. https://doi.org/10.1016/j.ejor.2023.09.019
Limmer, S., 2023. Bilevel large neighborhood search for the electric autonomous dial-a-ride problem. Transp. Res. Interdiscip. Perspect. 21, 100876. https://doi.org/10.1016/j.trip.2023.100876
Liu, Y., He, J., Tan, C., 2026. Energy replenishment optimization for inland liner electric container ships using benders decomposition. Transp. Res. E 211, 104867. https://doi.org/10.1016/j.tre.2026.104867
Luo, T., Song, Y., Li, G., 2023. An agent-based simulation study for escaping the "chicken-egg" dilemma between electric vehicle penetration and charging infrastructure deployment. Resour. Conserv. Recycl. 194, 106966. https://doi.org/10.1016/j.resconrec.2023.106966
Meherishi, L., van den Berg, P.L., Zuidwijk, R., 2026. Battery replenishment and repositioning problem in electric barge networks. Eur. J. Oper. Res. 335 (1), 83–104. https://doi.org/10.1016/j.ejor.2026.02.018
Mitici, M., Pereira, M., Oliviero, F., 2022. Electric flight scheduling with battery-charging and battery-swapping opportunities. EURO J. Transp. Logist. 11, 100074. https://doi.org/10.1016/j.ejtl.2022.100074
Neves, S.A., Marques, A.C., 2025. What has driven the adoption of BEV and PHEV in the EU? Res. Transp. Bus. Manag. 60, 101331. https://doi.org/10.1016/j.rtbm.2025.101331
Netherlands Enterprise Agency(RVO), 2025. The netherlands: List of fuels, january 2025.
van Oosterom, S., Mitici, M., 2023. Optimizing the battery charging and swapping infrastructure for electric short-haul aircraft-the case of electric flight in norway. Transp. Res. C: Emerg. Technol. 155, 104313. https://doi.org/10.1016/j.trc.2023.104313
Port of Rotterdam Authority, 2021. First emission-free inland shipping vessel on energy containers in service. News press release, 6 September 2021. https://www.portofrotterdam.com/en/news-and-press-releases/first-emission-free-inland-shipping-vessel-on-energy-containers-in-service.
Piedra de la Cuadra, R., Ortega, F.A., 2024. Bilevel optimization for the deployment of refuelling stations for electric vehicles on road networks. Comput. Oper. Res. 162, 106460. https://doi.org/10.1016/j.cor.2023.106460
Pi, Z., Wang, K., Wei, Y.-M., Huang, Z., 2024. Transitioning from gasoline to electric vehicles: electrification decision of automakers under purchase and station subsidies. Transp. Res. E: Logist. Transp. Rev. 188, 103640, 1–24. https://doi.org/10.1016/j.tre.2024.103640
Rahiminia, S., Mehrabi, A., Pourseyed Aghaee, M., Jamili, A., 2023. Adopting a bilevel optimization method for the freight transportation problem: a multi-objective programming approach. Transp. Res. Rec. 2677 (2), 490–504. https://doi.org/10.1177/03611981221107627
Rapson, D.S., Muehlegger, E., 2023. The economics of electric vehicles. Rev. Environ. Econ. Policy 17 (2), 274–294. https://doi.org/10.1086/725484
Rijkswaterstaat D., 2025. [Data set]. Retrieved https://downloads.rijkswaterstaatdata.nl/scheepvaart/goederenvervoer/.
Shao, L., Yang, J., Zhang, M., 2017. Subsidy scheme or price discount scheme? mass adoption of electric vehicles under different market structures. Eur. J. Oper. Res. 262, 1181–1195. https://doi.org/10.1016/j.ejor.2017.04.030
Shi, L., Sethi, S.P., Çakanyıldı rım, M., 2022. Promoting electric vehicles: reducing charging inconvenience and price via station and consumer subsidies. Prod. Oper. Manag. 31, 4333–4350. https://doi.org/10.1111/poms.13871
Springel, K., 2021. Network externality and subsidy structure in two-sided markets: evidence from electric vehicle incentives. Am. Econ. J.: Econ. Policy 13, 393–432. https://doi.org/10.1257/pol.20190131
Srivastava, A., Kumar, R.R., Chakraborty, A., Mateen, A., Narayanamurthy, G., 2022. Design and selection of government policies for electric vehicles adoption: a global perspective. Transp. Res. E: Logist. Transp. Rev. 161, 102726. https://doi.org/10.1016/j.tre.2022.102726
Sun, L., Zhang, Y., Ma, F., Ji, F., Xiong, Y., 2023. Energy and speed optimization of inland battery-powered ship considering dynamic electricity price and complex navigational environment. Energy Rep. 9, 293–304.
TNO, 2009. Methodologies for estimating shipping emissions in the Netherlands. TNO Report. TNO. https://www.tno.nl/media/2151/methodologies_for_estimating_shipping_emissions_netherlands.pdf.
Villa, D., Montoya, A., Herrera, A.M., 2020. The electric riverboat charging station location problem. J. Adv. Transp. 2020, 6527924. https://doi.org/10.1155/2020/6527924
Villamar, D., Aussel, D., 2024. A bilevel optimization approach of energy transition in freight transport: SOS1 method and application to the ecuadorian case. Comput. Manag. Sci. 21, 39. https://doi.org/10.1007/s10287-024-00520-3
Wang, S., Qi, J., Laporte, G., 2022. Governmental subsidy plan modeling and optimization for liquefied natural gas as fuel for maritime transportation. Transp. Res. B 155. https://doi.org/10.1016/j.trb.2021.11.003
Wang, Y., Fan, R., Lin, J., Xie, X., Zhang, W., Srinivasan, D., 2025. Subsidies for shared private electric vehicle chargers: a three-level stackelberg game analysis. Transp. Res. D: Transp. Environ. 140, 104626. https://doi.org/10.1016/j.trd.2025.104626
Yang, S., Ma, F., Sun, H., Sun, L., 2024. Optimal voyage scheduling of battery-swapping based inland all-electric ship fleet considering environmental factors and time-of-use electricity price. In: 2024 9th Asia Conference on Power and Electrical Engineering (ACPEE). https://doi.org/10.1109/ACPEE60788.2024.10532657

Zero Emission Services, 2021a. Zero emission services commences operation: first emission-free inland shipping vessel on energy containers in service. Press release. Accessed 18 August 2024.
Zero Emission Services, 2021b. ZESpack : battery and technology. https://zeroemissionsservices.nl/en/zespack-2/.
Zero Emission Services, 2025. ITG, ZES and Nedcargo drive emission-free inland shipping. https://zeroemissionservices.nl/en/itg-zes-and-nedcargo-drive-emission-free-inland-shipping/.
Zhang, J., Li, X., Jia, D., Zhou, Y., 2023. A bilevel programming for union battery swapping stations location-routing problem under joint distribution and cost allocation. Energy 272, 127152. https://doi.org/10.1016/j.energy.2023.127152
Zhang, Y., Sun, L., Fan, T., Ma, F., Xiong, Y., 2023. Speed and energy optimization method for the inland all-electric ship in battery-swapping mode. Ocean Eng. 284, 115234. https://doi.org/10.1016/j.oceaneng.2023.115234
Zhang, Y., Sun, L., Ma, F., Wu, Y., Jiang, W., Fu, L., 2022. Collaborative optimization of the battery capacity and sailing speed considering multiple operation factors for a battery-powered ship. World Electr. Veh. J. 13 (2), 40. https://doi.org/10.3390/wevj13020040
Zhang, Y., Sun, L., Sun, W., Ma, F., Xiao, R., Wu, Y., Huang, H., 2024. Bilevel optimal infrastructure planning method for the inland battery swapping stations and battery-powered ships. Tsinghua Sci. Technol. 29 (5). https://doi.org/10.26599/TST.2023.9010138