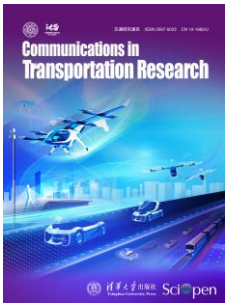
COMMTR

Communications in Transportation Research



RESEARCH ARTICLE

Emission-aware operations of flex-route transit vehicles: An exact logic-based Benders decomposition algorithm

Shuiwang Chen[1], Mingyang Li[2], Lingxiao WU[1,3]✉

1 Department of Aeronautical and Aviation Engineering, The Hong Kong Polytechnic University, Hong Kong 999077, China.

2 Department of Logistics and Maritime Studies, The Hong Kong Polytechnic University, Hong Kong 999077, China.

3 The Hong Kong Polytechnic University Shenzhen Research Institute, Shenzhen 518057, China.

✉ Corresponding author.

E-mail: ingxiao-leo.wu@polyu.edu.hk

Just Accepted

DOI 10.26599/COMMTR.2026.9640076

Received 6 September 2026

Revised 22 October 2026

Accepted 25 November 2026

Homepage

Cite this article: https://doi.org/10.26599/COMMTR.2026.9640076

© The Author(s) 2026.

This is an open access article under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0, http://creativecommons.org/licenses/by/4.0/), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

Published by Tsinghua University Press

https://www.sciopen.com/journal/2097-5023

Commun Transp Res

Research Article

Emission-aware operations of flex-route transit vehicles: An exact logic-based Benders decomposition algorithm

Shuiwang Chen[1], Mingyang Li[2], Lingxiao WU[1,3]✉

Cite this article: XXX

ABSTRACT: Flex-route transit combines a scheduled checkpoint backbone with demand-responsive pickup and delivery. We study an emission-aware operational planning problem in which vehicle assignment, passenger sequencing, capacity, time windows, and arc speeds must be coordinated. We formulate the routing layer on a checkpoint-expanded network and develop an exact logic-based Benders decomposition (LBBD) that separates route selection from independent vehicle-level speed subproblems. The baseline LBBD uses no-good cuts for speed-infeasible routes and optimality cuts for feasible routes. An enhanced LBBD adds piecewise-linear speed–emission bounds, load-dependent emission bounds, and reusable vehicle-level cuts to strengthen the routing master before and after subproblem evaluations. The speed subproblems are solved exactly as convex second-order cone programs. Across small instances, enhanced LBBD solves all cases to optimality and, on the more demanding small cases, is about 490 times faster than mixed-integer second order cone programming (MISOCP) and 270 times faster than baseline LBBD. On large instances, it reduces the mean optimality gap to 0.57%, compared with 6.47% for MISOCP and 9.26% for baseline LBBD. In practice, the proposed model can reduce emissions by up to 30% relative to the traditional distance-based routing model.

KEYWORDS: Flex-route transit; Emission reduction; Second-order conic programming; Logic-based Benders decomposition

Just Accepted

1 Introduction

Flex-route transit (FRT) is a semi-flexible public transportation service. It is intended for scenarios in which a fixed route is inefficient for dispersed demand, but a fully demand-responsive service may reduce the reliability and passenger consolidation advantages of scheduled transit (Shahin et al., 2024; Zheng et al., 2018). A FRT line follows a predetermined sequence of checkpoints. These checkpoints anchor the vehicle schedule of operations. Between consecutive checkpoints, the vehicle can leave the reference route to serve passenger pickup and dropoff locations. This mode integrates a predictable service schedules with the flexible operations to respond demand (Quadrifoglio et al., 2008, 2007).

FRT is often viewed as a sustainable alternative to private cars. Its actual environmental impact depends on how the service is designed and operated (Banister, 2008; Miller et al., 2016). In low-density areas, a flexible vehicle may travel many miles without passengers. It may also make detours to serve only a few passengers or require additional vehicles to maintain the checkpoint schedule. Coordinated passenger assignments and routes can help consolidate requests, reduce unnecessary empty travel, and use each vehicle more efficiently. Environmental performance depends on more than vehicle technology and route length. Dispatchers also make decisions before and during each trip. They decide which requests to service, which detours to allow, and what speeds vehicles should travel. FRT operators therefore need to consider emissions with service feasibility and coverage.

FRT operating decisions directly affect vehicle emissions because service constraints govern how vehicles travel. Passenger assignments determine the order of pickup and delivery stops, and any resulting detours increase travel distance and emissions. Tight checkpoint or passenger time windows force vehicles to travel faster to stay on schedule, which affects both travel time and emissions. Meanwhile, passenger boarding and alighting also change the vehicle's load and its load-related emissions. Road grade is another factor. The same load can produce different emissions on uphill and level roads. These relationships motivate routing models that account for energy use and emissions (Bektaş and Laporte, 2011; Demir et al., 2012). FRT operators must consider these factors with pickup-and-delivery precedence, vehicle capacity, checkpoint order, and hard time windows. The environmental operations of FRT must therefore account for its schedule, load profile, and speed on each road.

This motivates our research question about how a FRT operator to choose routes and control speeds when emissions are a primary objective. The operator must preserve the checkpoint backbone schedules and accommodate passenger pickup and delivery while selecting speeds that satisfy all time windows. Because routing choices dictate feasible speed schedules and speeds alter the overall cost of a route, decoupling these decisions fails to capture the actual operating trade-offs. We review research on flexible transit and speed optimization with emissions as an objective.

1.1. Literature review

FRT models combine scheduled checkpoints with demand-responsive deviations to service passengers (Mageean and Nelson, 2003; Qiu et al., 2015, 2014; Quadrifoglio et al., 2008, 2007; Shahin et al., 2024; Vansteenwegen et al., 2022; Zheng et al., 2019, 2018). Pickup-and-delivery and dial-a-ride studies provide the basic formulation of the

[1] Department of Aeronautical and Aviation Engineering, The Hong Kong Polytechnic University, Hong Kong, 999077, China. [2] Department of Logistics and Maritime Studies, The Hong Kong Polytechnic University, 999077, Hong Kong, China. [3] The Hong Kong Polytechnic University Shenzhen Research Institute, Shenzhen, 518057, China.

✉ Corresponding author. E-mail: lingxiao-leo.wu@polyu.edu.hk

Received: xxx xx, 20xx; Revised: xxx xx, 20xx; Accepted: xxx xx, 20xx

© The Author(s) 20xx. This is an open access article under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0, http://creativecommons.org/licenses/by/4.0/).

FRT routing problem (Chen et al., 2021; Molenbruch et al., 2023; Sartori Vieira et al., 2025). Recent FRT studies have also considered energy consumption (Li et al., 2025; Tang et al., 2026). However, most FRT formulations assume fixed speeds or fixed travel times. As a result, they do not capture the interaction among routing and emissions.

Emission-aware routing studies often link route decisions with speed choices. The pollution-routing literature jointly models routes, loads, speeds, and travel times. It shows that the shortest route may not minimize emissions (Bektaş and Laporte, 2011; Dabia et al., 2017; Fukasawa et al., 2018). Research on eco-driving routing also shows that speed affects both schedule feasibility and emissions (Lai et al., 2024; Ma et al., 2021; Miao et al., 2018; Wu et al., 2025). Other transportation studies also study this interactions by integrating route choices and vehicle speeds or arrival times (Akturk et al., 2025; Chen et al., 2024; Wang, 2016; Wang and Meng, 2012; Zhen et al., 2024; Zhuge et al., 2024, 2026). Aviation and maritime research account for factors such as cargo weight, weather, and schedule disruptions (Akturk et al., 2025; Wang and Meng, 2012; Zhen et al., 2024; Zhuge et al., 2026). Across these applications, both the route and the speed profile shape emissions.

Existing studies address these integrated problem through nonlinear optimization, dynamic programming, and branch-and-price algorithms (Dabia et al., 2017; Fukasawa et al., 2018, 2016). These approaches typically include route and speed variables in one integrated model. The others use linear approximations designed for specific speed-emission functions (Wang, 2016). Together, these studies show the value of coordinating discrete route choices with continuous speed decisions. However, they do not provide a decomposition framework for large-scale routing problems with load profile, hard time windows, and speed optimization.

1.2. Research gaps and contributions

Despite this progress, existing studies do not simultaneously address mandatory checkpoint topologies, paired passenger requests, passenger loads, hard time windows, endogenous arc-specific speeds, and load-dependent emissions. The literature also lacks an exact decomposition by vehicle that separates these routing and speed decisions while providing certified bounds and optimality gaps. To the best of our knowledge, no prior study develops a logic-based Benders decomposition for this FRT structure or decomposes routing and endogenous speed decisions within a vehicle-level exact framework. We address this gap by studying the joint routing and speed optimization problem for FRT operations with emissions as the objective. Specifically, this paper makes the following three main contributions:

Our first contribution is a modeling framework for environmentally responsible FRT operations (Section 3). We formulate a checkpoint-expanded pickup-and-delivery model that jointly determines request assignments, paired endpoint sequences, checkpoint-compatible routes, passenger loads, time windows, and arc speeds. The objective function minimizes emissions directly through a physical speed, load, and grade relationship. This formulation captures how service flexibility and schedule protection generate environmental costs. The checkpoint expansion maintains the original service structure and removes incompatible transitions prior to optimization, which creates a compact routing representation for the decomposition method.

Our second contribution is an exact logic-based Benders decomposition (LBBD) algorithm (Section 4). This method decomposes the problem into a discrete routing master and vehicle-level speed optimization subproblems. We strengthen the routing master problem and reformulate each nonlinear speed subproblem as an equivalent second-order cone program (SOCP). We derive no-good cuts for infeasible route patterns and full-pattern cost cuts for feasible patterns. We also develop an enhanced version that incorporates piecewise-linear speed–emission cuts, load-dependent emission bounds, and reusable vehicle-level cuts to tighten the routing master. We prove that both methods are exact and converge finitely.

Our third contribution is the numerical experiment and the operational implications (Section 5). We compare the baseline and enhanced LBBD methods with an exact mixed-integer second-order cone programming (MISOCP) method. The results show that the enhanced LBBD improves lower bound and improves optimality. We also conduct experiments to examine the environmental impacts of speed control, joint route and speed decisions, vehicle capacity, road grade, time-window width, and aerodynamic parameters. These results provide operational insights for FRT operators.

2 Problem description

2.1. Problem background

FRT operates between conventional fixed-route transit and fully demand-responsive service. Fig. 1 shows that a FRT line contains an ordered sequence of checkpoints. Every FRT vehicle must visit these checkpoints in the prescribed order and schedule. Between two consecutive checkpoints, the vehicle can depart from the reference line to serve passenger pickup and delivery locations. Hence, the checkpoint sequence remains fixed, whereas the passenger stops visited between checkpoints and their service order should be determined by the operator according to the demand.

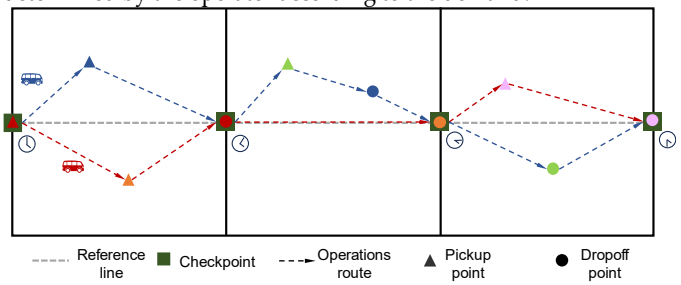


Fig.1. Illustration of FRT operations.

This operating mode combines schedule regularity with spatial flexibility. The checkpoints preserve the basic timetable of the transit line and provide common reference locations for the fleet. By contrast, passenger demand points require demand-responsive deviations from the reference line. A request may be served entirely between two consecutive checkpoints. In addition, passenger may remain onboard while the vehicle passes checkpoints. Therefore, checkpoints restrict the movement of vehicles but do not necessarily divide a passenger trip into independent segments. A feasible operating plan must coordinate all deviations with the remaining checkpoint schedule.

We represent the physical transportation system by a directed graph $G = (\mathcal{V}, \mathcal{A})$. Its node set contains the ordered checkpoints $\mathcal{C} = \{c_o, c_1, \cdots, c_H\}$ the pickup-node set $\mathcal{P}$, and the delivery-node set $\mathcal{D}$. The request set is denoted by $\mathcal{R}$. Each request $r \in \mathcal{R}$ is represented by a paired pickup and delivery $(p_r, d_r) \in \mathcal{P} \times \mathcal{D}$, and q_r denotes its party size. The service is operated by a homogeneous fleet $\mathcal{K} = \{1, \cdots, K\}$,

Just Accepted

where every vehicle has capacity Q, curb mass w, an origin o_k, and a destination z_k

For each vehicle k, an operated route is an ordered node sequence $R_k = \{i_0^k, i_1^k, \cdots, i_{n_k}^k\}$ starting from its origin and ending at its destination. The subsequence of checkpoint nodes in R_k must be exactly $(c_0, c_1, \cdots, c_H)$. Passenger nodes may be inserted between these checkpoints, subject to their segment compatibility. If vehicle k serves request r, p_r and d_r must occur in R_k, and p_r must precede d_r. The same request cannot be divided between different vehicles.

The passenger sequence determines the load carried along the route. Let b_i denote the load change at node i, where $b_{p_r} = q_r$ and $b_{d_r} = -q_r$. If l_i is the onboard load after serving node i, the load on a selected transition from i to j satisfies $l_j = l_i + b_j$. The vehicle starts and terminates empty, and its load must satisfy $0 \le l_i \le Q$ at every visited node. Consequently, the ordering of passenger endpoints not only determines route feasibility but also determines the passenger mass transported on each arc.

The route must also admit a feasible schedule. Each node $i \in \mathcal{V}$ has a service duration s_i and a time window $[t_i^e, t_i^l]$. Let t_i denote the departure time from node i. Early arrival is allowed and results in waiting, whereas service cannot depart later than t_i^l. Checkpoint time windows maintain the scheduled backbone, while the windows of non-checkpoint endpoints restrict the allowable timing of passenger deviations. Thus, the feasibility of an insertion depends not only on its local travel time but also on the remaining time available before all downstream checkpoints and passenger stops.

For every physical arc $(i, j) \in \mathcal{A}$, let d_{ij} denote its three-dimensional travel distance and θ_{ij} its road-grade angle. The operating speed on a selected positive-distance arc is a continuous decision $v_{ij} \in [\underline{v}, \overline{v}]$, and the corresponding travel time is d_{ij}/v_{ij}. Increasing speed may preserve the feasibility of a tight downstream time window, whereas it changes the emissions generated on that arc. Conversely, a lower speed may be environmentally preferable although it can consume the schedule slack required to accommodate passenger deviations. Route construction and speed selection therefore cannot be treated as independent operating decisions.

For a selected arc (i, j), let L be the onboard passenger load and v the operating speed. Following the comprehensive modal-energy model of (Barth and Boriboonsomsin, 2008), the energy per unit distance is

$$f(L, v; \theta_{ij}) = \frac{\beta_1}{v_{ij}} + g(w + ml_i)\left(\sin\theta_{ij} + \beta_2 \cos\theta_{ij}\right) + \frac{1}{2}\beta_3 v_{ij}^2,$$

where m is the passenger mass, g is gravitational acceleration, β_1 and β_2 are nonnegative coefficients, and β_3 is the quadratic speed coefficient. The energy used on the arc is therefore

$$F_{ij}^{\mathrm{E}} = d_{ij} f(L, v; \theta_{ij}).$$

For the electric vehicles commonly used in FRT operations, indirect emissions from energy use are calculated as

$$E = e \sum_{(i,j)\in\mathcal{A}} F_{ij}^{\mathrm{E}},$$

with the energy–emission factor

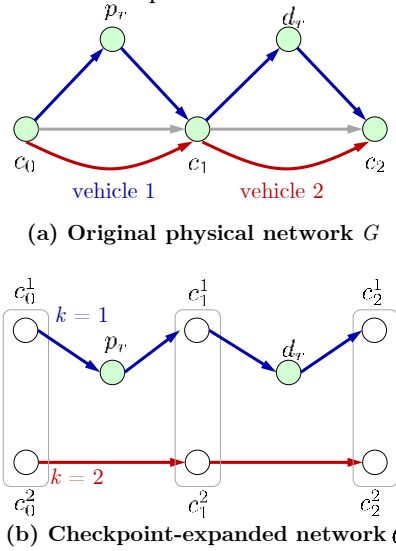
$$e = \frac{EF_{grid}}{3.6\,\eta_{btw}\eta_{chg}},$$

where EF_{grid} is the grid emission factor, 3.6 converts MJ to kWh, η_{btw} is battery-to-wheel efficiency, and η_{chg} is charging efficiency. Because this conversion is proportional to tractive energy, the objective also provides a direct proxy for electricity use and its variable cost for operators. Labor, fixed vehicle costs, and other nonenergy expenses are outside the present scope. Passenger service is protected by the hard pickup and delivery windows, so the emission objective is optimized only over plans that meet the promised service level. Arc emissions depend on road grade, operating speed, and the load determined by the preceding pickup and delivery sequence, as well as on distance. The problem is consequently to assign all requests to vehicles, construct checkpoint-compatible routes, schedule all visited nodes, propagate vehicle loads, and determine continuous arc speeds so as to minimize total emissions while satisfying pairing, precedence, capacity, checkpoint-order, and time-window requirements.

2.2. Checkpoint-expanded network

To streamline the model formulation, we introduce a checkpoint-expanded network. The expanded network $\tilde{G} = (\tilde{\mathcal{N}}, \tilde{\mathcal{A}})$ duplicates every mandatory checkpoint c_h for each vehicle, forming a vehicle-specific chain $(c_0^k, \ldots, c_H^k)$ while keeping passenger nodes shared. Each checkpoint chain anchors a particular vehicle k, and a mapping variable κ_i propagates this identity to the passenger nodes it serves. As illustrated in Fig. 2, this representation is equivalent to the original physical network. However, it allows the model to replace the traditional three-index routing variables x_{ij}^k with more compact two-index variables x_{ij}, exclude checkpoint-incompatible arcs in advance, and express model much more compactly.

The expanded representation has node complexity $O(N + KH)$ because the $2N$ passenger nodes are shared and the $H + 1$ checkpoint copies are created for each available vehicle. Candidate arcs have a quadratic upper bound before filtering, while checkpoint order, vehicle identity, and segment compatibility remove most incompatible arcs in advance. The resulting two-index structure is more compact than a vehicle indexed three-index formulation and supports the separation of the routing master from vehicle level speed subproblems (Li et al., 2025). However, the main limitations of this representation are its computational dependence on fleet size and checkpoint density, and the homogeneous fleet assumption.



(a) Original physical network G

(b) Checkpoint-expanded network $\tilde{G}$

Fig. 2. Checkpoint-expanded network.

By design, each vehicle k traces an ordered checkpoint chain $\tilde{C}_k$ bounded by origin o_k and destination z_k,whereas all passenger nodes are formulated as shared entities. Defining $\mathcal{O} = \{o_k : k \in K\}$ and $\mathcal{Z} = \{z_k : k \in K\}$ as the origin and destination sets, respectively, we denote the shared passenger nodes by $\tilde{U}$. The network topology is

Just Accepted

mapped via the outgoing and incoming arc sets $\mathrm{out}(i)$ and $\mathrm{in}(i)$ for any generic node i. Each component in the expanded network retains its original time windows, demands, service durations, distances, and energy coefficients. We denote its positive-distance arc set by $\tilde{\mathcal{A}}^+ = \{(i,j) \in \tilde{\mathcal{A}}: \tilde{\delta}_{ij} > 0\}$.

2.3. Decisions

To formulate the model, we rely on binary variables as follows:

$$x_{ij} = \begin{cases} 1, & \text{if arc } (i,j) \in \tilde{\mathcal{A}}^+ \text{ is used,} \\ 0, & \text{otherwise,} \end{cases}$$
$$y_i = \begin{cases} 1, & \text{if node } i \text{ is used,} \\ 0, & \text{otherwise.} \end{cases}$$

The integer variable $\kappa_i \in \{0, \dots, K\}$ records the route assignment:

$$\kappa_i = \begin{cases} k, & \text{if vehicle } k \text{ is assigned to node } i, \\ 0, & \text{otherwise.} \end{cases}$$

Additionally, nonnegative continuous variables t_i and l_i record the departure time and payload after service at node i.

3 Optimization Model

In this section, we present the emission-aware model for FRT operations. We formulate the model as follows:

$$\mathbf{P}: \min e \sum_{(i,j)\in\mathcal{A}^+} d_{ij}\left[\frac{\beta_1}{v_{ij}} + g(w + ml_i)\left(\sin\theta_{ij} + \beta_2 \cos\theta_{ij}\right) + \frac{1}{2}\beta_3 v_{ij}^2\right] x_{ij} \quad (1)$$

$$\text{s.t. } y_c = y_{o_k}, \quad \forall k \in \mathcal{K}, c \in \tilde{\mathcal{C}}_k, \quad (2)$$

$$\kappa_c = k y_{o_k}, \quad \forall k \in \mathcal{K}, c \in \tilde{\mathcal{C}}_k, \quad (3)$$

$$y_{o_k} \geq y_{o_{k+1}}, \quad k = 1,2,\dots,K-1, \quad (4)$$

$$y_i = 1, \quad \forall i \in \tilde{\mathcal{U}}, \quad (5)$$

$$y_i \leq \kappa_i \leq K y_i \quad \forall i \in \tilde{\mathcal{V}}, \quad (6)$$

$$\sum_{(i,j)\in\mathrm{out}(i)} x_{ij} = y_i, \quad \forall i \in \tilde{\mathcal{V}} \backslash \mathcal{Z}, \quad (7)$$

$$\sum_{(j,i)\in\mathrm{in}(i)} x_{ji} = y_i, \quad \forall i \in \tilde{\mathcal{V}} \backslash \mathcal{O}, \quad (8)$$

$$\kappa_{p_r} = \kappa_{d_r}, \quad \forall r \in \mathcal{R}, \quad (9)$$

$$t_{p_r} \leq t_{d_r}, \quad \forall r \in \mathcal{R}, \quad (10)$$

$$\kappa_i - \kappa_j \leq K(1 - x_{ij}), \quad \forall (i,j) \in \tilde{\mathcal{A}}, \quad (11)$$

$$\kappa_j - \kappa_i \leq K(1 - x_{ij}), \quad \forall (i,j) \in \tilde{\mathcal{A}}, \quad (12)$$

$$l_{o_k} = 0, l_{z_k} = 0, \quad \forall k \in \mathcal{K}, \quad (13)$$

$$0 \leq l_i \leq Q y_i, \quad \forall i \in \tilde{\mathcal{V}}, \quad (14)$$

$$l_j \geq l_i + b_i x_{ij} - Q(1 - x_{ij}), \quad \forall (i,j) \in \tilde{\mathcal{A}}, \quad (15)$$

$$l_j \leq l_i + b_i x_{ij} + Q(1 - x_{ij}), \quad \forall (i,j) \in \tilde{\mathcal{A}}, \quad (16)$$

$$t_i^e y_i \leq t_i \leq t_i^l y_i, \quad \forall i \in \tilde{\mathcal{V}}, \quad (17)$$

$$t_j \geq t_i + \frac{d_{ij}}{v_{ij}} x_{ij} + s_j x_{ij} - M_{ij}^t (1 - x_{ij}), \quad \forall (i,j) \in \tilde{\mathcal{A}}, \quad (18)$$

$$\underline{v} \leq v_{ij} \leq \overline{v}, \quad \forall (i,j) \in \tilde{\mathcal{A}}, \quad (19)$$

$$x_{ij} \in \{0,1\}, y_i \in \{0,1\}, \kappa_i \in \{0,1,\dots,K\}, t_i \geq 0, l_i \geq 0. \quad (20)$$

The objective function (1) minimizes total emissions over all selected positive-distance arcs. Constraints (2) activate the checkpoint copy for vehicle k whenever that vehicle is used. Constraints (3) assign the identity of vehicle k to every checkpoint on its chain. Constraints (4) order vehicle activation and remove equivalent solutions caused by identical vehicles. Constraints (5) require every shared pickup and delivery node to be visited. Constraints (6) set the lower and upper bound on the identity variable at used nodes. Constraints (7) require one selected outgoing arc from every used node except a destination. Constraints (8) require one selected incoming arc to every used node except an origin. Constraints (9) assign the pickup and delivery of request r to the same vehicle. Constraints (10) impose pickup-before-delivery precedence. Constraints (11) limit the increase in vehicle identity along a selected arc. Constraints (12) limit the decrease in vehicle identity along a selected arc. Together, these constraints propagate a constant vehicle identity along each selected route.

Constraints (13) set the load at every vehicle origin and destination to zero. Constraints (14) impose nonnegative load and limit load to vehicle capacity. Constraints (15) and (16) give the lower and upper load propagation relation on a selected arc. Constraints (17) impose the earliest and latest allowable time at a used node. Constraints (18) propagate departure time from node i to node j, including travel time, service time, and deactivation when the arc is not selected. Constraints (19) impose the minimum and maximum arc speed. Constraints (20) define the domains of all decision variables.

Although our model does not explicitly include a vehicle operating cost in the objective function, the operated fleet is endogenous through the activation variables $\boldsymbol{\kappa}$. A vehicle contributes curb mass emissions $d_{ij} g w \left(\sin\theta_{ij} + \beta_2 \cos\theta_{ij}\right)$ on every selected positive distance arc $(i,j) \in \mathcal{A}^+$ once activated ($x_{ij} = 1$). Therefore, increasing K enlarges the feasible set without forcing additional vehicles into a solution. The formulation therefore avoids an uncalibrated fixed activation penalty while retaining the physical emission effect of vehicle usage.

The departure time t_i is defined after service at node i, and the payload l_i is the load after service. Therefore, a selected arc(i,j) carries the departing load l_i, while unused nodes have zero time and load through the activation constraints. Model $\mathbf{P}$ is a mixed integer nonlinear program. A selected route fixes the node sequence, arc loads, and time window constraints on the admissible speeds. The speed choices then determine travel times and route emissions. Reciprocal speed and squared speed terms, along with products involving load, arc selection, and route activation, couple these discrete and continuous decisions. This structure makes direct optimization harder as the network and request set grow.

4 Logic-based Benders decomposition

Section 4.1 defines the routing master and speed subproblems. Sections 4.2 presents the baseline algorithm. Section 4.3 introduces enhanced valid inequalities. Section 4.4 gives the complete algorithm.

4.1. Model decomposition

4.1.1 Master problem

We decouple the speed decision and time-window constraints from the model $\mathbf{P}$, yielding a master problem $\mathbf{P}_\mathbf{M}^{(n)}$ at iteration n as follows:

$$\mathbf{P}_\mathbf{M}^{(n)}: \min \Theta$$
$$(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{\kappa}, \boldsymbol{l}) \in \mathcal{X},$$
$$(\Theta, \boldsymbol{x}) \in \mathcal{B}^{(n)},$$

where Θ estimates total emissions, $\mathcal{X} \coloneqq \{(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{\kappa}, \boldsymbol{l}) | (2) - (16), (20)\}$ denote the feasible space of routing and load decisions, and $\mathcal{B}^{(n)}$

Just Accepted

contains cuts generated in earlier iterations. The resulting problem can be viewed as a constraint programming model with a constant objective function. Nevertheless, this relaxation is quite weak, since it only identifies feasible routes and does not aim to minimize emissions. To improve the formulation, we develop a more tightened master problem based on Proposition 1.

Proposition 1. *For a positive-distance arc (i,j) and a fixed onboard load L, if time-window constraints are ignored, the speed-dependent part of the emission function is minimized at*

$$v^{\text{eco}} = \left(\frac{\beta_1}{\beta_3}\right)^{\frac{1}{3}}, \quad v^\star = \min\{\overline{v}, \max\{\underline{v}, v^{\text{eco}}\}\}.$$

The corresponding arc value

$$\underline{\epsilon}_{ij}(L) = e\, d_{ij}\left[\frac{\beta_1}{v_{ij}^\star} + g(w+mL)(sin\theta_{ij} + \beta_2 cos\theta_{ij}) + \frac{1}{2}\beta_3 (v_{ij}^\star)^2\right]$$

is a valid lower bound on the emission of that arc for every feasible speed.

Proof. For fixed L and (i,j), the only speed-dependent terms are $\phi(v) = \beta_1/v + \frac{1}{2}\beta_3 v^2$. Since $\phi''(v) = 2\frac{\beta_1}{v^3} + \beta_3 > 0$, $\phi(v)$ is strictly convex. Solving $\phi'(v) = -\beta_1/v^2 + \beta_3 v = 0$ gives v^{eco}. Strict convexity implies that the minimizer on the bounded interval $[\underline{v}, \overline{v}]$ is its projection $v^\star$. The remaining grade and mass term is independent of speed, so substituting $v^\star$ gives the stated lower bound. ■

In particular, if $v^{\text{eco}} \le \underline{v}$ under the adopted parameter setting, then $v^\star = \underline{v}$ which gives the minimum-speed result as a special case. To obtain a coefficient independent of the unknown master load, we use

$$\underline{\epsilon}_{ij} = \min_{0 \le L \le Q} \underline{\epsilon}_{ij}(L).$$

Proposition 2. *For every selected positive-distance arc (i,j), any feasible speed satisfies*

$$t_j \ge t_i + \frac{d_{ij}}{\overline{v}} + s_j.$$

With binary arc activation, a valid linear form is

$$t_j \ge t_i + \frac{d_{ij}}{\overline{v}} x_{ij} + s_j x_{ij} - M_{ij}^t\left(1 - x_{ij}\right), \quad \forall (i,j) \in \tilde{\mathcal{A}}.$$

Proof. For a selected arc, $v_{ij} \le \overline{v}$ implies $\frac{d_{ij}}{v_{ij}} \ge \frac{d_{ij}}{\overline{v}}$. Adding the service duration at node j gives $t_j \ge t_i + \frac{d_{ij}}{\overline{v}} + s_j$. When $x_{ij} = 0$, the sufficiently large constant $M_{ij}^t := t_i^l + \frac{d_{ij}}{\underline{v}} + s_j$ deactivates the implication, which yields the necessary time-window. ■

Build upon Propositions 1 and 2, we can initially strength the master problem $\mathbf{P}_{\mathbf{M}}^{(n)}$ as:

$$\begin{aligned}
\hat{\mathbf{P}}_{\mathbf{M}}^{(n)}:\ & \min \Theta \\
& (\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{\kappa}, \boldsymbol{l}) \in \mathcal{X}, \\
& (\Theta, \boldsymbol{x}) \in \mathcal{B}^{(n)}, \\
& \Theta \ge \sum\nolimits_{(i,j)\in\tilde{\mathcal{A}}^+} \underline{\epsilon}_{ij} x_{ij}\ , \\
& t_j \ge t_i + \frac{d_{ij}}{\overline{v}} x_{ij} + s_j x_{ij} - M_{ij}^t\left(1 - x_{ij}\right), \quad \forall (i,j) \in \widetilde{\mathcal{A}}, \\
& t_i \le t_i^l y_i, \quad \forall i \in \tilde{\mathcal{V}}.
\end{aligned}$$

4.1.2 Subproblem

Once the master solution $\hat{\boldsymbol{x}}$ is fixed, the expanded network identifies one ordered route for each active vehicle. The load recursion then fixes the onboard load on every traversed arc. The speed and departure time variables of one vehicle do not appear in another vehicle's constraints, and emissions are additive across vehicles. Each active vehicle can therefore be evaluated separately, with the complete route emission given by the sum of the vehicle level optima. This also allows the vehicle subproblems to be solved in parallel.

For a master solution $\hat{x}$, let vehicle k have fixed route $R_k(\hat{x}) = (i_0^k, \dots, i_{n_k}^k)$ and arc index set $\mathcal{J}_k = \{0, \cdots, n_k - 1\}$. The route fixes its load recursively as $\Lambda_0^k = 0$ and $\Lambda_{n+1}^k = \Lambda_n^k + b_{i_{n+1}^k}$. For compactness, write $d_n^k = d_{i_n^k i_{n+1}^k}$ and $\theta_\pi^k = \theta_{i_n^k i_{n+1}^k}$, and let $\mathcal{J}_k^+ = \{n \in \mathcal{J}_k : d_n^k > 0\}$ denote its positive-distance arcs. The zero-distance arcs require no speed decision. The speed optimization subproblem is formulated as:

$$\begin{aligned}
\mathbf{P}_{\mathbf{S}}^k:\ & \min e \sum_{n\in\mathcal{J}_k^+} d_\pi^k \left[\frac{\beta_1}{v_n^k} + g(w + m\Lambda_n^k)(\sin\theta_n^k + \beta_2 \cos\theta_n^k) + \frac{1}{2}\beta_3 (v_n^k)^2\right] \\
\text{s.t.}\ & t_{i_n^k}^e \le t_n^k \le t_{i_n^k}^l, \quad \forall\, n \in \{0, \dots, n_k\}, \\
& t_{n+1}^k \ge t_n^k + \frac{d_n^k}{v_n^k} + s_{i_{n+1}^k}, \quad \forall n \in J_k^+, \\
& t_{n+1}^k \ge t_n^k + s_{i_{n+1}^k}, \quad \forall n \in J_k \setminus J_k^+, \\
& \underline{v} \le v_n^k \le \overline{v}, \quad \forall n \in J_k^+,
\end{aligned}$$

where the first constraint enforces the time window at every visited node, the next two constraints propagate departure times along positive- and zero-distance arcs, respectively, and the last constraint restricts every operating speed to its admissible interval. Although the route and loads are now parameters, reciprocal speed remains in both the objective and the time constraints, and squared speed remains in the objective. We next show that these two nonlinear expressions have exact rotated second-order-cone representations.

Introduce nonnegative epigraph variables ρ_n^k and η_n^k for $\frac{1}{v_n^k}$ and $(v_n^k)^2$, respectively. We employ the following two Propositions to reformulate the model.

Proposition 3. *For $v_n^k > 0$ and $\rho_n^k \ge 0$, the rotated cone constraint*

$$\left\| \left(2, \frac{v_n^k}{\overline{v}} - \overline{v}\rho_n^k\right) \right\|_2 \le \frac{v_n^k}{\overline{v}} + \overline{v}\rho_n^k \tag{21}$$

is equivalent to $v_n^k \rho_n^k \ge 1$. Because ρ_n^k has a nonnegative objective coefficient and enters the time propagation constraint with a nonnegative coefficient, the conic problem admits an optimal solution satisfying $\rho_n^k = 1/v_n^k$.

Proof. Squaring both sides of (21) gives

$$4 + \left(\frac{v_n^k}{\overline{v}} - \overline{v}\rho_n^k\right)^2 \le \left(\frac{v_n^k}{\overline{v}} + \overline{v}\rho_n^k\right)^2,$$

which simplifies to $4 \le 4 v_n^k \rho_n^k$. Thus, the cone is equivalent to $v_n^k \rho_n^k \ge 1$. For fixed v_n^k, replacing ρ_n^k by $1/v_n^k$ preserves the cone, relaxes the corresponding time constraint, and does not increase the objective. Hence an optimal solution with equality always exists. ■

Proposition 4. *For $v_n^k > 0$ and $\eta_n^k \ge 0$, the rotated cone constraint*

$$\left\| \left(\frac{2v_n^k}{\overline{v}}, \frac{\eta_n^k}{\overline{v}^2} - 1\right) \right\|_2 \le \frac{\eta_n^k}{\overline{v}^2} + 1 \tag{22}$$

is equivalent to $\eta_n^k \ge (v_n^k)^2$. Since the coefficient of η_n^k in the objective is nonnegative, the conic problem admits an optimal solution satisfying $\eta_n^k = (v_n^k)^2$.

Proof. Squaring both sides of (22) gives

$$\frac{4(v_n^k)^2}{\overline{v}^2} + \left(\frac{\eta_n^k}{\overline{v}^2} - 1\right)^2 \le \left(\frac{\eta_n^k}{\overline{v}^2} + 1\right)^2,$$

which simplifies to $\frac{4(v_n^k)^2}{\overline{v}^2} \le \frac{4\eta_n^k}{\overline{v}^2}\overline{v}^2$ or equivalently $\eta_n^k \ge (v_n^k)^2$. For fixed v_n^k, replacing η_n^k by $(v_n^k)^2$ preserves feasibility and does not

Just Accepted

increase the objective. Hence an optimal solution with equality always exists. ■

Propositions 3 and 4 replace the two nonlinear epigraph relations without changing the optimal speed or objective value. Applying them to every positive-distance arc gives the following SOCP:

$$\mathbf{P}^k_{\mathbf{SOCP}}: \mathcal{G}_k = \min e \sum_{n\in J_k^+} d_v^k \left[\beta_1 \rho_n^k + g(w+\Lambda_n^k)(\sin\theta_n^k + \beta_2 \cos\theta_n^k) + \frac{1}{2}\beta_3 \eta_n^k\right]$$

$$\text{s.t.} \quad t_{i_n^k}^e \le t_n^k \le t_{i_n^k}^l, \quad \forall\, n \in \{0,\dots,n_k\}, \tag{23}$$

$$t_{n+1}^k \ge t_n^k + d_n^k \rho_n^k + s_{i_{n+1}^k}, \quad \forall n \in J_k^+, \tag{24}$$

$$t_{n+1}^k \ge t_n^k + s_{i_{n+1}^k}, \quad \forall n \in J_k \setminus J_k^+, \tag{25}$$

$$\underline{v} \le v_n^k \le \overline{v}, \quad \forall n \in J_k^+, \tag{26}$$

$$\frac{1}{\overline{v}} \le v_n^k \le \frac{1}{\underline{v}}, \quad \forall n \in J_k^+, \tag{27}$$

$$\underline{v}^2 \le \eta_n^k \le \overline{v}^2, \quad \forall n \in J_k^+, \tag{28}$$

$$\left\| \left(2, \frac{v_n^k}{\overline{v}} - \overline{v}\rho_n^k\right)\right\|_2 \le \frac{v_n^k}{\overline{v}} + \overline{v}\rho_n^k, \quad \forall n \in J_k^+, \tag{29}$$

$$\left\| \left(\frac{2v_n^k}{\overline{v}}, \frac{\eta_n^k}{\overline{v}^2} - 1\right)\right\|_2 \le \frac{\eta_n^k}{\overline{v}^2} + 1, \quad \forall n \in J_k^+, \tag{30}$$

$$t_n^k \ge 0,\ v_n^k \ge 0,\ \rho_n^k \ge 0,\ \eta_n^k \ge 0. \tag{31}$$

Since their objective coefficients are nonnegative, an optimal solution exists with both epigraphs tight..Constraints (23) – (25) retain the original time-window and sequencing requirements. Constraints (26) – (28) give bounds implied by the speed interval, while (29) – (30) impose the two epigraph relations. All remaining expressions are affine. Hence, problem $\mathbf{P}^k_{\mathbf{SOCP}}$ is a convex SOCP equivalent to $\mathbf{P}^k_{\mathbf{S}}$. It either certifies that the master route has no feasible speed schedule or returns its exact minimum emission E_k. The SOCP objective retains the complete emission expression, including the load and grade term. Constraint (29) enforces $v_n^k \rho_n^k \ge 1$ and constraint (30) enforces $\eta_n^k \ge (v_n^k)^2$. For comparison, embedding the SOCP representation of the speed-dependent terms in the full routing model and applying the exact convex-hull linearization to each product $l_i x_{ij}$ (see Equations (36)-(38)) yields an exact mixed-integer second-order cone program (MISOCP).

4.2. Baseline LBBD

The master problem deliberately retains only necessary timing conditions and lower estimates of emission. A route pattern selected by the master may therefore have no feasible assignment of speeds, and even a speed-feasible pattern may have a true emission well above the master estimate Θ. The baseline LBBD addresses these two outcomes with one feasibility cut and one optimality cut (Hooker and Ottosson, 2003). Both cuts are generated only after the selected pattern has been evaluated by the vehicle speed subproblems. Therefore, they require no dual information from those subproblems.

Suppose first that at least one vehicle subproblem is infeasible. Because the master pattern fixes the vehicle routes and their node sequences, no alternative speed choice can make this complete pattern feasible. Let $\Gamma(\hat{x}) = \{(i,j): \hat{x}_{ij} = 1\}$ denote the set of selected expanded arcs. The following no-good cut excludes the evaluated pattern:

$$\sum_{(i,j)\in\Gamma(\hat{x})} x_{ij} \le |\Gamma(\hat{x})| - 1. \tag{32}$$

The cut requires a future master solution to change at least one selected arc. It is conservative because it removes the infeasible pattern without making a claim about its individual arcs, any of which may still belong to another feasible route.

Suppose instead that all vehicle subproblems are feasible. Summing their optimal values gives the exact minimum emission $\mathcal{G}(\hat{x})$ attainable by the master route pattern. The baseline full-pattern cost cut is

$$\Theta \ge \mathcal{G}(\hat{x}) - (\mathcal{G}(\hat{x}) - \underline{\mathcal{G}})\left(|\Gamma(\hat{x})| - \sum_{(i,j)\in\Gamma(\hat{x})} x_{ij}\right). \tag{33}$$

When the same complete pattern is selected again, the mismatch term in (33) is zero and Θ must be at least $\mathcal{G}(\hat{x})$. If the pattern changes, the penalty relaxes the right-hand side toward the valid global bound $\underline{\mathcal{G}}$, thereby preserving validity. Thus, the baseline cuts guarantee exact treatment of every evaluated complete pattern. However, their limitation is locality. The no-good cut records only infeasibility of the full pattern, while the cost cut conveys little information to a solution that reuses only some evaluated arcs or one evaluated vehicle route.

4.3. Enhanced LBBD

The Enhanced LBBD retains both baseline cuts, and adds three complementary families of valid inequalities.

4.3.1 Piecewise-linear speed–emission cuts

The master problem $\hat{\mathbf{P}}_{\mathbf{M}}^{(n)}$ uses only the fastest possible traversal and therefore does not distinguish between different admissible travel times on a selected arc. Its estimate of schedule use and speed related emissions is therefore weak. To recover part of this relationship without reintroducing the nonlinear speed variable, let τ_{ij} denote arc travel time and write $v_{ij} = d_{ij}/\tau_{ij}$. The speed-dependent emission can then be expressed as the convex function

$$h_{ij}(\tau) = \frac{\beta_1}{d_{ij}}\tau + \frac{1}{2}\beta_3 \frac{d_{ij}^2}{\tau^2}.$$

The admissible interval is $\tau_{ij} \in [d_{ij}/\hat{v}_{ij}, d_{ij}./\hat{v}_{ij}]$. Let ξ_{ij} estimate this component in the master. For a set $\mathcal{T}$ of reference travel times, the supporting tangent at each $\bar{\tau} \in \mathcal{T}$ gives

$$\xi_{ij} \ge h_{ij}(\bar{\tau})x_{ij} + h'_{ij}(\bar{\tau})\left(\tau_{ij} - \bar{\tau}x_{ij}\right), \tag{34}$$

$$\frac{d_{ij}}{\overline{v}} x_{ij} \le \tau_{ij} \le \frac{d_{ij}}{\underline{v}} x_{ij}. \tag{35}$$

Constraints (35) ensures the activation of the tangent (34) only when $x_{ij} = 1$ and gives zero contribution when the arc is not selected. Taking several reference points constructs a piecewise-linear (PWL) lower envelope of h_{ij}. This approximation does not need to replicate the exact speed cost everywhere to remain effective. It ranks master patterns according to the travel-time resources they consume and transfers a valid portion of the continuous cost into every master solve.

4.3.2 Load-dependent emission bound

The preceding cuts represent only the speed-dependent part of arc emission. However, two routes using similar speeds may have different costs because they carry different passenger loads over different grades. Although the master already contains the load variable l_i, the emission contribution is active only when arc (i,j) is selected, creating the product $x_{ij}l_i$. Let $\lambda_{ij} = x_{ij}l_i$. Because x_{ij} is binary and $0 \le l_i \le Q$, its convex-hull linearization is

$$0 \le \lambda_{ij} \le Q x_{ij} \tag{36}$$

$$\lambda_{ij} \le l_i, \tag{37}$$

Just Accepted

$$\lambda_{ij} \geq l_i - Q(1 - x_{ij}). \quad (38)$$

Thus, λ_{ij} is zero on an unselected arc and equals the departure load l_i on a selected arc. Combining this exact product representation with (34) gives the aggregate master bound

$$\Theta \geq e \sum_{(i,j)\in\tilde{A}^+} d_v^k \left(\xi_{ij} + \frac{1}{2} g(\sin\theta + \beta_2 \cos\theta)(w x_{ij} + m\lambda_{ij})\right). \quad (39)$$

This inequality differentiates route patterns according to both speed-related and load-related emission. In particular, it penalizes carrying passengers over arcs for which the mass-and-grade coefficient is large, information that is absent from a purely distance- or minimum-speed-based lower bound. Because the load variables are already linked by the routing constraints in $\mathcal{X}$, the inequality strengthens the objective estimate without changing the underlying service decisions.

4.3.3 Vehicle-level Benders cuts

The full-pattern cut (33) loses most of its strength as soon as any vehicle route changes, even if several previously evaluated vehicle routes remain unchanged. This is especially restrictive in a multi-vehicle problem, where a neighborhood change for one vehicle creates a formally new global pattern. We therefore disaggregate the estimated emission by vehicle. Let Θ_k represent vehicle k's emission and let $\underline{\mathcal{G}}_k$ be a valid vehicle-level lower bound. For an evaluated route R_k, we add the following cuts:

$$\Theta \geq \sum_{k\in K} \Theta_k, \quad (40)$$

$$\Theta_k \geq \mathcal{G}_k(R_k) - \left(\mathcal{G}_k(R_k) - \underline{\mathcal{G}}_k\right)\left(|\Gamma_k(R_k)| - \sum_{(i,j)\in\Gamma_k(R_k)} x_{ij}\right). \quad (41)$$

If the same route R_k reappears, the mismatch term vanishes and (41) restores its exact evaluated cost. If another vehicle is rerouted, this information remains active through (40). The cut is therefore reusable across many global route patterns and is finer than the full-pattern cost cut. Symmetric copies can also transfer an evaluated route cost among identical vehicles when the expanded-network representation permits the corresponding relabeling.

The three enhancements are mutually reinforcing rather than substitutes. Arc-level inequalities provide broad but approximate information for unevaluated routes, whereas vehicle-level cuts provide exact but route-specific information after evaluation. The load bound connects the master load propagation to the emission estimate. Their validity is summarized next.

Proposition 5. *If $\bar{\tau}$ is in the admissible travel-time interval and $\underline{\mathcal{G}}_k$ is valid, then (34) – (41) are valid for every feasible expanded route.*

Proof. Since $f''(\tau) = 3\beta_3 d_{ij}^3/\tau^4 \geq 0$, each tangent underestimates f_{ij} on its admissible interval. Multiplying the tangent inequality by the binary selection variable and using the perspective form gives (34). An unselected arc has zero perspective contribution. For the load product, if $x_{ij} = 0$, (39) implies $0 \leq \lambda_{ij} \leq 0$, so $\lambda_{ij} = 0 = x_{ij} l_i$. If $x_{ij} = 1$, the inequalities imply $\lambda_{ij} \geq l_i$ and $\lambda_{ij} \leq l_i$, so $\lambda_{ij} = l_i$. Thus, the linearization is exact for binary. Finally, if an evaluated vehicle route reappears, the mismatch term in (41) is zero and the cut enforces its evaluated cost. If it differs in at least one arc, the right-hand side is no larger than the valid lower bound $\underline{\mathcal{G}}_k$. Summing valid vehicle bounds and adding the PWL and load bounds preserves validity. ■

4.4. Algorithmic framework

Algorithm 1 integrates the routing master, vehicle speed subproblems, and cut generation in a common exact procedure.

Algorithm 1: Logic-based Benders decomposition

Input: Instance data, enhancement flag ι_{enh}, tolerance ε_{opt}

Output: Best feasible plan S^* and bounds LB, UB.

Build the routing master and set $LB \leftarrow -\infty, UB \leftarrow +\infty$
if $\iota_{\text{enh}} = 1$ **then**
    Add valid inequalities (34) – (39)
end if
while time remains and $(UB - LB)/UB \geq \varepsilon_{\text{opt}}$ **do**
    Solve the master and update LB from its certified bound
    if the master has no incumbent route pattern **then**
        break
    end if
    Extract $\hat{x}$ and routes R_k; set $\mathcal{G}_{\text{cand}} \leftarrow 0$
    feasible←**true**
    for each active vehicle k **do**
        Solve the convex speed subproblem $\mathbf{P}_{\mathbf{SOCP}}^k$ for $R_k(\hat{x})$
        if the subproblem proves the route infeasible **then**
            feasible←**false; break**
        end if
        $\mathcal{G}_{\text{cand}} \leftarrow \mathcal{G}_{\text{cand}} + \mathcal{G}_k(R_k(\hat{x}))$
    end for
    if *feasible* = **false then**
        Add no-good cut (32) to the master $\hat{\mathbf{P}}_{\mathbf{M}}^{(n)}$
    else
        if $\mathcal{G}_{\text{cand}} < UB$ **then**
            $(UB, S^*) \leftarrow (\mathcal{G}_{\text{cand}}, \{R_k(\hat{x})\})$
        end if
        Add the full-pattern cost cut (33)
        if $\iota_{\text{enh}} = 1$ **then**
            Add vehicle cuts (40) – (41) and symmetric copies
        end if
    end if
end while
return S^*, UB, LB

At each iteration, the lower bound of the master problem $\hat{\mathbf{P}}_{\mathbf{M}}^{(n)}$ provides a global lower bound (LB). An upper bound (UB) is obtained only when every active vehicle subproblem is feasible. The sum of their exact optimal values then defines a feasible solution to the original model. The enhancement flag ι_{enh} distinguishes the two variants. Baseline LBBD uses the general lower bounds in $\hat{\mathbf{P}}_{\mathbf{M}}^{(n)}$ and adds feasibility and full pattern cuts. Enhanced LBBD adds arc and load inequalities before the first master solve. It also generates reusable vehicle-level cuts after feasible evaluations. Both variants use the same incumbent management, feasibility checks, and stopping rule. Proposition 6 establishes finite termination and exactness.

Proposition 6. *If every master problem and vehicle-level subproblem is solved exactly, both baseline and enhanced LBBD terminate after finitely many route-pattern evaluations.*

Proof. The expanded graph contains finitely many binary route patterns. If a route pattern is infeasible, Constraint (32) prevents the master problem from selecting this pattern again. If the pattern is feasible, Constraint (33) bounds the cost whenever this pattern recurs. Proposition 5 establishes the validity of the enhanced inequalities. These inequalities can strengthen the lower bound, whereas they cannot remove a feasible pattern. Since the number of route patterns is finite, the algorithm either terminates with time limitation or stop set after finitely many iterations. ■

5 Case study

Section 5.1 describes the experimental setup. Section 5.2 reports algorithmic results. Section 5.3 examines the operational implications.

5.1. Experimental setups

We generate the test instances from the Metropolitan Transportation

Just Accepted

Authority's Line 646 service, a benchmark used in prior FRT studies (Qiu et al., 2014; Quadrifoglio et al., 2007). The service region is a rectangle with 10 miles by 1 mile. It contains three equally spaced mandatory checkpoints, with scheduled departures at 0, 20, and 40 minutes. We place pickup-and-delivery requests in the two checkpoint segments. We generate the non-checkpoint time windows with five minutes. We consider ten instances for each $N = |\mathcal{R}| \in \{5,10,15,20\}$, with the corresponding fleet sizes $K \in \{1,2,3,4,5\}$. The $N = 5$ and $N = 10$ groups are treated as small instances, while $N = 15$ and $N = 20$ are treated as large instances. Instances ss-*N*-I is corresponding to the I-th instance within a small-scale (ss) group of N passengers.

All other parameters, including the checkpoint structure, vehicle characteristics, passenger mass, speed bounds, and time-window construction, follow Li et al., 2025 and Wu et al., 2024. The objective value produced by the model is tractive energy in MJ. We use $\eta_{\text{btw}} = 0.90$, $\eta_{\text{chg}} = 0.90$, and $EF_{\text{grid}} = 0.195$ kg CO_2/kWh, giving 0.06687 kg CO_2 per MJ of energy.

The algorithms are implemented in Python 3.9 and solved with Gurobi 13.0.2 on a PC equipped with an Intel Core i7 2.10 GHz processor and 32 GB of RAM. Each run uses one worker and one Gurobi thread. The overall time limit is 7,200 seconds, and an individual master solve is limited to 3,600 seconds.

5.2. Algorithmic performance

5.2.1 Small scale instance

The small-scale instances provide a controlled check of the decomposition and compare the direct MISOCP with both LBBD variants. Table 1 reports every instance and an average row for each request size. In the table, Obj. is the best feasible incumbent and LB is the certified lower bound, both in kg CO_2. Gaps are calculated by $\frac{UB-LB}{UB} \times 100\%$. Average rows report average values of the instance results. For $N = 5$, all three methods report zero gaps and the same objective values to the displayed precision. The average MISOCP runtime is 10.3 seconds, compared with 0.1 seconds for LBBD variants.

The distinction becomes clear at $N = 10$. MISOCP reaches the time limit on nine of ten instances and has an average gap of 0.10% with an average CPU time of 6,579.0 seconds. The baseline LBBD reaches the time limit on four instances and has an average gap of 0.72% with an average CPU time of 3,621.7 seconds. Enhanced LBBD closes all ten instances with zero displayed gap and an average CPU time of only 13.4 seconds. This reduction aligns with the underlying design of the algorithmic enhancements. The baseline master problem learns the cost of a route pattern only after its complete evaluation. In contrast, the enhanced master problem begins with valid arc and load emission bounds and continuously reuses vehicle level information across related global patterns.

Figure 3 shows the same distinction in the bound trajectories. The enhanced lower bound converges to the final optimum after only a few evaluations. In contrast, the baseline lower bound remains below the incumbent for numerous iterations and only rises in small increments near the end. The upper-bound trajectory shows that a feasible incumbent is usually found early. For the baseline method, the main computational burden is therefore certification.

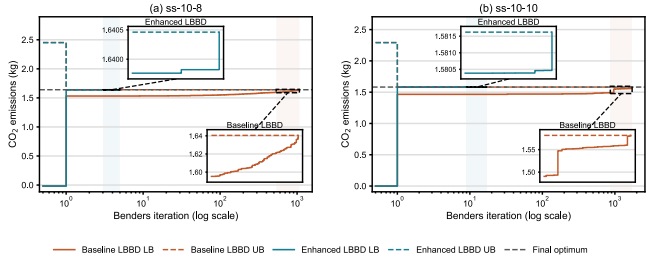


Fig. 3. Representative lower- and upper-bound trajectories for small-scale instances.

5.2.2 Large-scale instance

The large instances stress the master problem because the number of checkpoint compatible pickup and delivery patterns grows rapidly with customer size N. Table 2 reports computational performance of large-scale instances. For the $N = 15$ instances, MISOCP and baseline LBBD reach the 7,200 second limit in every instance and leave average gaps of 2.58% and 8.42%, respectively. Enhanced LBBD solves all ten instances with an average computational time of 919.0 seconds. The methods obtain the same incumbent values, while the enhanced master provides substantially stronger certificates of solution quality.

The $N = 20$ group is even more computational demanding. MISOCP and baseline LBBD reach the time limit on every instance with mean gaps of 6.47% and 9.26%, respectively. Enhanced LBBD certifies five of the ten instances and reduces the mean gap to 0.57%. Its average lower bound is 1.799 kg CO_2, compared with 1.690 kg for MISOCP and 1.640 kg for baseline LBBD. The mean incumbent is 1.809 kg for all three methods, although the incumbents differ in MISOCP and baseline LBBD reach the time limit on every instance with mean gaps of 6.47% and 9.26%, respectively. Enhanced LBBD certifies five of the ten instances and reduces the mean gap to 0.57%. Its average lower bound is 1.799 kg. Thus, the main advantage is stronger optimality certification rather than uniformly better feasible solutions. The enhanced method is not uniformly fast at this scale because solving the master problem can become the dominant computational cost.

Fig. 4 shows the representative trajectories. The enhanced lower bound makes a substantial initial jump toward the final objective before closing the remaining gap in a small number of evaluations. The baseline lower bound instead improves gradually over hundreds or thousands of iterations, even when its corresponding upper bound is already near the best feasible emission level. These cuts therefore improve the algorithm's capability of certifying optimality.

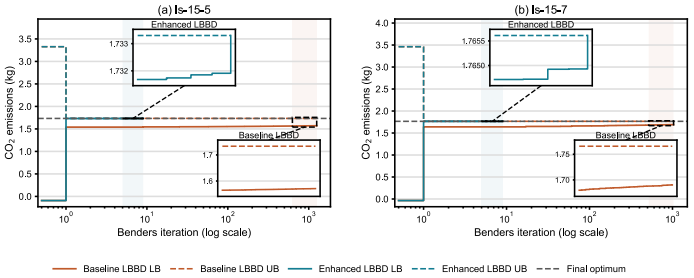


Fig. 4. Representative lower- and upper-bound trajectories for large-scale instances.

Just Accepted

Table 1. Computational performance on the small-scale instances

Instance	MISOCP				Baseline LBBD				Enhanced LBBD			
	Obj. (kg)	LB (kg)	Gap (%)	CPU (s)	Obj. (kg)	LB (kg)	Gap (%)	CPU (s)	Obj. (kg)	LB (kg)	Gap (%)	CPU (s)
ss-5-1	0.923	0.923	0.00	9.5	0.923	0.923	0.0	0.1	0.923	0.923	0.0	0.1
ss-5-2	0.835	0.835	0.00	17.6	0.835	0.835	0.0	0.0	0.835	0.835	0.0	0.1
ss-5-3	0.799	0.799	0.00	4.0	0.798	0.798	0.0	0.0	0.798	0.798	0.0	0.1
ss-5-4	0.847	0.847	0.00	5.8	0.847	0.847	0.0	0.0	0.847	0.847	0.0	0.1
ss-5-5	0.846	0.846	0.00	6.3	0.845	0.845	0.0	0.0	0.845	0.845	0.0	0.1
ss-5-6	0.838	0.838	0.00	4.8	0.838	0.838	0.0	0.1	0.838	0.838	0.0	0.1
ss-5-7	0.877	0.877	0.00	3.8	0.877	0.877	0.0	0.1	0.877	0.877	0.0	0.1
ss-5-8	0.982	0.982	0.00	11.3	0.982	0.982	0.0	0.1	0.982	0.982	0.0	0.1
ss-5-9	0.858	0.858	0.00	33.1	0.859	0.859	0.0	0.3	0.859	0.859	0.0	0.1
ss-5-10	0.866	0.866	0.00	7.3	0.865	0.865	0.0	0.0	0.865	0.865	0.0	0.1
Avg.	**0.867**	**0.867**	**0.00**	**10.3**	**0.867**	**0.867**	**0.0**	**0.1**	**0.867**	**0.867**	**0.0**	**0.1**
ss-10-1	0.969	0.969	0.00	990.1	0.969	0.969	0.0	0.1	0.969	0.969	0.0	0.2
ss-10-2	1.725	1.722	0.21	7200.0	1.725	1.707	1.0	7,200.0	1.725	1.725	0.0	13.7
ss-10-3	1.118	1.115	0.19	7200.0	1.117	1.117	0.0	0.7	1.117	1.117	0.0	0.3
ss-10-4	1.798	1.796	0.09	7200.0	1.798	1.769	1.6	7,200.0	1.798	1.798	0.0	22.0
ss-10-5	1.676	1.674	0.15	7200.0	1.676	1.662	0.8	7,200.0	1.676	1.676	0.0	15.1
ss-10-6	1.176	1.176	0.01	7200.0	1.176	1.176	0.0	2.3	1.176	1.176	0.0	0.5
ss-10-7	1.260	1.260	0.02	7200.0	1.260	1.260	0.0	2.2	1.260	1.260	0.0	0.5
ss-10-8	1.640	1.640	0.00	7200.0	1.640	1.640	0.0	2,298.4	1.640	1.640	0.0	13.2
ss-10-9	1.636	1.632	0.25	7200.0	1.636	1.575	3.7	7,200.0	1.636	1.636	0.0	34.2
ss-10-10	1.582	1.580	0.10	7200.0	1.582	1.582	0.0	5,113.8	1.582	1.582	0.0	34.5
Avg.	**1.458**	**1.456**	**0.10**	**6579.0**	**1.458**	**1.446**	**0.7**	**3,621.7**	**1.458**	**1.458**	**0.0**	**13.4**

Table 2. Computational performance on the large-scale instances

Instance	MISOCP				Baseline LBBD				Enhanced LBBD			
	Obj. (kg)	LB (kg)	Gap (%)	CPU (s)	Obj. (kg)	LB (kg)	Gap (%)	CPU (s)	Obj. (kg)	LB (kg)	Gap (%)	CPU (s)
ss-5-1	1.614	1.597	1.05	7200.0	1.614	1.489	7.8	7,200.0	1.614	1.614	0.0	334.6
ss-5-2	1.758	1.738	1.17	7200.0	1.758	1.648	6.3	7,200.0	1.758	1.758	0.0	665.8
ss-5-3	1.694	1.640	3.14	7200.0	1.694	1.547	8.6	7,200.0	1.694	1.694	0.0	1,078.4
ss-5-4	1.734	1.716	1.02	7200.0	1.734	1.563	9.9	7,200.0	1.734	1.734	0.0	648.1
ss-5-5	1.733	1.704	1.68	7200.0	1.733	1.571	9.4	7,200.0	1.733	1.733	0.0	713.8
ss-5-6	1.775	1.713	3.47	7200.0	1.775	1.629	8.2	7,200.0	1.775	1.775	0.0	532.1
ss-5-7	1.766	1.752	0.76	7200.0	1.765	1.690	4.3	7,200.0	1.765	1.765	0.0	153.4
ss-5-8	2.051	1.883	8.20	7200.0	2.051	1.800	12.3	7,200.0	2.051	2.051	0.0	629.2
ss-5-9	1.824	1.805	1.04	7200.0	1.824	1.667	8.6	7,200.0	1.824	1.824	0.0	456.1
ss-5-10	1.871	1.792	4.26	7200.0	1.871	1.705	8.9	7,200.0	1.871	1.871	0.0	3,978.4
Avg.	**1.782**	**1.734**	**2.58**	**7200.0**	**1.782**	**1.631**	**8.4**	**7,200.0**	**1.782**	**1.782**	**0.0**	**919.0**
ss-10-1	1.817	1.662	8.50	7200.0	1.817	1.625	10.6	7,200.0	1.817	1.811	0.3	7,200.0
ss-10-2	1.831	1.709	6.69	7200.0	1.831	1.660	9.4	7,200.0	1.831	1.831	0.0	1,809.9
ss-10-3	1.782	1.701	4.53	7200.0	1.782	1.630	8.6	7,200.0	1.782	1.781	0.0	7,200.0
ss-10-4	1.740	1.683	3.26	7200.0	1.740	1.604	7.8	7,200.0	1.740	1.740	0.0	2,681.6
ss-10-5	1.726	1.636	5.22	7200.0	1.727	1.579	8.6	7,200.0	1.727	1.727	0.0	7,109.8
ss-10-6	1.821	1.675	8.02	7200.0	1.820	1.628	10.6	7,200.0	1.822	1.769	2.9	7,200.0
ss-10-7	1.748	1.707	2.38	7200.0	1.748	1.636	6.4	7,200.0	1.748	1.748	0.0	4,413.0
ss-10-8	1.929	1.647	14.59	7200.0	1.929	1.652	14.3	7,200.0	1.929	1.929	0.0	6,188.5
ss-10-9	1.891	1.789	5.40	7200.0	1.890	1.739	8.0	7,200.0	1.891	1.872	1.0	7,200.1
ss-10-10	1.804	1.694	6.06	7200.0	1.804	1.652	8.4	7,200.0	1.804	1.776	1.5	7,200.1
Avg.	**1.809**	**1.690**	**6.47**	**7200.0**	**1.809**	**1.640**	**9.3**	**7,200.0**	**1.809**	**1.798**	**0.6**	**5,820.3**

5.3. Operational implications

The operational study contains six experiments. It covers joint route and speed optimization, discrete speed menus, capacity linked curb mass, road grade, time window width, and aerodynamic sensitivity.

5.3.1 Value of integrated model

To quantify the value of coordinating route selection with endogenous speed decisions, we use a two-stage benchmark. The first stage solves a classical shortest-distance routing problem based on a standard distance-minimizing objective function

$$\min \sum_{i,j\in\mathcal{A}^+} d_{ij}x_{ij}.$$

The second stage then evaluates this first-stage optimal route using the isolated speed optimization model $\mathbf{P}^k_{\mathbf{SOCP}}$. In contrast to this sequential process, our integrated framework optimizes the physical route and the continuous speed profile simultaneously. Fig. 5 presents the resulting objective values for every comparable instance.

Several benchmark runs do not produce a feasible speed schedule. This shows that routing based only on distance can conflict with strict passenger time windows. Even in feasible cases, the sequential benchmark leads to higher emissions in both size groups. Mean emissions increase by 37.7% for $N = 15$ and 30.8% for $N = 20$. This shows that distance alone cannot accurately reflect energy use and emissions in this problem. A longer route can offer a better load sequence and arrival schedule, allowing the vehicle to travel at speeds with lower emissions. In contrast, joint route and speed planning considers both schedule feasibility and emissions, especially when passenger requests have strict time windows.

5.3.2 Benefit of continuous speed control

We examine the impact of speed control policies on the emissions. Figure 6 compares the continuous-speed model with three discrete-speed policies with speed increments Δv of 10, 5, and 1 mph.

Just Accepted

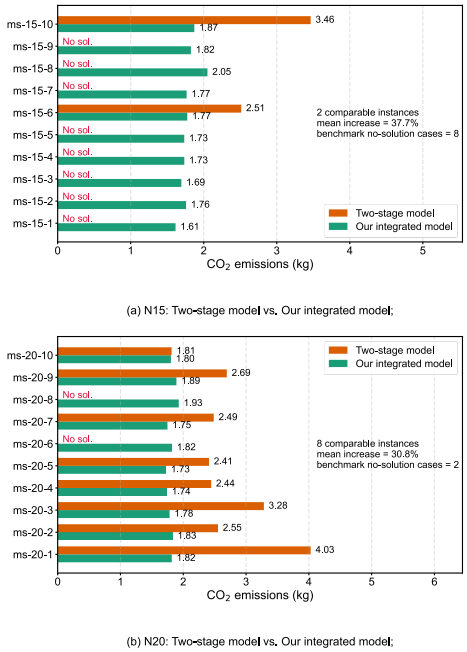


Fig. 5. Integrated routing and speed planning versus a two-stage distance-based benchmark.

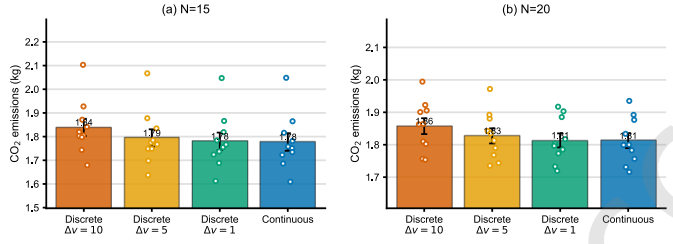


Fig. 6. CO_2 emissions under continuous and discrete speed policies (dots indicate individual instance values).

For both $N = 15$ and $N = 20$, the 10-mph grid yields the highest average emissions. Reducing the step size to 5 mph narrows this gap. The 1-mph grid performs nearly identical to our continuous speed model. This is because a coarse grid can enforce a vehicle traveling away from the optimal speed that balances the reciprocal travel-time term and the quadratic speed term. Although a coarse speed is easier to implement, its resulting emission penalty becomes prominent in larger instances.

5.3.3 Impact of capacity

We study the impact of vehicle capacity on emissions, with Q varying from 3 to 8. We define the vehicle curb mass as

$$w(Q; c) = 6350 + c(Q - 5)\text{kg}, c = \{0,100,200,500\}\text{kg per seat.}$$

Figure 7(a) – (b) shows that the emission benefit of greater capacity depends on the mass increment. For $N = 15$, increasing capacity from 5 to 8 gives mean emissions of 1.548, 1.602, and 1.763 kg CO_2 for $c = 100$, 200 and 500, compared with 1.782 kg at $Q = 5$. The reductions are 13.11%, 10.09%, and 1.04%. Under fixed mass, the corresponding mean at $Q = 8$ is 1.495 kg. For $N = 20$, the $Q = 8$ means are 1.800, 1.865, and 2.060 kg CO_2, compared with 1.809 kg at Q=5. The smallest coefficient yields a 0.50% reduction, while the two larger coefficients yield increases of 3.08% and 13.88%. For $c = 500$, the mean at $Q = 8$ is slightly higher than the $Q = 7$ mean of 1.752 kg for $N = 15$. Thus, the linked-mass curve is not monotonic.

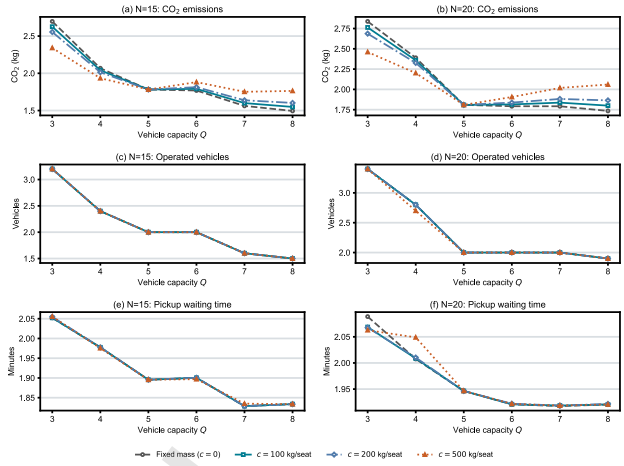


Fig. 7. Impact of capacity on emissions and operations.

Figure 7(d) shows that mean fleet use for $N = 20$ falls from 2.0 at $Q = 5$ to 1.9 at $Q = 8$ under all three positive coefficients. However, using fewer vehicles does not necessarily reduce emissions if the vehicles are heavier. This result reflects a trade-off between consolidating routes and the mass-related energy use of the vehicles that remain in service. A smaller mass increment preserves more of the capacity benefit. A larger increment can reduce or reverse that benefit. The preferred capacity therefore depends on vehicle design, demand, and service constraints. The waiting-time panels show service outcomes that operators can assess alongside emissions.

5.3.4 Impact of road grade

We evaluate the effect of terrain by generating six scenarios with maximum absolute grades of 0%, 1%, 2%, 3%, 4%, and 5% for every existing large instance. The flat scenario is the reference and all demand, service, fleet, and energy parameters remain unchanged. We compare emissions and physical directed arc sets. Route dissimilarity is measured as $100(1 - |A_g \cap A_0|/|A_g \cup A_0|)$, averaged over the ten instances in each request group, with vehicle labels ignored.

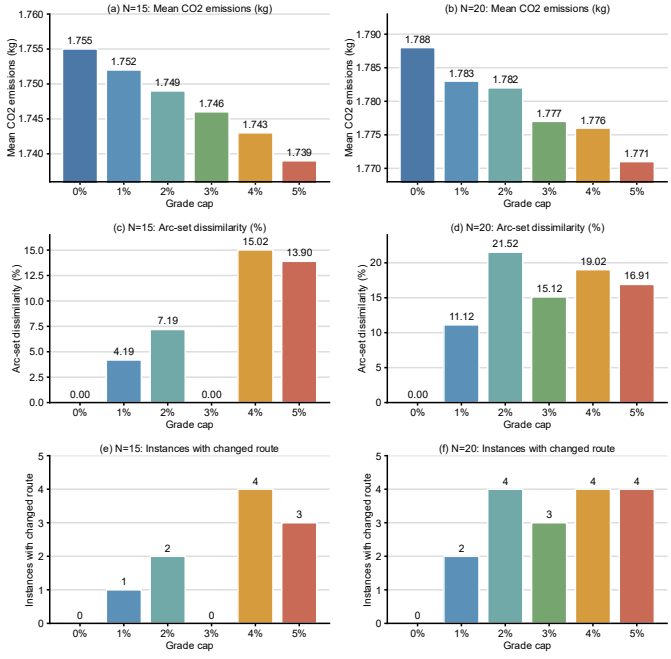


Fig. 8. Road grade sensitivity of CO_2 emissions and route changes.

Figure 8 shows that at a 5% grade cap mean emissions change from 1.755 to 1.739 kg CO_2 for $N = 15$ and from 1.788 to 1.771 kg CO_2 for

Just Accepted

$N = 20$. The reductions are 0.89% and 0.95%. Routes change in 3 of 10 in $N = 15$ instances and 4 of 10 in $N = 20$ instances. Mean directed arc dissimilarity is 13.90% and 16.91%, while the dissimilarity reaches 21.52% at a 2% cap in $N = 20$ instance.

The relatively small net emission change reflects the offset between uphill energy use and downhill recovery when route endpoints have matched elevations. The route changes show that terrain information affects the physical routing decisions. Therefore, terrain data should be considered when route structure is sensitive to elevation.

5.3.5 Impact of time window

We vary customer time window widths over 1, 3, 5, 7, and 10 minutes for the existing instances. The five-minute width is the reference, and all other demand, terrain, fleet, and service settings remain unchanged. The standard solution procedure is applied to the same ten instances in each group, and emissions, route distance, operated vehicles, and pickup waiting time are recorded.

Figure 9 shows that tightening the width from five minutes to one-minute increases mean emissions by 14.78% for $N = 15$ and 15.13% for $N = 20$. Mean route distance increases by about 10.3% and 11.0%, and two instances in each group require a third vehicle. Widening the width to ten minutes reduces mean emissions by 3.95% and 3.34%. The waiting measure starts at the opening of the window, so widening the window can increase recorded waiting while its center remains fixed.

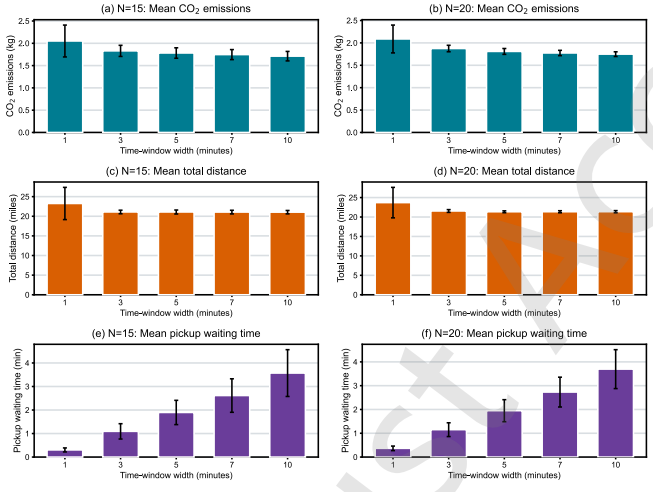


Fig. 9. Impact of time window on system performance.

5.3.6 Impact of aerodynamic sensitivity

The coefficient β_3 reflects the vehicle's aerodynamic characteristics and its sensitivity to speed. To evaluate the impact of this aerodynamic term, we multiplied β_3 by a scaling factor $\lambda \in \{0.5, 0.75, 1, 1.5, 2\}$ across all large-scale instances.

Figure 10 shows that mean emissions increase monotonically with the multiplier. For $N = 15$, the mean rises from 1.632 to 2.079 kg CO_2 as λ increases from 0.5 to 2. For $N = 20$, it rises from 1.659 to 2.107 kg CO2. These changes are approximately 27.4% and 27.0%.

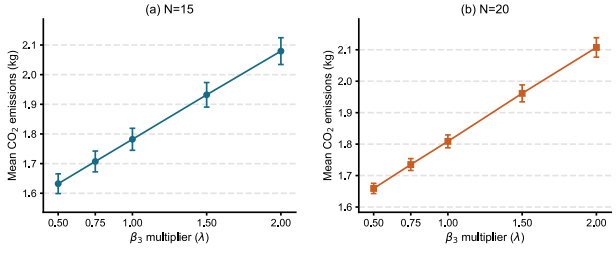


Fig. 7. Sensitivity of emissions to the aerodynamic coefficient.

These results demonstrate that aerodynamic efficiency is a critical lever in green vehicle routing. Because emissions are highly sensitive to this parameter, integrating specific vehicle aerodynamic profiles into the planning phase can yield much more accurate emissions estimates and more eco-friendly speed profiles.

6 Conclusion

We formulate an emission-aware flex-route transit model that combines scheduled checkpoints with demand-responsive service. The model captures vehicle assignment, passenger precedence, vehicle capacity, time windows, and speed-dependent tractive energy on a checkpoint-expanded network.

To solve the proposed model, we develop an exact logic-based Benders decomposition (LBBD) method and an enhanced variant. The baseline method develops standard no-good optimality and feasibility cuts. The enhanced method adds piecewise-linear approximation cuts, load-dependent emission cuts, and aggregated vehicle-level cuts. The computational results show that these enhancements tighten the master problem, improve gaps by 9%, and reduce computation time by two orders of magnitude. The operational results also show that integrated route and speed optimization can reduce emissions by nearly 30% compared with the traditional distance-based method.

We identify three directions for future research. First, future studies can extend this framework to handle stochastic demand, uncertain travel times, and heterogeneous fleets. Second, it is recommended to develop a branch-and-price algorithm to handle larger networks. Finally, we could also use demand and behavior data to calibrate passenger choices, and then optimize the schedules.

Author contributions

Shuiwang Chen: Writing – original draft, Validation, Software, Methodology, Formal analysis, Data curation. **Mingyang Li:** Writing – review & editing, Formal Analysis, Methodology. **Lingxiao WU:** Conceptualization, review & editing, Supervision, Methodology, Funding acquisition.

Replication and data sharing

The source codes and replication package is available on ETS data under ID ETS-Data-26-00112.

Acknowledgements

This paper is supported by Shenzhen Science and Technology Program, China (Grant number: JCYJ20240813162012016) and The Hong Kong Polytechnic University (Grant number: JSS2627-019).

Declaration of competing interest

The authors have no competing interests to declare that are relevant to the content of this article.

References

Akturk, D., Candogan, O., Gupta, V., 2025. Managing Resources for Shared Micromobility: Approximate Optimality in Large-Scale Systems. Management Science 71, 5676–5695. https://doi.org/10.1287/mnsc.2022.02023

Banister, D., 2008. The sustainable mobility paradigm. Transport

Just Accepted

Policy, New Developments in Urban Transportation Planning 15, 73–80. https://doi.org/10.1016/j.tranpol.2007.10.005

Barth, M., Boriboonsomsin, K., 2008. Real-World Carbon Dioxide Impacts of Traffic Congestion. Transportation Research Record 2058, 163–171. https://doi.org/10.3141/2058-20

Bektaş, T., Laporte, G., 2011. The Pollution-Routing Problem. Transportation Research Part B: Methodological, Supply chain disruption and risk management 45, 1232–1250. https://doi.org/10.1016/j.trb.2011.02.004

Chen, S., Wu, L., Ng, K.K.H., Liu, W., Wang, K., 2024. How airports enhance the environmental sustainability of operations: A critical review from the perspective of Operations Research. Transportation Research Part E: Logistics and Transportation Review 183, 103440. https://doi.org/10.1016/j.tre.2024.103440

Chen, X., Wang, Yinhai, Wang, Yong, Qu, X., Ma, X., 2021. Customized bus route design with pickup and delivery and time windows: Model, case study and comparative analysis. Expert Systems with Applications 168, 114242. https://doi.org/10.1016/j.eswa.2020.114242

Dabia, S., Demir, E., Woensel, T.V., 2017. An Exact Approach for a Variant of the Pollution-Routing Problem. Transportation Science (INFORMS) 51, 607–628. https://doi.org/10.1287/trsc.2015.0651

Demir, E., Bektaş, T., Laporte, G., 2012. An adaptive large neighborhood search heuristic for the Pollution-Routing Problem. European Journal of Operational Research 223, 346–359. https://doi.org/10.1016/j.ejor.2012.06.044

Fukasawa, R., He, Q., Santos, F., Song, Y., 2018. A Joint Vehicle Routing and Speed Optimization Problem. INFORMS Journal on Computing 30, 694–709. https://doi.org/10.1287/ijoc.2018.0810

Fukasawa, R., He, Q., Song, Y., 2016. A disjunctive convex programming approach to the pollution-routing problem. Transportation Research Part B: Methodological 94, 61–79. https://doi.org/10.1016/j.trb.2016.09.006

Hooker, J.N., Ottosson, G., 2003. Logic-based Benders decomposition. Math. Program., Ser. A 96, 33–60. https://doi.org/10.1007/s10107-003-0375-9

Lai, D., Costa, Y., Demir, E., Florio, A.M., Van Woensel, T., 2024. The pollution-routing problem with speed optimization and uneven topography. Computers & Operations Research 164, 106557. https://doi.org/10.1016/j.cor.2024.106557

Li, M., Tang, J., Feng, T., 2025. An exact method to solve the flex-route transit operational planning problem considering energy consumption. International Journal of Production Research 63, 6155–6177. https://doi.org/10.1080/00207543.2025.2469287

Ma, B., Hu, D., Chen, X., Wang, Y., Wu, X., 2021. The vehicle routing problem with speed optimization for shared autonomous electric vehicles service. Computers & Industrial Engineering 161, 107614. https://doi.org/10.1016/j.cie.2021.107614

Mageean, J., Nelson, J.D., 2003. The evaluation of demand responsive transport services in Europe. Journal of Transport Geography 11, 255–270. https://doi.org/10.1016/S0966-6923(03)00026-7

Miao, C., Liu, H., Zhu, G.G., Chen, H., 2018. Connectivity-based optimization of vehicle route and speed for improved fuel economy. Transportation Research Part C: Emerging Technologies 91, 353–368. https://doi.org/10.1016/j.trc.2018.04.014

Miller, P., de Barros, A.G., Kattan, L., Wirasinghe, S.C., 2016. Public transportation and sustainability: A review. KSCE J Civ Eng 20, 1076–1083. https://doi.org/10.1007/s12205-016-0705-0

Molenbruch, Y., Braekers, K., Eisenhandler, O., Kaspi, M., 2023. The Electric Dial-a-Ride Problem on a Fixed Circuit. Transportation Science 57, 594–612. https://doi.org/10.1287/trsc.2023.1208

Qiu, F., Li, W., Haghani, A., 2015. A methodology for choosing between fixed-route and flex-route policies for transit services. Journal of Advanced Transportation 49, 496–509. https://doi.org/10.1002/atr.1289

Qiu, F., Li, W., Zhang, J., 2014. A dynamic station strategy to improve the performance of flex-route transit services. Transportation Research Part C: Emerging Technologies 48, 229–240. https://doi.org/10.1016/j.trc.2014.09.003

Quadrifoglio, L., Dessouky, M.M., Ordóñez, F., 2008. Mobility allowance shuttle transit (MAST) services: MIP formulation and strengthening with logic constraints. European Journal of Operational Research 185, 481–494. https://doi.org/10.1016/j.ejor.2006.12.030

Quadrifoglio, L., Hall, R. W., & Dessouky, M. M. (2006). Performance and design of mobility allowance shuttle transit services: bounds on the maximum longitudinal velocity. Transportation Science, 40(3), 351-363.

Sartori Vieira, F., Sörensen, K., Vanstenwegen, P., 2025. Optimization of a semiflexible demand-responsive feeder bus system using variable neighborhood search. International Transactions in Operational Research 32, 3413–3440. https://doi.org/10.1111/itor.13616

Shahin, R., Hosteins, P., Pellegrini, P., Vandanjon, P.-O., Quadrifoglio, L., 2024. A survey of Flex-Route Transit problem and its link with Vehicle Routing Problem. Transportation Research Part C: Emerging Technologies 158, 104437. https://doi.org/10.1016/j.trc.2023.104437

Tang, J., Yang, H., Li, M., 2026. The flex-route transit operational planning problem with electric vehicles. Transportation Research Part C: Emerging Technologies 188, 105701. https://doi.org/10.1016/j.trc.2026.105701

Vansteenwegen, P., Melis, L., Aktaş, D., Montenegro, B.D.G., Sartori Vieira, F., Sörensen, K., 2022. A survey on demand-responsive public bus systems. Transportation Research Part C: Emerging Technologies 137, 103573. https://doi.org/10.1016/j.trc.2022.103573

Wang, S., 2016. Fundamental properties and pseudo-polynomial-time algorithm for network containership sailing speed optimization. European Journal of Operational Research 250, 46–55. https://doi.org/10.1016/j.ejor.2015.10.052

Wang, S., Meng, Q., 2012. Sailing speed optimization for

Just Accepted

container ships in a liner shipping network. Transportation Research Part E: Logistics and Transportation Review 48, 701–714. https://doi.org/10.1016/j.tre.2011.12.003

Wu, F., Adulyasak, Y., Cordeau, J.-F., 2024. Modeling and Solving the Traveling Salesman Problem with Speed Optimization for a Plug-In Hybrid Electric Vehicle. Transportation Science (INFORMS) 58, 562–577. https://doi.org/10.1287/trsc.2023.0247

Wu, F., Ye, H., Bektaş, T., Dong, M., 2025. New and tractable formulations for the eco-driving and the eco-routing-and-driving problems. European Journal of Operational Research 321, 445–461. https://doi.org/10.1016/j.ejor.2024.10.005

Zhen, L., Zhuge, D., Zhang, S., Wang, S., Psaraftis, H.N., 2024. Optimizing Sulfur Emission Control Areas for Shipping. Transportation Science (INFORMS) 58, 614–638. https://doi.org/10.1287/trsc.2023.0278

Zheng, Y., Li, W., Qiu, F., 2018. A slack arrival strategy to promote flex-route transit services. Transportation Research Part C: Emerging Technologies 92, 442–455. https://doi.org/10.1016/j.trc.2018.05.015

Zheng, Y., Li, W., Qiu, F., Wei, H., 2019. The benefits of introducing meeting points into flex-route transit services. Transportation Research Part C: Emerging Technologies 106, 98–112. https://doi.org/10.1016/j.trc.2019.07.012

Zhuge, D., Wu, J., Zhen, L., Wang, S., & Wang, Y. (2026). Liner fleet deployment and speed optimization under emission reduction technologies. Transportation Science, 60(3), 508-526.

Author biography



Shuiwang Chen received his Bachelor's and Master's degrees in transportation from Southwest Jiaotong University in Chengdu, China. He has been pursuing his Ph.D. in transportation at The Hong Kong Polytechnic University since 2024. His research interests include large-scale combinatorial optimization and decomposition algorithms.



Mingyang Li received the B.S. degree in vehicle engineering from Dalian Jiaotong University (DJTU), Dalian, China, in 2019, and the Ph.D. degree at the Central South University, Changsha, China in 2025. He is currently a postdoc fellow at The Hong Kong Polytechnic University. His research interests include vehicle routing problem.



Lingxiao Wu is currently an Assistant Professor with the Department of Aeronautical and Aviation Engineering, The Hong Kong Polytechnic University. His research interests include the design and implementation of methods that combine operations research and data analytics for solving problems arising in logistics management and transportation. His research has appeared in leading international journals, including Operations Research, Transportation Research—Part B: Methodological, Transportation Research—Part E: Logistics and Transportation, and Transportation Science.

Just Accepted