Optimizing multi-modal ride-matching with transfers

Stokkink, Patrick; de Palma, André; Geroliminis, Nikolas

DOI
10.1016/j.trc.2025.105274

Publication date
2025

Document Version
Final published version

Published in
Transportation Research Part C: Emerging Technologies

Transportation Research Part C 179 (2025) 105274

Contents lists available at ScienceDirect

Transportation Research Part C

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

Optimizing multi-modal ride-matching with transfers

Patrick Stokkink [a],*, André de Palma [b], Nikolas Geroliminis [c]

[a] *Department of Engineering Systems and Services, Faculty of Technology, Policy and Management, TU Delft, Netherlands*
[b] *THEMA, CY Cergy Paris University, France*
[c] *École Polytechnique Fédérale de Lausanne (EPFL), Urban Transport Systems Laboratory (LUTS), Switzerland*

ARTICLE INFO

Keywords:
Multi-modal
Ride-matching
Ride-sharing
Transfers
Integer programming
Simulated annealing

ABSTRACT

One of the limitations of ride-sharing is that matched drivers and riders need to have similar itineraries and desired arrival times for ride-sharing to be competitive against other transport modes. By allowing a single transfer at a designated transfer hub, their itineraries need to be only partially similar, and therefore more matching options are created. In this paper, we develop an optimal matching approach that matches riders to drivers, taking into account multi-modal routing options to model competition and collaboration between multiple modes of transport. We allow for transfers between modes and between multiple drivers. We model this as a path-based integer programming problem and we develop a simulated annealing algorithm to efficiently solve realistic large-scale instances of the problem.

Our analysis indicates that a single transfer hub can reduce significantly the average generalized cost of riders and the total vehicle hours traveled by creating efficient matches. As opposed to previous studies, our work shows that ride-sharing not only attracts former public transport users but also former private car users. By allowing for intermodal transfers and by choosing the cost parameters such that transfers are favorable, itineraries where commuters use their car first, before sharing a ride on the second part of their journey, becomes an appealing alternative. Multi-modal ride-matching with transfers has the potential to increase ride-sharing, reduce the number of vehicle hours traveled in private cars, and reduce the number of cars that are present in urban areas during peak hours of congestion.

1. Introduction

Transport is one of the main sources of CO_2 emissions and possibly the source individual commuters have the most direct influence on. Ride-matching, also known as carpooling or ride-sharing, as an alternative to traveling alone by car is known to reduce CO_2 emissions and traffic congestion in large-scale networks (Alisoltani et al., 2021). Thereby, it is expected to reduce car ownership in the long term. Nevertheless, while ride-sharing (not to be confused with ride-hailing which is like a taxi service) is considered a feasible solution to fight congestion and car ownership, its market share remains small with average car occupancy in the US being only 1.5 in 2019 (U.S. Department of Energy (DOE), 2022).

One of the reasons that current carpooling systems are not successful is the lack of a central operator who can match riders in a multi-modal system with different itineraries in a reliable and efficient way for all parties involved. An important operational limitation of direct ride-sharing is that a pairing of drivers and riders needs to be found with matchable itineraries. This means that a driver needs to be able to pick up and drop off the matched rider without deviating too much from their original route. In

* Corresponding author.
E-mail addresses: patrick.stokkink@epfl.ch (P. Stokkink), andre.de-palma@cyu.fr (A. de Palma), nikolas.geroliminis@epfl.ch (N. Geroliminis).

https://doi.org/10.1016/j.trc.2025.105274
Received 24 June 2024; Received in revised form 5 July 2025; Accepted 7 July 2025
Available online 5 August 2025
0968-090X/

addition to this, desired arrival times of the rider and the driver need to be similar. Dissimilar matches increase the costs of drivers and riders, such that ride-sharing is no longer competitive with private or public transport.

A potential solution for this is to allow riders to transfer between drivers and between multiple modes of transportation. In this work, we consider the possibility of a single transfer, since more than one transfer makes ride-sharing a less appealing alternative. By allowing transfers, a larger set of potential matches is available for drivers and riders since the itineraries only need to be partially similar, as they spend only a part of their trip together. Transfers may also promote family carpooling at least for part of their journey. As families share their origin but do not necessarily share their destination, they can spend the first part of their trip together before one of them transfers. According to Li et al. (2007), family carpooling makes up nearly 75% of all carpools.

The framework presented in this paper has the potential to enhance commuter mobility and improve the accessibility of transport networks. Consequently, a platform capable of facilitating the matching of riders and drivers in a multi-modal system could prove compelling for cities, public transport organizations, or private entities. Cities may find value in investing in such a platform as it aligns with their objective of enhancing the mobility of residents. For public transport operators, a multi-modal platform could be advantageous for more effectively integrating their services by providing first or last-mile connections to other modes of transportation. Private organizations may also find this platform appealing, whether through advertising within the application or by charging a nominal fee for their matching services.

1.1. State of the art

In ride-sharing, individual commuters share a ride for a part of their journey, which reduces the time a driver is traveling with a partially empty vehicle. Teal (1987) provides an early definition of carpooling and distinguishes between different types of carpoolers. Shaheen and Cohen (2019) give an overview of the various shared-ride services that exist in the modern day. The two most important ones are ride-sharing (also known as carpooling, where commuters that have a predefined trip purpose share a ride) and ride-hailing (also known as ride-sourcing, which is more similar to a taxi service). Whereas studies have shown that ride-hailing generally leads to an increase in congestion (Beojone and Geroliminis, 2021; Schaller, 2021), ride-sharing generally reduces congestion by increasing vehicle occupancy and thereby reducing the number of vehicles on the road (Caulfield, 2009; Gurumurthy et al., 2019; de Palma et al., 2022a). Ride-sharing may lead to environmental and societal benefits but brings forth many optimization challenges. For a review of those challenges, the reader is referred to Agatz et al. (2012).

One of the most important challenges is the matching of drivers and riders. The literature has focused on finding optimal matchings (Özkan and Ward, 2020; de Palma et al., 2022a,b) as well as stable matchings (Wang et al., 2018; Yan et al., 2021). In a stable matching, no rider and driver can improve their current situation by matching with each other. Matching problems can be classified as static or dynamic problems. In static problems, all information on possible drivers and riders is known in advance. In dynamic problems, information is made available gradually and the matching is made partially based on the available drivers and riders. In this work, we consider a static setting, where the complete set of potential drivers and riders is known in advance, but uncertainty exists about the travel times.

The matching problem can be extended to include transfers between various vehicles. Herbawi and Weber (2012b) model the ride-matching problem with transfers and time windows and use a genetic algorithm to solve this problem. Masoud and Jayakrishnan (2017a) consider multi-hop ride-matching where a driver can carry multiple riders and riders can join multiple drivers. Various alternatives solution approaches to the multi-hop ride-sharing have been considered, such as a chromosome-based approach (Cheikh and Hammadi, 2016), a station-first algorithm (Xu et al., 2020), and a genetic algorithm (Herbawi and Weber, 2012a). Huang et al. (2018) include carpooling in the trip planning of commuters next to public transport. Commuters are allowed to transfer between drivers or between modes. Lu et al. (2020) consider ride-sharing with transfers in short-notice evacuations such as during natural or man-made disasters. Ride-sharing with transfers carries many similarities with multi-stage crowd-shipping. In that case, packages are transported by drivers rather than riders, and these may also be transferred from one driver to another (Raviv and Tenzer, 2018; Chen et al., 2018, 2019; Li et al., 2024). Despite the similarities, we note that riders and packages are very different due to their perceived inconvenience from transfers and waiting time.

In the literature, carpooling has been modeled both as a competitor of public transport (Li et al., 2021) or as a complement to public transport (Kong et al., 2020). The former considers public transport as an alternative mode of transport (de Palma et al., 2022a,b). In this case, carpooling can reduce public transport users and therefore has negative societal effects. The latter considers public transport as a feeder to carpooling or carpooling as a feeder to public transport (Masoud et al., 2017; Ma et al., 2019; Kumar and Khani, 2021). In that case, the two services may help to improve each other and form a competitive alternative against private car usage. In this work, we consider both alternatives simultaneously to properly consider the interaction of the two transport modes when riders are allowed to make transfers. Integrated on-demand mobility systems such as ride-sharing and vehicle-sharing have previously been considered in an integrated framework with other modes of transportation by Stiglic et al. (2018) and Enzi et al. (2024). Whereas (Stiglic et al., 2018) only consider ride-sharing as a feeder to public transit, we allow transfers to and from ride-sharing and even between two ride-sharing vehicles. Enzi et al. (2024) consider a vehicle sharing system for business trips, rather than a ride-sharing framework which is considered in this work. Chen et al. (2024) provide a synthesis of collaboration in multi-modal shared mobility systems.

Multi-modality is becoming more common for first- and last-mile transportation. Xu et al. (2024) consider a multi-modal hub-based system that combines ride-sourcing with public transit. Lee et al. (2024) assess the resilience of such a multi-modal system under disruptions. Yan et al. (2024) consider a neural-network-based approach to enhance recommendations in multi-modal

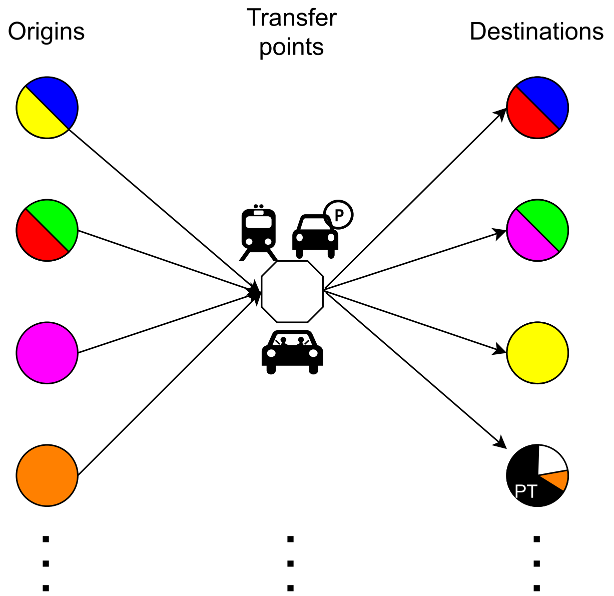

Fig. 1. Graphic illustration of transfers. Colors represent individual commuters. The lower-left half of a circle represents the driver and the upper-right half represents the rider. If drivers are not matched to a rider, the full circle is colored by the driver or rider.

transport. Liu et al. (2024) integrate ride-sharing into a multi-modal traffic model and Tang et al. (2024) consider the partitioning of a multi-modal transportation network using such traffic models, in this case through a 3D macroscopic fundamental diagram.

In this paper, we consider a multi-modal ride-matching problem. Riders are matched to drivers with the goal of minimizing the total costs associated with the commute of riders. For this, we consider their travel time, with mode-dependent cost parameters. Additional mode-dependent costs may also be considered, such as parking and fuel costs for private car usage and public transport fares. Thereby, we consider penalties for waiting at transfer points and schedule delay penalties for early and late arrivals at the final destination. We consider the schedule delay structure for commuters as defined previously by Vickrey (1963) and Small (1982) and used in a carpooling framework by de Palma et al. (2022b).

1.2. Contribution and organization of the paper

In this paper, we consider a ride-matching framework with inter and intra-modal transfers. In our framework, riders are able to match with multiple drivers sequentially and drivers are able to match with multiple riders both sequentially and simultaneously if the capacity of their car allows. Riders can transfer between modes or drivers at designated transfer hubs as depicted in Fig. 1. Transfer hubs have connections to public transport services and offer parking opportunities for riders that use their own car to reach the transfer hub. We model the multi-modal ride-matching problem with transfers as a path-based integer programming problem. We develop a heuristic algorithm based on simulated annealing to solve realistic large-scale instances efficiently. We summarize the main contributions of this paper below:

- We develop a framework for multi-modal transport of riders that considers public transport, solo driving, and ride-sharing. The uniqueness of this framework lies in the explicit modeling of the matches between riders and drivers, as well as the scheduling preferences and the transfers between modes and between drivers.
- We derive theoretical optimality conditions for the matching as well as the departure times. These optimality conditions allow for more efficient cost computations and dominance properties, which improve the overall computational performance of the developed methodology.
- We formulate the multi-modal ride-matching problems with multiple transfer hubs as a path-based integer programming problem and we develop a Simulated Annealing (SA) algorithm for large-scale instances. The construction of neighborhood solutions is tailored to the path-based formulation to enhance computational efficiency.

The rest of this paper is organized as follows. Section 2 provides a formulation and a description of the multi-modal ride-matching problem with transfers, which is formulated as an integer linear programming problem. In Section 3 we describe our solution approach. We develop a simulated annealing algorithm that exploits the difference between capacitated and uncapacitated modes to reduce the search space. The results are discussed in Section 4, where we evaluate the performance of our method and the effect of a multi-modal system with transfers on a toy network and a larger realistic case study based on the city of Chicago. The paper is concluded in Section 5. A notational glossary (Table 3), proofs of theorems, remarks, and details on the cost definitions are relegated to the Appendix.

2. Problem description and formulation

In this section, we describe the multi-modal ride-matching problem. We provide the modeling assumptions and the mathematical formulation in Section 2.1. We describe the way the costs and parameters are computed in Section 2.2 and describe the departure time choices of riders and drivers in Section 2.3.

2.1. Assumptions and mathematical formulation

The matching approach is based on a set of predefined rider paths. We consider a static setting in the sense that all drivers and riders submit their trip details in advance. Let $\mathcal{I}$ be the set of riders, $\mathcal{J}$ the set of drivers, and $\mathcal{H}$ the set of transfer hubs. The set of riders can be split into two subsets $\mathcal{I}^c$ and $\mathcal{I}^{nc}$ according to car ownership. Those in $\mathcal{I}^c$ own a car which they may use if they are not matched to a driver, whereas those in $\mathcal{I}^{nc}$ do not own a car and will therefore take public transport if they are not matched to a driver. Although the number of possible matches can get large, it is still polynomial in the number of riders, drivers, and transfer hubs. Note that, if the number of transfers is not limited to one, the number of potential matches would increase exponentially in the number of transfer hubs. Given that the possible number of matches for riders is polynomial, we generate all possible paths in advance. We let the drivers set the departure times, such that the costs of rider paths are independent and such that the cost of a path is independent of the total matching which allows us to determine the costs à-priori. The problem then reduces to selecting the optimal set of rider paths, taking into account that drivers may carry multiple riders both sequentially and simultaneously, as long as their paths are compatible and the capacity of the car is not exceeded.

Relaxing the assumption on a single transfer would lead to an exponential number of paths. This has been considered by Stokkink et al. (2024) for a crowd-shipping system where parcels can be transferred between couriers. A similar approach as the one by Stokkink et al. (2024), using column and row generation, can be taken for the ride-matching problem with transfers. The results of Stokkink et al. (2024) indicate the limited benefits of allowing for more than one transfer. In our case, we deal with people rather than parcels. Drivers and riders are less flexible in terms of additional travel and/or waiting time and their inconvenience perceived by making transfers. Due to these aspects, we have chosen to restrict the number of transfers to a maximum of one transfer per rider. This has also been quantified by Olszewski and Krukowski (2012) and Sil et al. (2022), who study the discomfort of transfers in public transport and inter-modal transport systems, respectively. Masoud and Jayakrishnan (2017b) and Hou et al. (2012) show in a multi-hop ride-sharing framework, that even when allowing for more than one transfer, this is not frequently used.

The model is based on the following set of assumptions:

(A1) Travel times are exogenous and time-independent and there is no congestion.
(A2) Drivers determine their departure time that minimizes their own costs. If matched, a rider must agree to the departure time of the driver.
(A3) The matching is determined by a central operator but drivers and riders only accept the match if their costs are lower than that of their solo-travel alternatives. The objective of the central operator is to minimize the costs of all riders.
(A4) (a) Drivers can only perform a pickup in their departure zone or at a transfer hub and only perform a drop-off at a transfer hub or in their arrival zone.
  (b) Riders can only be picked up in their departure zone or at a transfer hub and only be dropped off at a transfer hub or in their arrival zone.
(A5) Drivers can reach a transfer hub as long as their detour is at most τ time units (riders have no such constraint, as long as their costs are minimized)

Here, Assumption (A1) is required for the formulation to hold, whereas Assumptions (A2)–(A5) can be relaxed without significantly influencing the formulation and the solution framework. For Assumptions (A2)–(A5), we consider that they add efficiency to the system and make the operational platform more appealing and acceptable to riders and drivers. In Assumption (A4) we use the concept of a departure and arrival zone. A departure (arrival) zone is a region within which a driver can pick up (drop off) a rider without a significant detour that is within reasonable proximity from their exact origin (destination). In practice, this can be interpreted as a block of houses, a census tract, or even a community within a city. According to Assumption (A4), drivers only stop outside their departure or arrival zones at transfer hubs.

Every individual has an origin o_i, a destination d_i and a desired arrival time t_i^*. Let $\mathcal{K}$ be the set of rider paths and let binary parameters $e_{ik} = 1$ if rider path k corresponds to rider i, and 0 otherwise. Binary parameters $a_{jk} = 1$ if driver j contributes to rider path k, and 0 otherwise. The cost of rider path k is denoted by c_k. Our model aims to minimize the total costs of riders and does not account for the costs of drivers. Since they determine the departure time, they do not incur any additional scheduling delay costs by sharing a ride. Other than that, drivers are assumed to be fully compensated for the inconvenience of sharing their car with others and the minor detour that may be involved with picking up and dropping off passengers. The design of compensation schemes for drivers is outside the scope of this work. Let decision variable $x_k = 1$ if rider path k is chosen and 0 otherwise. Let q_j be the capacity of the car of driver j, that is, the maximum number of riders driver j is able to transport at the same time. We let the driver only perform pickups at their own origin or the transfer hub and only perform drop-offs at their own destination or the transfer hub. We distinguish between direct trips that take a rider directly from their origin to their destination, and indirect trips that pass through a transfer hub. This means we can identify the following three types of trips, for which the binary parameter a_{jk} is adapted to denote the trip type and the transfer hub that is used.

- Direct trip: $a^0_{jk} = 1$ if driver j contributes to rider path k through a direct trip.
- First leg of indirect trip: $a^{1h}_{jk} = 1$ if driver j contributes to rider path k through a first-leg trip to transfer hub h.
- Second leg of indirect trip: $a^{2h}_{jk} = 1$ if driver j contributes to rider path k through a second-leg trip from transfer hub h.

We use decision variable y_{jh} to define through which transfer hub driver j is going. Similar to ride-sharing paths, public transport paths and paths where riders use their own private car have a corresponding cost c_k. Since no driver is involved in these paths all a^l_{jk} parameters are equal to 0. For a multi-modal path where one leg is a ride-sharing leg, only the corresponding a^l_{jk} is 1, and the others all remain zero. We formulate the matching problem as follows:

$$(\mathbf{P1})\ minimize \sum_{k\in\mathcal{K}} c_k x_k \tag{1a}$$

such that

$$\sum_{k\in\mathcal{K}} e_{ik} x_k = 1 \qquad \forall i \in \mathcal{I} \tag{1b}$$

$$\sum_{k\in\mathcal{K}} a^0_{jk} x_k \le q_j \left(1 - \sum_{h\in\mathcal{H}} y_{jh}\right) \qquad \forall j \in \mathcal{J} \tag{1c}$$

$$\sum_{k\in\mathcal{K}} a^{1h}_{jk} x_k \le q_j y_{jh} \qquad \forall j \in \mathcal{J}, h \in \mathcal{H} \tag{1d}$$

$$\sum_{k\in\mathcal{K}} a^{2h}_{jk} x_k \le q_j y_{jh} \qquad \forall j \in \mathcal{J}, h \in \mathcal{H} \tag{1e}$$

$$\sum_{h\in\mathcal{H}} y_{jh} \le 1 \qquad \forall j \in \mathcal{J} \tag{1f}$$

$$x_k \in \{0,1\} \qquad \forall k \in \mathcal{K} \tag{1g}$$

$$y_{jh} \in \{0,1\} \qquad \forall j \in \mathcal{J}, h \in \mathcal{H} \tag{1h}$$

The objective (1a) is to minimize the cost of all matches. Every rider needs to be matched to exactly one driver, which is enforced by Constraints (1b). Feasibility of the solution from the perspective of a driver is enforced through Constraints (1c)–(1e). Drivers serve as a shared resource among riders, meaning their availability to one rider depends on whether they are already matched with another. This dependency is explicitly accounted for in these constraints. The feasibility of the solution from the perspective of a rider is enforced directly on the set of paths $\mathcal{K}$. That is, the set $\mathcal{K}$ only contains paths that are feasible for a rider. On every leg, a driver $j \in \mathcal{J}$ may have at most q_j riders in his/her car, which is enforced jointly by Constraints (1c), (1d) and (1e). A driver may either serve riders directly from his/her origin to his/her destination or through a transfer hub, but not both. This means that when a driver j makes an indirect trip, he/she can carry q_j riders on the first leg and q_j riders on the second leg. The set of riders on both legs may be partially similar, but it is possible that a driver carries $2q_j$ unique passengers on his/her full trip. Constraints (1f) ensure that a driver only makes a stop at one hub. These constraints also ensure that the first and second legs of a driver are compatible. That is, the first leg ends at the same transfer hub as the second leg starts.

A feasible solution to this problem always exists as long as every rider has access to public transport. As public transport capacity is unlimited, every rider will have a corresponding public transport path. A situation where public transport is unavailable or highly undesirable can be modeled by assigning an arbitrarily high cost to the corresponding path. The solution in which every rider uses public transport as a direct path between their origin and destination is then always feasible. The optimal solution obtained by solving **P1** is not necessarily unique, as multiple solutions may lead to the same objective value.

According to Proposition 1, the set of paths $\mathcal{K}$ can be reduced by removing strictly dominated paths. A dominated path is a path where, by replacing one or multiple legs with a solo leg (either public transport or solo driving), the cost can be reduced. By reducing the number of paths in $\mathcal{K}$, the computation time to solve **P1** can be reduced.

Proposition 1 (*Strictly Dominated Paths*)**.** *Consider a rider $i \in \mathcal{I}$ and a path $k_1 \in \mathcal{K}$. Let path $k_2 \in \mathcal{K}$ of rider i be a copy of path k_1 where one or multiple legs are replaced by direct or indirect legs where the rider travels alone (either by car or public transport). If $c_{k_2} < c_{k_1}$, then path k_1 is strictly dominated by path k_2 and can therefore be omitted from $\mathcal{K}$.*

2.2. Computation of costs and parameter values

In this subsection, we describe the costs and parameters of all rider paths. We distinguish between direct paths and indirect paths that go through a transfer hub. We consider three potential modes for riders: solo driving (SD) for those riders that own a private car, public transport (PT), and ride-sharing (RS). Each mode can be used as a direct path, or a combination of two modes can form an indirect path. Given that solo driving is not possible as a second-leg mode after public transport or ride-sharing (because they left their car at home) we have 7 potential mode choice combinations for indirect paths.

We consider the following cost components and the corresponding parameters. The value of time spent in transport may differ depending on the mode of transport. Therefore, we consider α^{SD}, α^{PT}, and α^{RS}, representing the value of time for solo driving, public transport and ride-sharing, respectively. Every mode also has an associated fixed (access) cost (ψ^{SD}, ψ^{PT}, and ψ^{RS}) and a

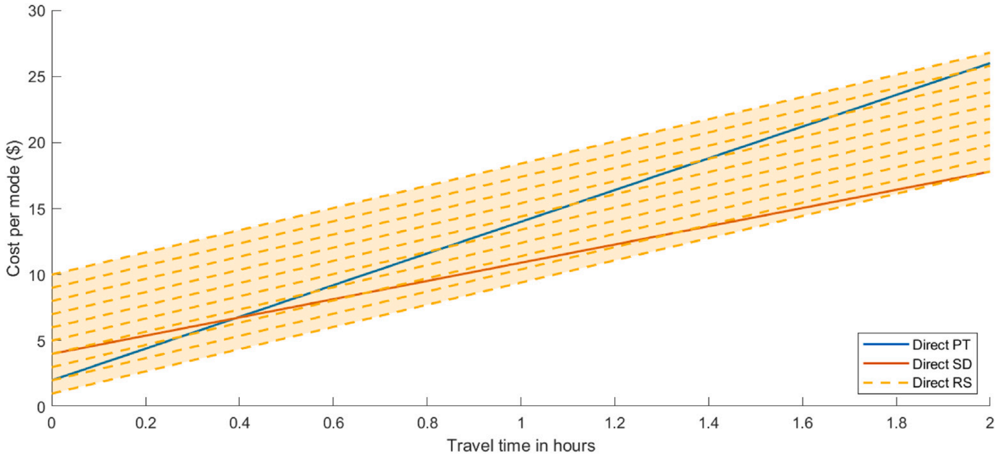

Fig. 2. Example of cost for direct modes (PT = Public transport, SD = Solo Driving) and travel times. Costs of ride-sharing (RS) are given for multiple values of the schedule delay costs, changing the intercept of the function.

variable cost per unit of time spent in that mode (ϕ^{SD}, ϕ^{PT}, and ϕ^{RS}). In the case of indirect modes, the fixed costs are scaled by a factor $f < 1$, reflecting the general reduction in access costs when coordinating transfers. To incorporate schedule delay preferences, we consider β the penalty for every unit of time an individual is early, and γ the penalty for every unit of time an individual is late. Waiting time is penalized by α^{wait}. Travel time between o and d is defined as $tt(o,d)$. We highlight that for the sake of notation, these parameters are all homogeneous. However, the formulation allows for fully heterogeneous parameter values among all individuals. In the latter, we use linear functions of earliness, lateness, and waiting time with respect to time. The computation of the cost provided in Appendix B can be generalized to non-linear cost functions without changing the problem formulation in **P1**. Schedule delay penalties for commuters are incorporated in a way that is consistent with Small (1982) and Arnott et al. (1993) and has been previously used in a ride-sharing framework by de Palma et al. (2022b).

The generalized costs consist of in-vehicle costs depending on the mode, possible physical costs such as fuel or a public transport ticket, waiting penalties, and schedule delay penalties. The exact cost formulation depends on the specific type of path and the departure time choice. When driving themselves, riders can leave at any time t. For the sake of tractability, we consider a set of discrete time intervals $t \in \mathcal{T}$ at which a rider can leave. The optimal departure times of carpooling drivers are also mapped to the closest discrete time interval $t \in \mathcal{T}$. Drivers as well as riders that are using their own cars determine their departure time in advance. A detailed description of departure time choices is provided in Section 2.3. For completeness, the exact cost definitions for each type of path are given in Appendix B.

The generalized costs depend heavily on the chosen mode and the total time spent in that mode. Clearly, the fixed (entry) costs and the variable costs per unit of travel time may differ for each mode. An example of the cost for different direct modes and travel times is given in Fig. 2. The access costs of public transport are relatively low (i.e., the cost associated with reaching the closest bus stop or train station), and costs go up relatively fast given the high value of time of users in public transport and the relatively high fares. In contrast, the utilization of a private car stands in opposition to this. The fixed costs of this are relatively high (i.e., costs for parking) but the costs go up slower because the value of time in a private car is generally lower and fuel costs are typically lower than public transport fares. Based on this, we observe that for shorter trips it is typically favorable to use public transport, whereas for longer trips it is typically favorable to take the car.

In Fig. 2 the fixed costs associated with ride-sharing are those associated with scheduling delay penalties. This may move the yellow curve up and down (indicated by the dotted lines), based on the difference between the desired arrival time of the driver and the rider. The value of time spent ride-sharing is typically somewhat between that of public transport and solo driving. Whether ride-sharing is the cheapest alternative therefore depends on the travel time, as well as how good the match is that can be found in terms of desired arrival time. A sensitivity analysis on the fixed and variable cost components of these modes is performed in Section 4.2.3. Clearly, specific parameter settings may eliminate certain modes completely. For example, if the fixed and variable costs of public transport are both higher than those of solo driving, public transport will not be used. This explains why users with specific features or living in specific areas are unlikely to use certain modes. Given the absence of spatial and socio-demographic data of user preferences in these specific modes, including segmented user groups, is outside the scope of this work, but is marked as an interesting direction of future research. The path-based structure of the mathematical model allows for efficient incorporation of user-specific parameters, without changing the general structure of the model.

For the indirect modes that are considered in this paper, the costs are a linear combination of the direct costs displayed in Fig. 2 depending on the times spent in each mode. For this, we also note that as parking costs can be reduced (or completely free) at park-and-ride spots, this may reduce the fixed part of solo driving when it is combined with another mode.

We note that in the computation of the costs, we only considered the costs of the riders. The reason for this is that due to the inflexibility of the drivers concerning departure time, drivers do not incur any extra scheduling delay costs with a rider. We can neglect the payment of riders to drivers since these are direct money transfers and therefore do not change the solution of the optimization problem. For example, riders that save a percentage of their (expected) costs can share it with the driver. The compensation schemes are outside the scope of this paper but deserve future research attention.

2.3. Departure time choice

In this work, we assume that drivers choose their departure time such that they minimize their own generalized cost. The reason for this is that coordination of departure times in a complex system where riders match with multiple drivers and drivers match with multiple riders is difficult both theoretically and in practice. However, there are some special cases for which the departure times can be determined optimally. In this section, we discuss those special cases and the jointly optimal departure times of matches.

Without transfers, the optimal departure time has a closed form solution. Consider a direct match where a single driver takes a group of riders directly from their origin to their destination. In case lateness is penalized heavier than earliness, the jointly optimal departure time is the minimal departure time of all matched individuals, as formally given in Theorem 1 (the proof is relegated to Appendix C). In this case, everyone is either on-time or early and no one is late. Every rider is matched to at most one driver and therefore the problem can be decomposed over the groups of riders and driver that share a ride altogether. The optimal departure time can be determined independently for every group.

Theorem 1 (*Optimal Departure Time for a Single Leg*). *Let all riders and the driver have identical origins and destinations. Let a driver with desired arrival time t_0^* be matched to N riders with desired arrival times $t_1^* \dots t_N^*$. With $\max(\beta_0, \beta_1, \dots \beta_N) < \min(\gamma_0, \gamma_1, \dots \gamma_N)$, the jointly optimal departure time is equal to $t^o = \min(t_0^*, \dots, t_N^*)$.*

With a transfer, the problem complexifies as more coordination is required. Consider a set of riders with identical destinations d that make a transfer at the transfer hub $h \in \mathcal{H}$. Also, consider a driver with an identical destination as the riders that only performs a ride-sharing trip between the same hub h and destination d. According to Theorem 2, the jointly optimal departure time on the second leg is a function of the arrival time of riders at the second leg, as well as the minimal desired arrival time. According to Theorem 3, the optimal departure on the first leg for a driver that takes one rider on the first leg and another rider on the second leg depends on the desired arrival time of all three individuals. The optimal departure time on the first leg also depends on the rider on the second leg, although they are not directly involved.

Theorem 2 (*Optimal Departure Time for a Second Leg Trip*). *Consider N riders $k_1, \dots k_N$ from origins $o_{k_1}, \dots, o_{k_N}$ who transfer at hub h to their identical destination d, and a driver i from hub h to the same destination d, with their desired arrival times $t_{k_1}^*, \dots t_{k_N}^*$ and t_i^*. Let all individuals have identical cost parameters $\alpha, \alpha^{\text{wait}}, \beta, \gamma$, with $\beta < \gamma$ and $\alpha^{\text{wait}} < \gamma$. We let t_1 be the last departure time for the first leg among all riders and the driver. The joint optimal departure time for the second leg t_2^o is a function of the departure time for the first leg t_1 which is defined as follows:*

$$t_2^o(t_1) = \begin{cases} \max(t_1, \min(t_i^*, t_{k_1}^*, \dots, t_{k_N}^*)) & \text{if } \alpha^{\text{wait}} \leq \beta \\ t_1 & \text{if } \alpha^{\text{wait}} > \beta \end{cases} \tag{2}$$

Theorem 3 (*Optimal Departure Time on First Leg*). *Consider a driver i from o to d passing through hub h, one rider j from o to hub h and one rider k from hub h to d, with their desired arrival times t_i^*, t_j^*, t_k^*. Let all individuals have identical cost parameters $\alpha, \alpha^{\text{wait}}, \beta, \gamma$, with $\beta < \alpha^{\text{wait}} < \gamma$. The joint optimal departure time for the first leg $t_1{}^\circ = \min(t_i^*, t_j^*, t_k^*)$. The joint optimal departure time of the second leg can then be determined according to Theorem 2.*

The results of Theorems 2 and 3 also emphasize the difficulty of coordination in more complex matching systems. If the driver takes another group of riders on his first leg, coordination of departure times with this group influences the departure time of the second group. Similarly, these riders may be matched to a second driver after making a transfer, therefore also influencing his departure time and vice versa. For a large system where many drivers and riders are (indirectly) connected to each other, determining the jointly optimal departure time is a complex problem to solve and difficult to implement both theoretically and in practice. Therefore, in this work, we consider that the driver is in charge of determining the departure times. In current ride-sharing systems such as BlaBlaCar, the driver is also in charge of determining the departure time, and the rider is forced to adapt if they are matched.

3. Solution approach

Although small and medium-scale instances of the problem can be solved with a commercial solver, as the size of the problem increases the number of paths increases and with that, the construction and solving time also increase. To efficiently obtain high-quality solutions we develop a heuristic algorithm based on Simulated Annealing (SA). SA is a probabilistic optimization method that aims to find a global optimum when multiple local optima exists. It uses the concept of temperature, where worse solutions are more likely to be accepted if the tempartume is high. At each iteration, a random neighboring solution is selected. After evaluating its quality, the move is accepted according to a probability depending on the solution quality and current temperature. Accepting worse solutions allows for a more extensive search of the solution space. The temperature slowly decreases to zero as to only accept improvements in the final iterations. The techniques used in this algorithm were independently introduced by various authors, among

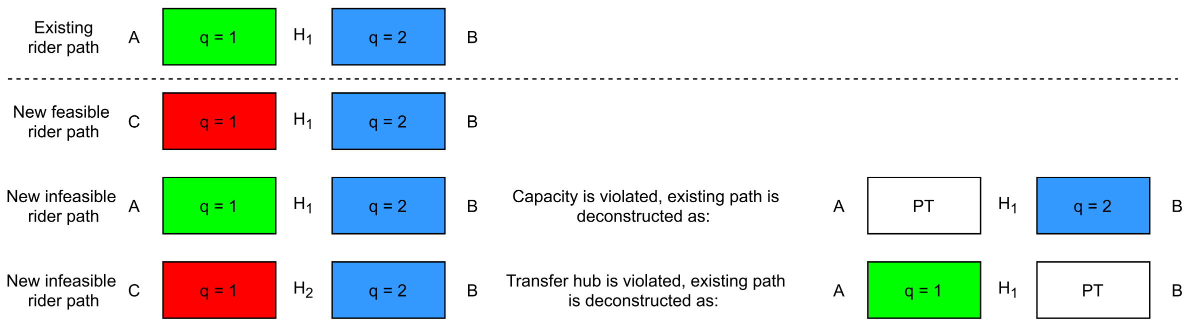

Fig. 3. Example of path deconstruction.

which Pincus (1970) and Kirkpatrick et al. (1983), and was first named Simulated Annealing in Kirkpatrick et al. (1983)[1]. This optimization algorithm has been shown to perform well for matching problems (Sasaki and Hajek, 1988; Bertsimas and Tsitsiklis, 1993) and in ride-sharing frameworks (Jung et al., 2016; Yu et al., 2018).

Our algorithm builds on the difference between capacitated (i.e., ride-sharing) and uncapacitated modes (i.e., public and private transport) in our multi-modal framework. When ride-sharing is involved in a rider path, one rider path can influence the feasibility of another due to the capacity constraints in (1c)–(1f). On the other hand, the use of uncapacitated modes is always allowed and does not influence the feasibility of the solution through the capacity constraints. This allows us to construct a partial solution based only on the ride-sharing legs with the SA algorithm and extend this partial solution by locally optimizing the uncapacitated modes on the empty legs. That is, the empty legs of a partial solution are filled in by either public or private transport. As a consequence, during the execution of the SA algorithm, we only keep track of the transfer hub and driver(s) that every rider uses and we do not need to keep track of the uncapacitated modes explicitly.

The SA algorithm is initialized by a simple construction heuristic where every rider is assigned to their first-best direct path. That is, a rider is matched to the best available driver, takes public transport, or drives solo, depending on which option leads to the lowest cost. The initial solution is referred to as X_0. We keep track of the best solution obtained, which we denote as X_{best}. After initialization, $X_{\text{best}} \leftarrow X_0$. The objective of a solution X is defined as $obj(X)$ which is equivalent to the sum of the costs of all rider paths in X as denoted in Eq. (1a). In every iteration of the algorithm t, a neighborhood solution Y_t is obtained through a set of neighborhood moves. If $obj(Y_t) < obj(X_t)$, the neighborhood solution is an improvement to the current solution and therefore $X_{t+1} \leftarrow Y_t$. If $obj(Y_t) > obj(X_t)$, the neighborhood solution is accepted with a probability $exp\left(\frac{obj(X_t)-obj(Y_t)}{T}\right)$, such that $X_{t+1} \leftarrow Y_t$. If the solution is not accepted $X_{t+1} \leftarrow X_t$. Here, T denotes the temperature. The temperature T is initialized at T_0 and is decreased by a cooling rate ρ in every iteration until the minimum temperature T_f is reached. If the accepted solution X_{t+1} is the best solution that has been obtained so far, we set $X_{\text{best}} \leftarrow X_{t+1}$.

A neighboring solution is obtained by selecting a random rider and finding the best corresponding rider path. Possibly, the addition of this rider path will make other rider paths infeasible. We prioritize the new rider path and deconstruct other rider paths until the solution is feasible. Deconstruction is the removal of a ride-sharing leg that causes the violation of a constraint, which is then filled up by an uncapacitated (i.e., public or private transport) leg to ensure feasibility of the rider path. An example of such a deconstruction is given in Fig. 3. If the new rider path is added that uses the same driver that is already used but does not violate capacity or transfer hub constraints, no deconstruction is needed. If the new rider path is added that violates the capacity constraint of a driver, this driver is removed from the existing rider path. Similarly, if the addition of the new rider path leads to the violation of a transfer hub constraint (i.e., the involved driver has to perform a pickup at H_1 and H_2 simultaneously), the involved driver is removed from the existing rider path. In the existing rider path, the former driver is replaced by a solo leg (either public or private transport). Due to the absence of capacity constraints on solo legs, this replacement always leads to a feasible solution.

For every potential rider path, we obtain the new total cost after adding this path and possibly deconstructing other paths. The best potential rider path is selected as the next neighbor. The objective of the selected neighbor is compared to the previous objective and accepted according to the SA approach.

Identifying violations and deconstructing paths can be computationally costly. To improve the computational efficiency of the algorithm we make two observations: (1) We look for the single best potential rider path as a new neighbor, (2) As the best driver path is added for a specific rider, removing a driver from this path always increases the cost of the solution. That is, deconstruction mostly leads to an increase of the total cost of the solution. Using these observations, if a potential path without considering deconstruction is already outperformed by another potential path, the deconstruction is omitted and the potential path is ignored.

In practice, the algorithm shows to speed up in later iterations. As a better solution is obtained, more drivers are occupied and therefore less feasible alternatives are to be checked for every rider. Due to this, the first iterations are more time-consuming than the last. The algorithm is executed for a fixed number of iterations after which the best solution X_{best} is returned.

[1] The name is derived from annealing in metallurgy and material science. The technique involves controlled heating and cooling of a material to alter its physical properties.

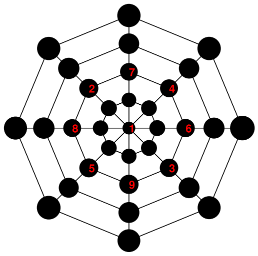

Fig. 4. Circular city with the distribution of origins (left) and destinations (right) and in red the index of the transfer hub. The size of the node shows the density of trips starting (left) and ending (right) in each node.

4. Results

4.1. Case study of toy network

We evaluate our model on a circular city consisting of 33 nodes, as depicted in Fig. 4. Every rider and driver has an origin and destination at one of the 33 nodes. Origins are more likely to be in the suburbs (the outer rings) whereas destinations are more likely to be in the city center. Transfer hubs can be at any of the nodes in the network. In our analysis, we use at most 9 hubs that are always added in the same order. The index of the hub is given in red in Fig. 4. Finding the optimal hubs is an interesting direction of future research, but is outside the scope of this work. Drivers can perform a pick-up or a drop-off at one of the transfer hubs, but only if their shortest path between origin and destination already passes through this hub. Drivers do not make any detours. We consider 500 drivers and 500 riders. Out of those riders, 75% own a car which they may use to drive themselves. Desired arrival times are drawn from a truncated normal distribution with a mean at 8:00 and a standard deviation of 1 h. The distribution is truncated such that we only allow desired arrival times between 7:00 and 9:00. Throughout this section, an indirect path with a transfer from mode m_1 to mode m_2 is denoted as $m_1 \rightarrow m_2$, for ease of notation.

The parameter settings are homogeneous among the entire population and are defined as follows. The value of time spent in a car α^{car} is equal to 6.4 [\$/h]. The value of time in public transport α^{pt} is higher and is set equal to 12.0 [\$/h]. In addition to this, public transport has a fixed cost ϕ^{pt} of 2.0 per trip. Earliness and lateness are penalized with β and γ equal to 3.9 [\$/h] and 15.21 [\$/h] respectively, independently of the mode of transport. Waiting time is penalized by α^{wait} which is equal to 13.5 [\$/h] such that $\beta < \alpha^{\text{car}} < \alpha^{\text{pt}} < \alpha^{\text{wait}} < \gamma$, consistent with the literature (Small, 1982). Fuel costs ϕ^{fuel} are equal to 4 [\$/h] and parking costs ϕ^{park} are equal to 1.5\$. The percentage of riders owning a car is set to 75%. In this case study, drivers only pass through transfer hubs if it does not increase their travel time (i.e., $\tau = 0$).

A lower bound on the average cost per rider in the discrete formulation is found when all riders find their perfect match (i.e., when origins, destinations, and desired arrival times of drivers and riders are identical). In this case, they only incur travel costs α^{car}. Given an average commuting time of 1 h and 15 min in the synthetic data, the lower bound is $\frac{75}{60}\alpha^{\text{car}} = 8\$$. On the other hand, an upper bound is found when all riders use public transport. In this case, they all incur travel costs α^{pt} and the fixed cost ϕ^{pt}. Given an average commuting time of 1 h and 15 min in the synthetic data, the upper bound is $\frac{75}{60}\alpha^{\text{pt}} + \phi^{\text{pt}} = 17\$$. We note that stronger bounds can be found by incorporating the portion of riders owning a car, by incorporating their private transport alternative.

All integer programming problems are implemented in Java with CPLEX version 12.6.3.0. All problems are solved to optimality and can be solved within a matter of seconds or minutes, depending on the exact problem configurations.

4.1.1. Influence of transfers on modal split, costs, and VHT

We consider the influence transfers make on the modal split, the average cost per individual, and the total Vehicle Hours Traveled (VHT). We vary the number of transfer hubs in the system between 0 and 9. The results are an average of 10 randomly simulated instances. The results are displayed in Fig. 5 where 5a displays the modal split of riders, 5b displays the average cost of riders, and 5c displays the VHT as a percentage of the VHT when ride-sharing is not available. Clearly, when there are no transfer hubs, the only possible mode choices are direct ride-sharing, solo driving, and public transport. By opening transfer hubs, a modal shift to the other modes is observed. Especially the number of riders ride-sharing on two separate legs and the number of riders using their own car on the first leg and ride-sharing on the second leg increases drastically. The reason for this is that by using a transfer, more options exist for matching to a driver with the same destination and a similar desired arrival time, at the cost of waiting at the transfer hub. The number of direct matches may be limited as the origin and destination of the rider and driver need to be identical and the desired arrival time needs to be relatively similar. Fig. 5c displays that the total VHT by riders in their private car significantly decreases by 30% when allowing transfers, which has a direct influence on emissions.

Fig. 5b displays how the costs change by opening transfer hubs. By using a single transfer hub in the center of the network, the average cost decreases from 12.80\$ to 12.10\$. Increasing the number of transfer hubs allows for a further decrease in the average cost, but not nearly as substantial as for the first hub in the center. When all 9 hubs are opened, the average cost decreases to 11.70\$. To put these numbers into the right perspective, we compare them to the lower and upper bounds defined in Section 4.1. The upper bound is strengthened by using the portion of riders that own a car. The upper bound is 14.00\$ and the lower bound is

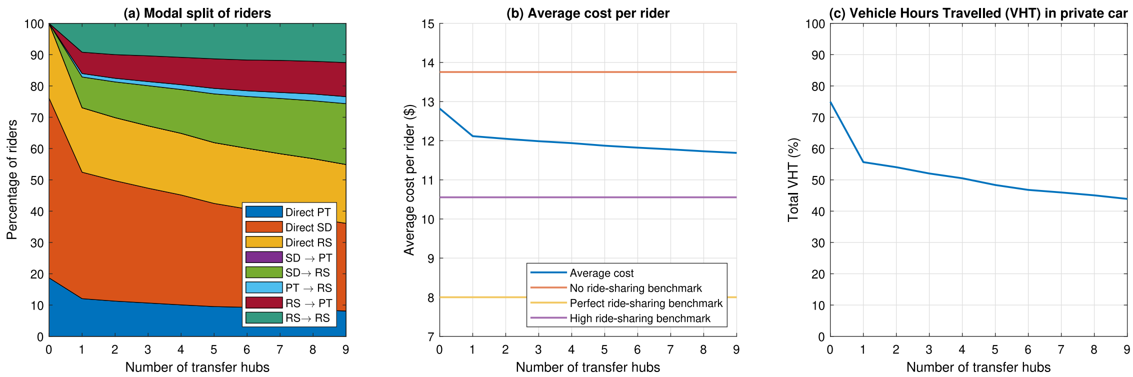

Fig. 5. Statistics for a varying number of hubs. VHT is given as a percentage of the VHT when ride-sharing is not available as a mode. SD = Solo Drive, PT = Public Transport, RS = Ride-Share.

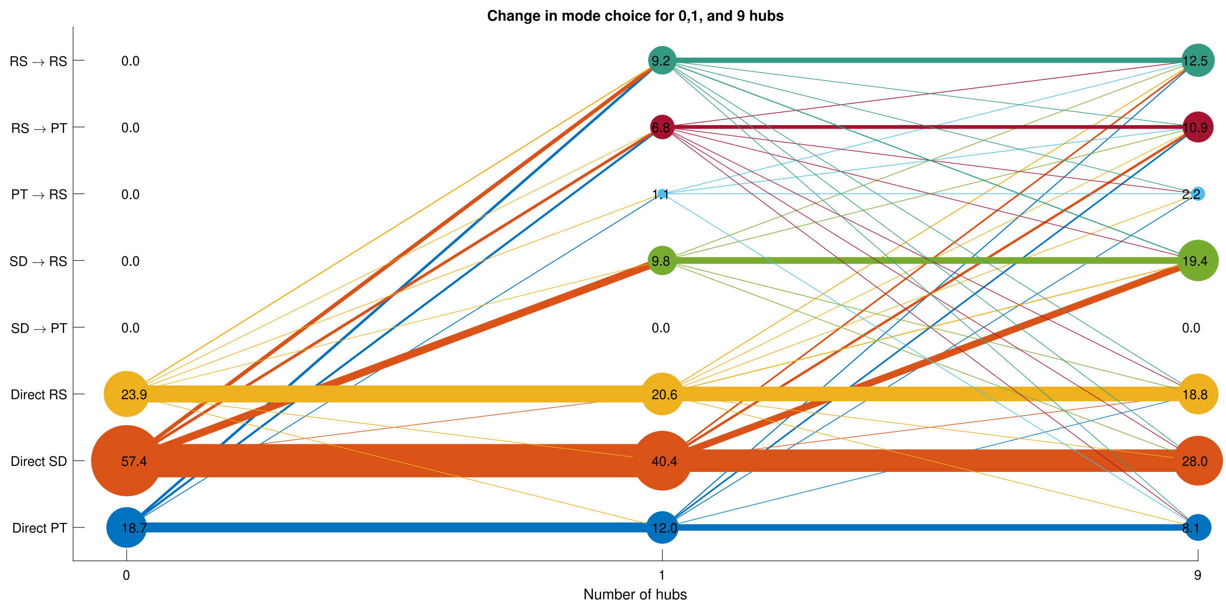

Fig. 6. Change in mode choice for 1, 2, and 9 transfer hubs. The vertical axis displays the modes, horizontal axis displays the number of hubs. The size of a bubble (and the number inside that bubble) depicts the number of riders using that mode and the thickness of the lines depicts how many riders change from one mode to another when the number of hubs changes. SD = Solo Drive, PT = Public Transport, RS = Ride-Share.

8.00$. We see that when using 9 hubs, the improvement from the no ride-sharing upper bound is doubled compared to when zero hubs are used. Thereby, the objective is almost 20% closer to the lower bound of the cost compared to when zero hubs are used. We emphasize that this lower bound is only attained if every rider can find a perfect match. Therefore, attaining this lower bound is highly unlikely in realistic scenarios where the number of drivers is not infinitely large. For example, when the number of drivers is 2500 (5 drivers for every rider) the costs only decrease to 10.60$ (the purple line in Fig. 5b).

Note that as the number of private vehicles used decreases, it is expected to have a further decrease in travel times due to a decrease in congestion. We do not include this effect in our analysis as travel times are exogenous, but in reality, the system could create even higher social benefits.

Fig. 6 displays how the mode choice changes when the number of hubs changes. Although the majority of the mode choices remain the same, some significant movements can be observed. For example, riders that drove their own car without transfer hubs mostly change to ride-share on both legs or to use their own car on the first leg and ride-share on the second leg. Riders that used to take public transport without transfer hubs, change either to ride-share on both legs or on a single leg while using public transport on the other. Former ride-sharers may change to any of the modes, abandoning their direct ride. As an effect of these changes, we also observe some riders that used their own car or public transport move towards a direct ride-share and vice-versa.

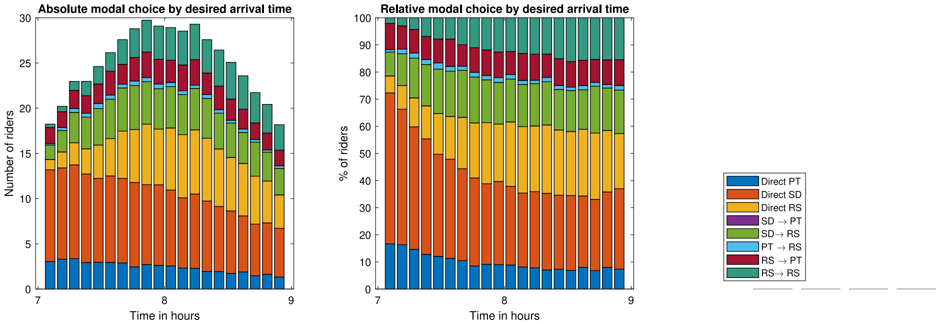

Fig. 7. Distribution of riders by desired arrival time and mode. SD = Solo Drive, PT = Public Transport, RS = Ride-Share.

4.1.2. Spatio-temporal distribution of riders

In this section, we evaluate the spatial–temporal distribution of riders. First, we classify riders by their desired arrival time and the mode they use to commute. The results are displayed in Fig. 7 where the left-hand panel displays the number of riders using every mode and the right-hand panel displays the proportion of riders using every mode (i.e., scaled by the number of riders with that desired arrival time). To obtain these results, 100 simulated instances have been used with 1 hub in the center and 4 on the second ring road (identified by 1–5 in Fig. 4).

It is clear that the proportion of riders traveling solo is the highest in the tails. The reason for this is that the number of potential matches with identical origins and destinations and similar desired arrival times is low since the number of individuals here is rather low. This effect is more apparent for riders with an early desired arrival time. When these riders match to a driver, it is highly likely that the desired arrival time of the driver is later than that of the rider, and therefore the rider will suffer from lateness. As lateness is penalized heavier than earliness, the effect is more apparent at the start of the morning commute than it is at the end. As the value of β approaches the value of γ the distribution gets more symmetric. At the peak of the rush hour (i.e., around 8:00 when most commuters have their desired arrival time), the number of ride-sharers is the highest. We see a skewness towards later desired arrival times, which follows the same reasoning as stated before. By changing the number of hubs, the modal share of each mode changes as described in Section 4.2.1. The shape of the distribution on the other hand stays roughly the same while being shifted either up or down depending on the mode.

For a more detailed analysis of ride-sharing with transfers, we look at the portion of riders that share a ride with a transfer distributed by origin and destination. We consider a network with a single hub in the center of the network. The results are analyzed in more detail by disaggregating over both origin and destination. Given that the network is symmetric in all interior roads, we only distinguish between the four rings, but not the nodes on the ring. That is, the network can be rotated without changing the distribution. The results are shown in Fig. 8. Riders that share a ride at a transfer generally have an origin at one side of the transfer hub and a destination on the opposite side, approximately. The reason for this is that the detour imposed by the transfer hub is relatively small for those origin–destination combinations. Furthermore, we observe there is a higher concentration of origins and destinations closer to the center.

4.1.3. Heuristic versus exact approach

We evaluate the performance of the SA heuristic on the circular city case study, for which the exact solution can be obtained by CPLEX within a reasonable amount of time. We evaluate the performance based on the optimality gap (i.e., the difference between the objective value obtained by the heuristic and that obtained by the exact MILP solver), the computation time of both methods, and the number of variables in the MILP formulation. For the SA heuristic, the parameters are tuned to $T_0 = 3$, $T_f = 0.001$, $\rho = 0.999$ and 50,000 iterations are used. The results are displayed in Table 1 for various numbers of drivers and riders, capacity of drivers, and number of hubs.

When the number of hubs is equal to zero ($|H| = 0$) the problem reduces to a simple matching algorithm of matching riders to drivers, with public and private transport alternatives. In this case, we see that the heuristic obtains an optimality gap below 0.2% and even obtains the optimal solution in some cases. The optimality gap increases with the number of hubs, as more optimization and coordination between drivers and riders is required to obtain the optimal solution. However, for 5 hubs the maximum optimality gap is still 3.6% and on average 2.7%.

In return for a small optimality gap, the computation time is reduced significantly. The solution time of CPLEX is between 5 and 30 times higher than that of the SA heuristic, even on small instances. We note that for larger instances, such as those discussed in the next subsections, CPLEX cannot find a feasible solution or is not even able to construct the formulation due to the size of the instance. In general, we do not observe an effect of the number of riders and drivers on the computation time of the heuristic (as opposed to the computation time of the CPLEX solver, which increases substantially). We also observe an increase in the number of variables (paths) that are needed for the mathematical formulation. The number of variables increases with the number of hubs and the number of drivers. This implies that for larger instances, such as those considered in the next sections, the exact formulation runs into memory issues which forbid CPLEX to construct the problem.

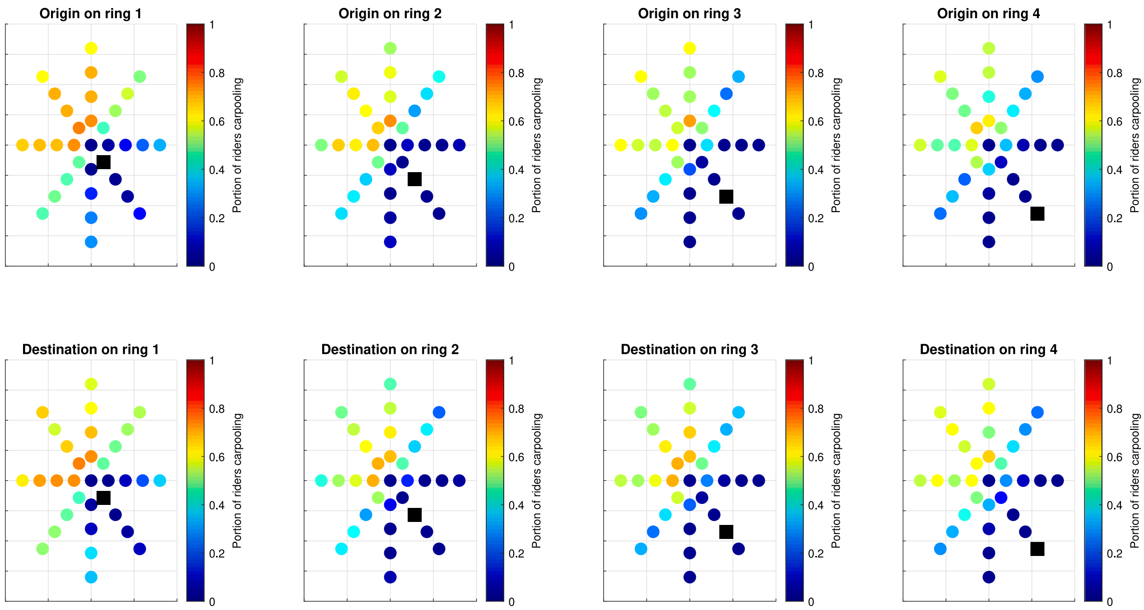

Fig. 8. Proportion of riders sharing a ride with a transfer for a network with a single transfer hub in the center. The destination and origin, for the top and bottom respectively, are marked by a black square.

Table 1
Comparison of optimality gap and computation time of heuristic versus exact solver.

$\|I\| = \|J\|$	max q	$\|H\|$	Optimality gap (%)	Time CPLEX (s)	Time SA (s)	Variables CPLEX
500	1	0	0.00	2.6	0.4	1 052
500	1	1	1.23	3.1	0.5	7 238
500	1	5	1.70	2.9	0.6	10 016
500	2	0	0.00	2.5	0.2	1 052
500	2	1	1.76	2.2	0.4	7 238
500	2	5	2.26	2.7	0.4	10 016
500	3	0	0.00	2.2	0.3	1 052
500	3	1	1.70	2.4	0.3	7 238
500	3	5	2.67	2.7	0.2	10 016
1000	1	0	0.08	9.7	1.0	2 647
1000	1	1	2.53	11.6	2.0	37 317
1000	1	5	3.27	13.7	2.2	51 710
1000	2	0	0.21	8.9	0.3	2 647
1000	2	1	3.04	10.8	0.8	37 317
1000	2	5	3.61	13.6	1.0	51 710
1000	3	0	0.09	9.3	0.3	2 647
1000	3	1	2.49	11.1	0.6	37 317
1000	3	5	3.12	13.0	0.8	51 710

Note: The first three columns denote the instance settings: the number of riders and drivers, the maximum capacity, and the number of transfer hubs, respectively. The fourth column denotes the optimality gap between the CPLEX solver and the SA heuristic. The fifth and sixth column denote the solution time of both approaches and the seventh column denotes the number of variables created for the CPLEX model.

4.2. *Case study of Chicago, USA*

To obtain results on a more realistic network, we use a case study of the city of Chicago, USA. Chicago is one of the biggest cities in the USA, with an area of more than 600 km^2 (230 sq mi) and a population of more than 2.7 million. We use data provided by the City of Chicago (2010) to establish 77 nodes based on the communities in the cities. We consider a fully connected graph, where the distances between nodes are obtained as Euclidian distances between the geographical centers of the communities. The origin–destination data of commuters is based on the use of ride-hailing vehicles, provided by City of Chicago (2022). The historical dataset has been used to construct demand rates of origin–destination pairs. In turn, this has been used to randomly generate instances of commuters. Commuters are generated according to a Poisson process with the obtained demand rate. The distribution of origins and destinations is displayed in Fig. 9.

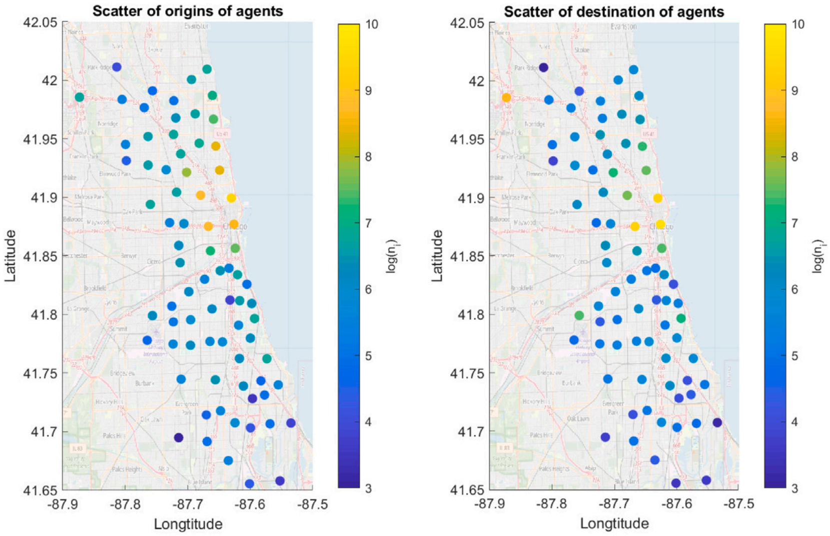

Fig. 9. Distribution of origins and destinations in Chicago (color depicts the number of commuters in log scale).

We use a maximum of five transfer hubs. Nodes that are selected as transfer hubs are identified by the Chicago public transport system as "Park-and-Ride" options. As these locations naturally have parking spaces and a connection to public transport, they meet the physical requirements for a transfer hub.

The parameter values used in this case study are similar to those described in Section 4.1, but tuned to fit this specific case study. The value of time spent in a car α^{SD} is equal to 6.4 [$/h]. The value of time spent ride-sharing, α^{RS}, is slightly higher and is set at 8.9 [$/h]. The value of time in public transport α^{PT} is higher and is set equal to 10.0 [$/h]. In addition to this, public transport has a fixed cost ϕ^{pt} of 2.0 per trip. Earliness and lateness are penalized with β and γ equal to 4.9 [$/h] and 15.21 [$/h] respectively, independently of the mode of transport. Waiting time is penalized by α^{wait} which is equal to 13.5 [$/h]. The percentage of riders owning a car is set to 75%. In this case study, drivers detour at most 10 min (τ = 10 min) to pass through transfer hubs. The fixed and variable costs of each mode are defined as follows. For public transport, the fixed cost ψ^{PT} is 2.0$ and the variable cost ϕ^{PT} is 2.0 [$/h]. For solo driving in a private car, the fixed cost ψ^{SD} is 3.5$ and the variable cost ϕ^{SD} is 0.5 [$/h]. For carpooling, the fixed cost ψ^{RS} is 2.5, while the variable cost ϕ^{RS} is negligible and the multiplication factor $f = 0.75$. Although all these parameters are chosen to give a reasonable representation of a real city, we note that many of these parameters depend on the exact case study. Therefore, a sensitivity analysis on some of the most influential parameters is performed in Section 4.2.3.

The SA heuristic is used to solve the instances. For this, the parameters are tuned to $T_0 = 3$, $T_f = 0.001$, $\rho = 0.9999$ and 100,000 iterations are used.

4.2.1. Influence of transfers on modal split, costs, and VHT

Similar to the analysis we performed for the circular city network, we vary the number of transfer hubs and analyze the influence of transfers on the modal split, the average cost of travelers, and the total vehicle hours traveled by private car in the network. The results are displayed in Fig. 10 where transfer hubs are added in an arbitrary order. In these experiments, approximately 5500 riders and 5500 drivers are used. In our framework, riders may change modes and departure times, whereas the behavior of drivers remains almost unchanged. The only change in driver behavior is that they may make a small detour of τ = 10 min to reach a transfer hub. Our experiments show that, on average, the detour per driver is not substantial. Given the small number of drivers carrying passengers that transfer, and the low maximum tour value of 10 min, average trip durations of drivers do not increase significantly across the tested configurations.

When no transfers are allowed, the majority of the commuters use public transport for their commute. This is caused by the share of users that do not own a car and by the relatively high portion of short-distance trips for which public transport is generally cheaper, as shown in Fig. 2. Less than 10% of all riders share a ride with a driver during their commute. The average cost for commuters is 5.80$ and the total VHT in a private car is equal to 1100 h.

When the number of transfer hubs increases, we see a significant modal shift to the $SD \rightarrow RS$ mode. This means that a large share of the commuters use their car to reach a transfer hub and travel with another commuter from there. Another smaller share of commuters uses ride-sharing on both legs but use a different driver on each of these legs. By using transfers, the average costs of riders are reduced from 5.80$ to 5.30%. The total VHT in a private car is also reduced significantly from 1100 to approximately 800 h. However, as the number of transfer hubs is increased further the total VHT increases again to 900 h. This shows that because the majority of commuters who use the transfer hub take their private car to reach the hub, an increase in ride-sharing is not guaranteed to reduce the total VHT. It is important to note here that because riders only use their car during the first leg of the trip,

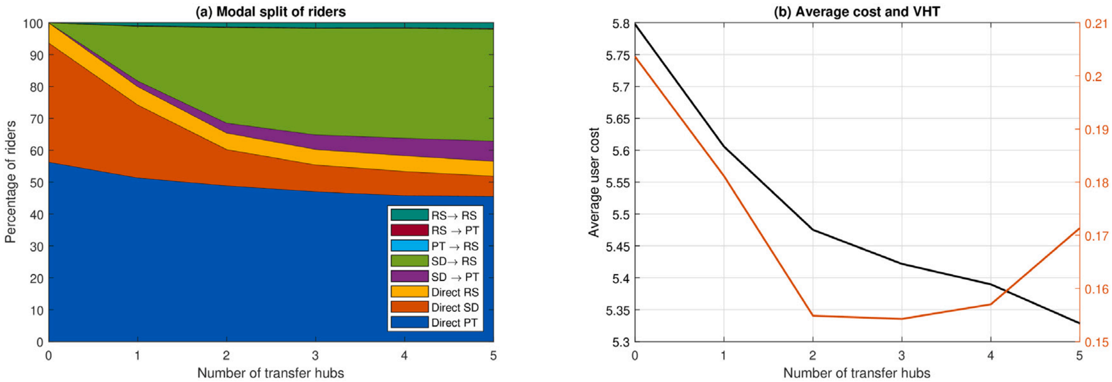

Fig. 10. Modal split, average cost, and VHT in private car for a varying number of transfer hubs in the Chicago network.

which is typically in a suburban area, the number of vehicles in the city center is also reduced substantially. Since the city center is usually the most problematic in terms of congestion, this forms a substantial advantage.

The results suggest that a good pricing/subsidy scheme for ride-sharing needs to be developed. Ride-sharing needs to be targeted at those commuters that would abandon their car. In our current analysis, we see that this is successful, but only up to a certain extent. We also see that public transport usage is reduced by 10%, which is generally undesirable. Designing a compensation scheme for ride-sharing with and without transfers is outside the scope of this work, but is marked as an important direction of further research. Similar schemes for ride-hailing systems to fight these negative effects have been studied by Hryhoryeva and Leclercq (2024).

4.2.2. Network analysis

To better understand the influence of geographical properties on mode choices, we perform a network analysis of the results. Fig. 11 displays the portion of commuters who use a transfer on their commute according to their origins and destinations (left and middle panel) and the number of commuters that pass through the five transfer hubs (right panel). The results indicate that the transfer hub closest to the city center (Lat 41.84, Lon −87.65) is the most used. The reason for this is that driving alone is generally cheaper, but ride-sharing excludes the use of expensive parking costs. Therefore, the closer to the CBD a commuter can transfer, the better off they usually are. This result may have implications for parking options. Parking in the city center is usually more restricted. Therefore, the price of parking at a transfer hub in the city center may be more expensive in reality than the price of parking outside the city center.

We observe that commuters who live and/or work in the city center are typically less inclined to use a transfer. The reason for this is that for their short trips, public transport is usually a cheap alternative and there are abundant direct ride-sharing options available. For origins and destinations further from the city center, more commuters tend to use a transfer during their commute. An exception to this can be found for commuters who have a destination in the far south of the network. The reason for this is that, because of the distribution of riders and drivers in the network, finding a match with a destination in that area is rather difficult. In this case, commuters are better off traveling the complete trip alone.

To support these results, we evaluate the correlation between trip distance and mode choice in Figs. 12 and 13. Fig. 12 considers the original distribution of trips, whereas 13 considers a distribution that is biased towards medium and long-distance trips. For the biased distribution, the frequency of trips shorter than 6 min is reduced by 35% and the frequency of trips longer than 60 min is increased by 100%. The first panel displays the distribution of the trip length, measured as the minimum number of hours required to get from origin to destination. The second panel displays a boxplot of the mode choice in the absence of transfers and the third panel displays the distribution when all five transfer hubs are used.

The results of the second panel in Fig. 12 confirm the theory described in Section 2.2. Long-distance trips are mainly performed by car (except for those commuters who do not own a car and who have to resort to ride-sharing or public transport), whereas short-distance trips are mostly performed by public transport. Ride-sharing is mostly used for medium-distance trips. This is even more apparent when more medium and long-distance trips are generated in Fig. 13.

When five transfer hubs are used, we observe that commuters with a large trip length shift to modes that use a transfer. The longest trips are those that connect ride-sharing and public transport, whereas the shortest trips are still direct solo driving trips. For the shortest trip, however, the sample size is too small to draw solid conclusions. The biased distribution mainly reduces PT usage and increases ride-sharing with and without a transfer.

4.2.3. Sensitivity analysis

To evaluate the effect of different system configurations and parameter settings, we perform a detailed sensitivity analysis with respect to the cost of solo driving (ψ^{SD}) and public transport usage (ψ^{PT} and ϕ^{PT}). For every scenario we evaluate two important system performance indicators: the average vehicle hours traveled per rider in every mode and the percentage of riders that make a

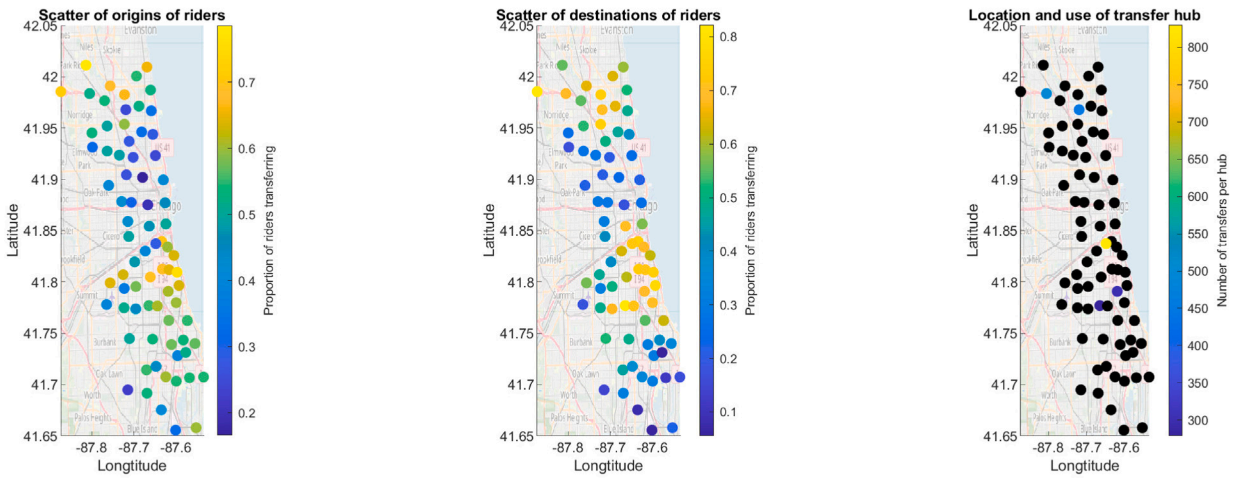

Fig. 11. Network analysis.

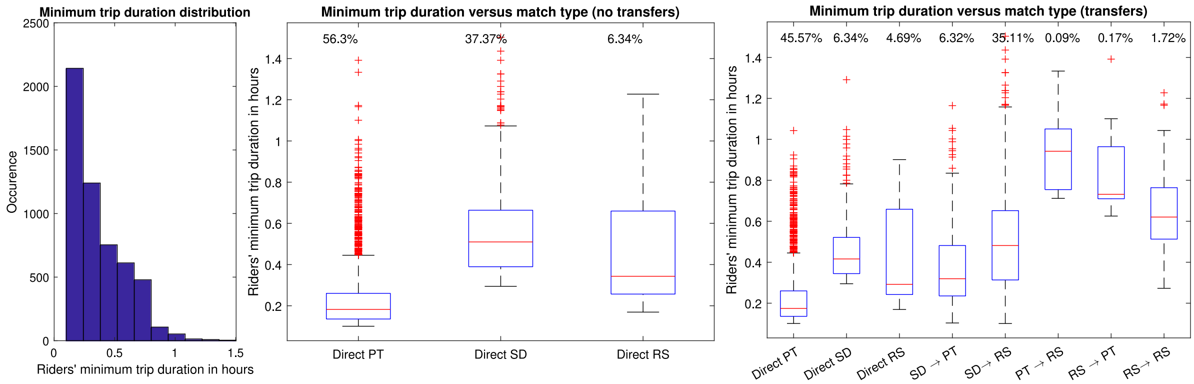

Fig. 12. Box plot of trip length versus mode choice with original distribution.

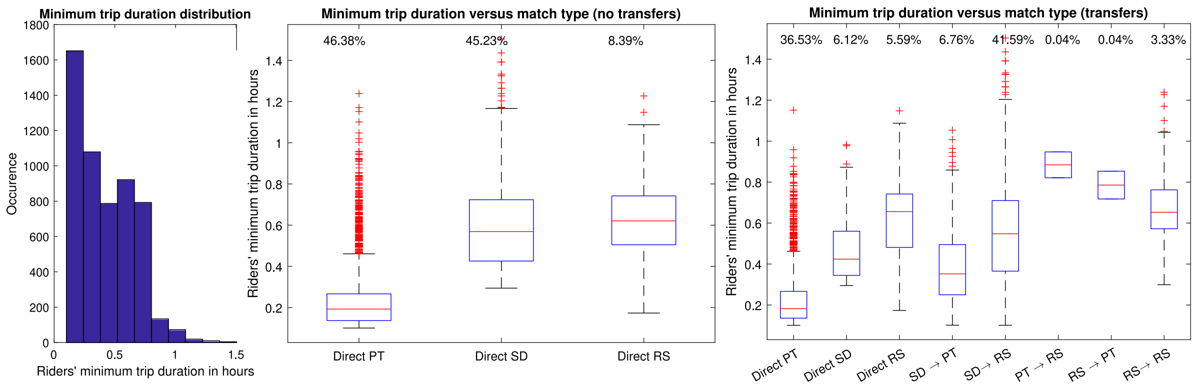

Fig. 13. Box plot of trip length versus mode choice with biased distribution.

transfer on their trip. We evaluate the effect of varying fixed costs of car usage in Fig. 14. This can be influenced by varying parking fees, for example. In Fig. 15 we evaluate the effect of the fixed and variable costs of public transport. By considering these different configurations, we aim to indicate how the results may change across cities with varying road and public transport networks and the associated costs.

The sensitivity analysis in Fig. 14 shows that for ψ^{SD} between \$1.5 and \$3.0 the VHT and percentage of riders making a transfer remains unchanged, while it significantly increases for higher values of ψ^{SD}. The reason for this is that for riders owning a car,

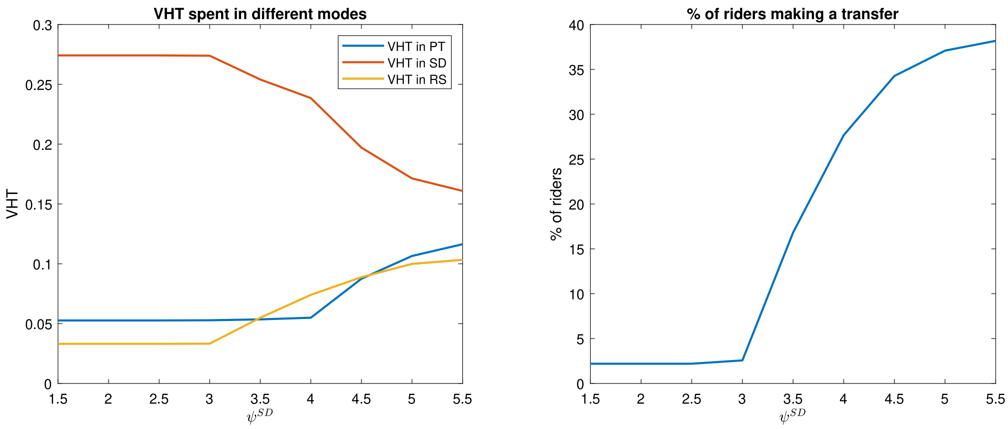

Fig. 14. Sensitivity analysis of VHT and transfers to ψ^{SD}.

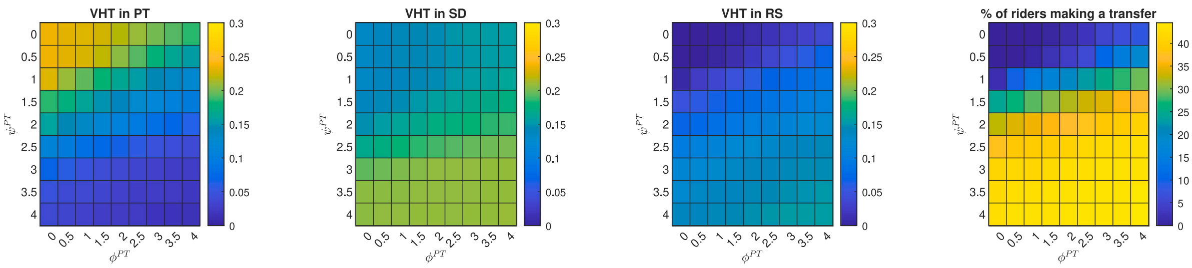

Fig. 15. Sensitivity analysis of VHT and transfers to ψ^{PT} and ϕ^{PT}.

Table 2
Sensitivity analysis on vehicle capacity.

Capacity (# riders)	Average cost	VHT in PT	VHT in SD	VHT in RS	% of riders transfer
1	5.33	0.11	0.17	0.10	37.09
2	5.28	0.10	0.17	0.11	42.35
3	5.27	0.09	0.17	0.12	43.39

using their car is always the dominant option when fixed costs are that low. When the fixed costs increase even further, the use of the private car decreases substantially and is replaced by ride-sharing first and also by public transport after. It also constitutes to a significant increase in the number of transfers. Similar conclusions can be drawn from the heat maps in Fig. 15. An increase in the costs of public transport (both fixed and variable) decreases public transport ridership. At the same time, we see higher VHTs for solo driving and ride-sharing and an increase in the percentage of riders that make a transfer.

Finally, we perform a sensitivity analysis on the capacity of vehicles. We test capacities of 1, 2 and 3 riders per vehicle, for which the results are displayed in Table 2. As the capacity of vehicles increases, ride-sharing becomes more attractive and the percentage of riders making a transfer on their way also increases. However, the increase in VHT for ride-sharing is countered by a decrease in VHT for public transport, while the VHT in private vehicles remains constant. This highlights the adverse effects of promoting ride-sharing, as the new users are former public transport users, rather than former car users.

5. Conclusion

In this paper, we introduced the multi-modal ride-matching problem with transfers. Ride-sharing, public transport, and private cars were modeled as complementary first- or last-mile modes, as well as competitive modes. Riders can change between two modes as well as between two drivers at designated transfer hubs. These transfer hubs have connections to public transport and have sufficient parking spaces for those reaching the transfer hub with their private car.

The problem is modeled as a path-based integer programming problem. All feasible paths for each rider and their corresponding costs are determined in advance. The selection of paths is optimized to minimize costs while ensuring feasible matches for drivers. The size of the problem can be drastically reduced by removing strictly dominated paths. Thereby, the computation of costs can be enhanced by considering optimality conditions on departure times of drivers and riders, which can then also help to reduce the size of the problem.

We developed an SA heuristic to solve realistic large-scale instances to near-optimality. The heuristic constructs new paths to improve upon the current solution and deconstructs existing paths to obtain feasible solutions. The computational speed of the heuristic is significantly improved by observing suboptimalities in an early stage to avoid costly computations. The heuristic reduces the computation time by 80 to 95% and attains an optimality gap of less than 3.3% and on average 2.3% on small-scale instances. Furthermore, it allows us to solve realistic instances within a reasonable amount of time.

The performance of the designed system and algorithm are evaluated in two case studies. The first is a small-scale case study of a toy circular city network, for which the problem can be solved to optimality within a reasonable amount of time. The second case study is more realistic and is based on the city of Chicago. The results show that with a limited number of transfer hubs, both the average cost per rider and the vehicle hours traveled can be reduced by more than 20%. When desired arrival times are drawn from a truncated normal distribution, ride-sharing is mostly observed around the center of the considered time window whereas private transportation is more popular in the tails of the time window, when fewer potential matches are available. The modal split is highly dependent on the difference between the fixed and variable costs of modes, as indicated by our sensitivity analysis. In addition to this, trip length and the location of the origin and destination of riders plays a big part in the mode choice.

Contrary to previous studies, our results show that ride-sharing does not only attract riders who were previously using public transport, but it also reduces private car usage by 20%. Without transfers, public transport is mostly used for short-distance trips whereas long-distance trips are performed using a personal vehicle. When allowing for transfers, we observe that a large portion of these long-distance trips are replaced by inter-modal trips that combine ride-sharing with public or private transport. Nevertheless, an increase in vehicle capacity shows that new ride-sharing users are drawn from public transport, rather than private transport. Due to the various simplifying assumptions on the decision-making of drivers and riders, further research is necessary to accurately quantify the true modal shift.

Further research is needed to incorporate travel-time uncertainty in the problem. Uncertainties may influence the cost associated to a match, but also the feasibility of a match. Through this, exogenous or endogenous congestion can be incorporated into the framework. Thereby, future research may focus on jointly modeling the morning and evening commute and the influence this has on ride-matching solutions. Other interesting research directions are the design of optimal transfer hub locations and the optimal design of incentives for drivers and riders.

CRediT authorship contribution statement

Patrick Stokkink: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization. **André de Palma:** Writing – original draft, Supervision, Methodology, Investigation, Formal analysis, Conceptualization. **Nikolas Geroliminis:** Writing – original draft, Supervision, Methodology, Investigation, Formal analysis, Conceptualization.

Acknowledgments

André de Palma is grateful to HARMONIC (ANR-PEPR, decarbonization project, France) and AFFINITE (Projet-ANR-20-CE22-0014) projects for their financial support. The authors sincerely thank the editor and referees for their valuable and insightful comments.

Appendix A. Notational glossary

See Table 3.

Appendix B. Computation of costs

In this appendix, we describe the costs for every type of path. We consider separately all 10 types of paths (3 direct and 7 indirect). For the sake of notation, we define $t^*_j(h)$ as the desired arrival time of driver j at transfer hub h if he travels through that hub. This is simply computed as $t^*_j(h) = t^*_j - tt(h, d_j)$. The three direct paths are described below:

Direct PT

Every rider $i \in \mathcal{I}$ has the option to take public transport. Public transport has a fixed cost plus a variable term per unit of time traveled. A rider $i \in \mathcal{I}$ that takes public transport incurs a cost:

$$c_k = (\alpha^{PT} + \phi^{PT})tt(o_i, d_i) + \psi^{PT} \tag{3}$$

Direct SD

Every rider that owns a car $i \in \mathcal{I}^c$ also has the option to drive from origin to destination directly. In that case, on top of his value of time, drivers pay for fuel consumption and parking at the destination. Departure time choices are made to minimize the costs. In the case of deterministic travel times, this means that they arrive exactly at their desired arrival time, and as such schedule delay penalties are zero. The costs for such a path are defined as:

$$c_k = (\alpha^{SD} + \phi^{SD})tt(o_i, d_i) + \psi^{SD} \tag{4}$$

Table 3
Notational glossary.

Sets	
$\mathcal{H}$	Set of transfer hubs (indexed h)
$\mathcal{I}$	Set of riders (indexed i)
$\mathcal{I}^c$	Set of riders owning a car
$\mathcal{I}^{nc}$	Set of riders not owning a car
$\mathcal{J}$	Set of drivers (indexed j)
$\mathcal{K}$	Set of rider paths (indexed k)
$\mathcal{T}$	Set of discrete time intervals (indexed t)
Parameters	
a^0_{jk}	Binary parameter indicating if driver $j \in \mathcal{J}$ contributes to rider path $k \in \mathcal{K}$ through a direct trip
a^{1h}_{jk}	Binary parameter indicating if driver $j \in \mathcal{J}$ contributes to rider path $k \in \mathcal{K}$ through a first-leg trip to transfer hub $h \in \mathcal{H}$
a^{2h}_{jk}	Binary parameter indicating if driver $j \in \mathcal{J}$ contributes to rider path $k \in \mathcal{K}$ through a second-leg trip from transfer hub $h \in \mathcal{H}$
c_k	Generalized cost of rider path $k \in \mathcal{K}$
d_i	Destination of individual $i \in \mathcal{I} \cup \mathcal{J}$
e_{ik}	Binary parameter indicating if rider path $k \in \mathcal{K}$ corresponds to rider $i \in \mathcal{I}$
f	Factor to multiply the fixed cost of a mode when it is used during an indirect trip. Typically, $f < 1$
q_j	Capacity of the car of driver $j \in \mathcal{J}$
o_i	Origin of individual $i \in \mathcal{I} \cup \mathcal{J}$
$tt(\cdot,\cdot)$	travel time between two nodes in the network
α^{SD}	Cost per time unit spent driving alone
α^{PT}	Cost per time unit spent in public transport
α^{RS}	Cost per time unit spent sharing a ride
α^{wait}	Cost per time unit spent waiting at a transfer hub
β	Cost per time unit arriving early at the destination
γ	Cost per time unit arriving late at the destination
τ	Maximum detour a driver is willing to make
ϕ^{SD}	Variable cost per time unit for driving alone
ϕ^{PT}	Variable cost per time unit for public transport
ϕ^{RS}	Variable cost per time unit for sharing a ride
ψ^{SD}	Fixed (access) cost for driving alone
ψ^{PT}	Fixed (access) cost for public transport
ψ^{RS}	Fixed (access) cost for sharing a ride
Decision variables	
x_k	Binary decision variable indicating if rider path $k \in \mathcal{K}$ is selected
y_{jh}	Binary decision variable indicating if driver j travels through transfer hub j

Direct RS

For every rider $i \in \mathcal{I}$, a direct match can be found with a driver $j \in \mathcal{J}$ if $o_i = o_j$ and $d_i = d_j$. As the driver selects the departure time to minimize his/her own cost, the arrival time at the final destination is equal to his/her desired arrival time of the driver, possibly imposing schedule delay costs on the rider. A rider is penalized for earliness by β and for lateness by γ. The notation $(\cdot)^+ = \max(0,\cdot)$, which means that either earliness or lateness is positive, but not both at the same time. Only if $t^*_i = t^*_j$, the rider arrives exactly on time, and therefore both earliness and lateness will be zero. For a match between $i \in \mathcal{I}$ and $j \in \mathcal{J}$, $e_{ik} = 1$, $a^0_{jk} = 1$ and all other parameters are equal to 0. The cost of this direct match is as follows:

$$c_k = (\alpha^{RS} + \phi^{RS})tt(o_i, d_i) + \beta(t^*_i - t^*_j)^+ + \gamma(t^*_j - t^*_i)^+ + \psi^{RS} \tag{5}$$

The seven indirect paths are described below:

Indirect RS → RS

We consider a rider $i \in \mathcal{I}$ and two drivers $j_1, j_2 \in \mathcal{J}$ where j_1 takes i on the first leg and j_2 takes i on the second leg with a transfer at transfer hub h. Similar to before, this is only feasible if $o_i = o_{j_1}$, $d_i = d_{j_2}$. Thereby, $t^*_{j_1}(h) \le t^*_{j_2}(h)$ to ensure that the rider is dropped off at the transfer hub before the scheduled pickup. The hub h needs to deviate at most τ min from the shortest path of both drivers j_1 and j_2. In this case, $a^{1h}_{j_1 k} = 1$ and $a^{2h}_{j_2 k} = 1$. The cost for the rider i is then defined as follows:

$$\begin{aligned} c_k = {} & (\alpha^{RS} + \phi^{RS})[tt(o_i, h) + tt(h, d_i)] + \alpha^{\text{wait}}[t^*_{j_2}(h) - t^*_{j_1}(h)] \\ & + \beta[t^*_i - t^*_{j_2}(h) - tt(h, d_i)]^+ + \gamma[t^*_{j_2}(h) + tt(h, d_i) - t^*_i]^+ + 2f\psi^{RS} \end{aligned} \tag{6}$$

Indirect RS → PT

For a path where only the first leg is a ride-sharing leg, a rider knows in advance when he will be picked up at the transfer hub and can therefore adjust his departure time on the first leg to the departure on the second leg. If rider $i \in \mathcal{I}$ and driver $j \in \mathcal{J}$ are matched with a transfer at hub $h \in \mathcal{H}$, they must share their origin $o_i = o_j$ and driver j may deviate at most τ min from his/her shortest path to reach hub h. In this case, $a^{1h}_{j_1 k} = 1$. The cost of this match is:

$$c_k = (\alpha^{RS} + \phi^{RS})tt(o_i, h) + (\alpha^{PT} + \phi^{PT})tt(h, d_i) + \beta[t^*_i - t^*_j(h) - tt(h, d_i)]^+ + \gamma[t^*_j(h) + tt(h, d_i) - t^*_i]^+ + f(\psi^{RS} + \psi^{PT}) \quad (7)$$

Indirect PT → RS

We consider the path where the rider shares a ride on the second leg and takes public transport on the first leg. If rider $i \in \mathcal{I}$ and driver $j \in \mathcal{J}$ are matched with a transfer at hub $h \in \mathcal{H}$, they must share their destination $d_i = d_j$ and driver j may deviate at most τ min from his/her shortest path to reach hub h. In this case, $a^{2h}_{j_2 k} = 1$. The cost of this match is:

$$c_k = (\alpha^{PT} + \phi^{PT})tt(o_i, h) + (\alpha^{RS} + \phi^{RS})tt(h, d_i) + \beta[t^*_i - t^*_j(h) - tt(h, d_i)]^+ + \gamma[t^*_j(h) + tt(h, d_i) - t^*_i]^+ + f(\psi^{PT} + \psi^{RS}) \quad (8)$$

Indirect SD → RS

We consider an indirect path where a rider drives alone on the first leg and shares a ride on the second leg. We note that the arrival time at the transfer hub should be coordinated, similar to an indirect ride-sharing match. Let rider $i \in \mathcal{I}$ be matched to driver $j \in \mathcal{J}$ on the second leg and let the departure time of rider i on the first leg be equal to $t \in T$. The arrival time of the rider at the transfer hub $h \in \mathcal{H}$ is then equal to $t + tt(o_i, h)$. In a deterministic setting, the rider can optimize their departure time t to arrive exactly on time at the transfer point and therefore incur no waiting time. Again, driver j may deviate at most τ min from his/her shortest path to reach hub h. In this case, $a^{2h}_{j_2 k} = 1$. The cost is defined as follows:

$$c_k = (\alpha^{SD} + \phi^{SD})tt(o_i, h) + (\alpha^{RS} + \phi^{RS})tt(h, d_i) + \alpha^{\text{wait}}[t^*_j(h) - t - tt(o_i, h)] + \beta[t^*_i - t^*_j(h) - tt(h, d_i)]^+ + \gamma[t^*_j(h) + tt(h, d_i) - t^*_i]^+ + f(\psi^{SD} + \psi^{RS}) \quad (9)$$

Indirect SD → PT

In case a rider $i \in \mathcal{I}$ drives their own car on the first leg and transfers to public transport at hub $h \in \mathcal{H}$, the cost is defined as follows:

$$c_k = (\alpha^{SD} + \phi^{SD})tt(o_i, h) + (\alpha^{PT} + \phi^{PT})tt(h, d_i) + f(\psi^{SD} + \psi^{PT}) \quad (10)$$

Indirect PT → PT and SD → SD

For completeness, we introduce two indirect alternatives where the same mode is used on both legs. Clearly, in the deterministic case, a direct option with that same mode would always perform at least as well and therefore these alternatives will never be used in practice. The cost for an indirect public transport path and an indirect solo drive path for individual i with a transfer at hub h are given as follows:

$$c_k = (\alpha^{PT} + \phi^{PT})(tt(o_i, h) + tt(h, d_i)) + 2f\psi^{PT} \quad (11)$$

$$c_k = (\alpha^{SD} + \phi^{SD})(tt(o_i, h) + tt(h, d_i)) + 2f\psi^{SD} \quad (12)$$

Appendix C. Theorems and proofs

Theorem 1 (*Optimal Departure Time for a Single Leg*)**.** *Let all riders and the driver have identical origins and destinations. Let a driver with desired arrival time t^*_0 be matched to N riders with desired arrival times $t^*_1 \dots t^*_N$. With $\max(\beta_0, \beta_1, \dots \beta_N) < \min(\gamma_0, \gamma_1, \dots \gamma_N)$, the jointly optimal departure time is equal to $t^o = \min(t^*_0, \dots, t^*_N)$.*

Proof. Without loss of generality, travel time is set equal to 0. Let t^o be the jointly optimal departure time and let $C(t)$ be the joint total cost for all drivers and passengers. Then, $t^o = \arg\min_t C(t)$. Without loss of generality, we sort the desired arrival times such that $t^*_0 \le t^*_1 \le \dots \le t^*_N$. The cost $C(t)$ is then defined as follows:

$$C(t) = \sum_{i \in \{0,\dots,N\} | t^*_i < t} \beta_i(t - t^*_i) + \sum_{i \in \{0,\dots,N\} | t^*_i \ge t} \gamma_i(t^*_i - t) \quad (13)$$

The function $C(t)$ is piece-wise linear and therefore the optimal departure time t^o has to be at one of the breakpoints $\{t^*_0, \dots t^*_N\}$. A graphic example is displayed in Fig. 16. Given that $\max(\beta_0, \beta_1, \dots \beta_N) < \min(\gamma_0, \gamma_1, \dots \gamma_N)$, it can be shown by contradiction that $t^o = \min(t^*_0, \dots, t^*_N)$. □

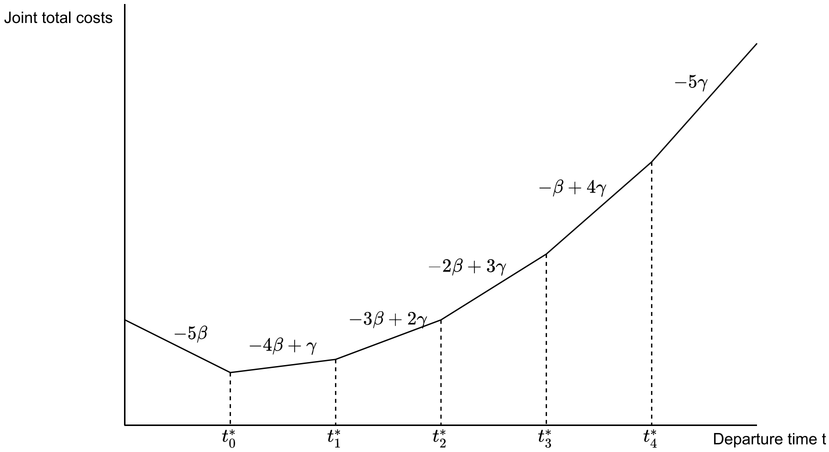

Fig. 16. Theorem 1 example with $N = 4$ riders and homogeneous β and γ.

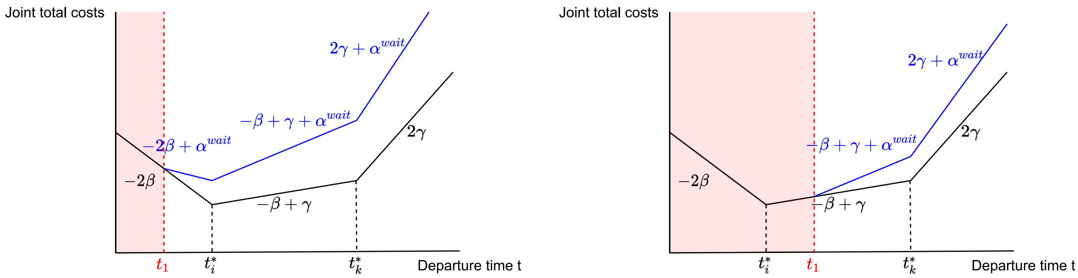

Fig. 17. Schedule delay and waiting time.

Theorem 2 (*Optimal Departure Time for a Second Leg Trip*). *Consider N riders $k_1, \dots k_N$ from origins $o_{k_1}, \dots, o_{k_N}$ who transfer at hub h to their identical destination d, and a driver i from hub h to the same destination d, with their desired arrival times $t^*_{k_1}, \dots t^*_{k_N}$ and t^*_i. Let all individuals have identical cost parameters $\alpha, \alpha^{\text{wait}}, \beta, \gamma$, with $\beta < \gamma$ and $\alpha^{\text{wait}} < \gamma$. We let t_1 be the last departure time for the first leg among all riders and the driver. The joint optimal departure time for the second leg t^o_2 is a function of the departure time for the first leg t_1 which is defined as follows:*

$$t^o_2(t_1) = \begin{cases} \max(t_1, \min(t^*_i, t^*_{k_1}, \dots, t^*_{k_N})) & \text{if } \alpha^{\text{wait}} \le \beta \\ t_1 & \text{if } \alpha^{\text{wait}} > \beta \end{cases} \tag{2}$$

Proof. Without loss of generality, travel time is equal to 0. Let $t^o_2(t_1)$ be the jointly optimal departure time on the second leg given the latest departure time t_1 on the first leg and let $C_2(t_1, t_2)$ be the joint total cost on the second leg for all drivers and passengers where t_1 is the latest departure time on the first leg and t_2 is the departure time on the first leg. Then, $t^o_2(t_1) = \arg\min_{t_2} C(t_1, t_2)$. Clearly, leaving before the last passenger has arrived makes the match infeasible and therefore attains a cost of ∞. This implies that $t^o_2(t_1) \ge t_1$. Therefore, in the remainder of this proof, we disregard the period before t_1 and the costs of waiting during that period.

Without loss of generality, we sort the desired arrival times such that $t^*_{k_1} \le \dots \le t^*_{k_N}$. The cost $C_2(t_1, t_2)$ is then defined as follows:

$$C_2(t_1, t_2) = (N+1)\alpha^{\text{wait}}(t_2 - t_1) + \sum_{k \in \{i, k_0, \dots, k_N\} | t^*_k < t} \beta(t - t^*_k) + \sum_{k \in \{i, k_0, \dots, k_N\} | t^*_k \ge t} \gamma_i (t^*_k - t) \tag{14}$$

Using the same reasoning as in Theorem 1, earliness is jointly preferred over lateness. In this case, we have an additional trade-off between earliness and waiting time. If waiting time is penalized more than earliness, it is best to leave immediately after everyone has arrived such that $t^o_2(t_1) = t_1$. Otherwise, it might be better to wait. We separately consider the cases where (i) $t_1 > \min(t^*_i, t^*_{k_1}, \dots, t^*_{k_N})$ and (ii) $t_1 \le \min(t^*_i, t^*_{k_1}, \dots, t^*_{k_N})$. This is graphically depicted in Fig. 17. In case (i), at least one matched individual is already late and therefore waiting longer is definitely not desirable as $\alpha^{\text{wait}} < \gamma$. In that case, $t^o_2(t_1) = t_1$. In case (ii), it is better to wait at the transfer hub until $t^o_2(t_1) = \min(t^*_i, t^*_{k_1}, \dots, t^*_{k_N})$, applying the reasoning from Theorem 1. Combining these individual cases leads to the optimal departure time on the second leg as given in Eq. (2). □

Theorem 3 (*Optimal Departure Time on First Leg*). *Consider a driver i from o to d passing through hub h, one rider j from o to hub h and one rider k from hub h to d, with their desired arrival times t_i^*, t_j^*, t_k^*. Let all individuals have identical cost parameters $\alpha, \alpha^{\text{wait}}, \beta, \gamma$, with $\beta < \alpha^{\text{wait}} < \gamma$. The joint optimal departure time for the first leg $t_1{}^{\circ} = \min(t_i^*, t_j^*, t_k^*)$. The joint optimal departure time of the second leg can then be determined according to Theorem 2.*

Proof. Without loss of generality, travel time is set equal to 0. Let $t_2^o(t_1)$ be the jointly optimal departure time on the second leg given the departure time t_1 on the first leg and let $C(t_1, t_2)$ be the joint cost of the driver and the two riders on both legs. Because $\beta < \alpha^{\text{wait}}$, according to Theorem 2, $t_2^o(t_1) = t_1$ given the relationship between the parameters. We can therefore determine $C(t_1) = C(t_1, t_2^o(t_1))$ which is only composed of schedule-delay costs and not waiting, given the immediate departure from the transfer hub. By applying Theorem 1 we obtain $t_1^o = \min(t_i^*, t_j^*, t_k^*)$. □

Data availability

Open source data was used.

References

Agatz, N., Erera, A., Savelsbergh, M., Wang, X., 2012. Optimization for dynamic ride-sharing: A review. European J. Oper. Res. 223 (2), 295–303.

Alisoltani, N., Leclercq, L., Zargayouna, M., 2021. Can dynamic ride-sharing reduce traffic congestion? Transp. Res. B: Methodol. 145, 212–246.

Arnott, R., De Palma, A., Lindsey, R., 1993. A structural model of peak-period congestion: A traffic bottleneck with elastic demand. Am. Econ. Rev. 161–179.

Beojone, C.V., Geroliminis, N., 2021. On the inefficiency of ride-sourcing services towards urban congestion. Transp. Res. C: Emerg. Technol. 124, 102890.

Bertsimas, D., Tsitsiklis, J., 1993. Simulated annealing. Statist. Sci. 8 (1), 10–15.

Caulfield, B., 2009. Estimating the environmental benefits of ride-sharing: A case study of Dublin. Transp. Res. D: Transp. Environ. 14 (7), 527–531.

Cheikh, S.B., Hammadi, S., 2016. Multi-Hop ridematching optimization problem: Intelligent chromosome agent-driven approach. Expert Syst. Appl. 62, 161–176.

Chen, X., Deng, H., Guan, S., Han, F., Zhu, Z., 2024. Cooperation-oriented multi-modal shared mobility for sustainable transport: Developments and challenges. Sustainability 16 (24), 11207.

Chen, Y., Guo, D., Xu, M., Tang, G., Zhou, T., Ren, B., 2019. PPtaxi: Non-stop package delivery via multi-hop ridesharing. IEEE Trans. Mob. Comput. 19 (11), 2684–2698.

Chen, W., Mes, M., Schutten, M., 2018. Multi-hop driver-parcel matching problem with time windows. Flex. Serv. Manuf. J. 30 (3), 517–553.

City of Chicago, 2010. Boundaries - Census tracts - 2010. https://data.cityofchicago.org/Facilities-Geographic-Boundaries/Boundaries-Census-Tracts-2010/5jrd-6zik.

City of Chicago, 2022. Transportation network providers - trips (2018 - 2022). https://data.cityofchicago.org/Transportation/Transportation-Network-Providers-Trips-2018-2022-/m6dm-c72p.

de Palma, A., Javaudin, L., Stokkink, P., Tarpin-Pitre, L., 2022a. Ride-sharing with inflexible drivers in the Paris metropolitan area. Transp. 1–24.

de Palma, A., Stokkink, P., Geroliminis, N., 2022b. Influence of dynamic congestion with scheduling preferences on carpooling matching with heterogeneous users. Transp. Res. B: Methodol. 155, 479–498.

Enzi, M., Parragh, S.N., Pisinger, D., 2024. Modeling and solving a corporate vehicle-sharing problem combined with other modes of transport. EURO J. Transp. Logist. 13, 100122.

Gurumurthy, K.M., Kockelman, K.M., Simoni, M.D., 2019. Benefits and costs of ride-sharing in shared automated vehicles across Austin, Texas: Opportunities for congestion pricing. Transp. Res. Rec. 2673 (6), 548–556.

Herbawi, W., Weber, M., 2012a. Modeling the multihop ridematching problem with time windows and solving it using genetic algorithms. In: 2012 IEEE 24th International Conference on Tools with Artificial Intelligence. Vol. 1, IEEE, pp. 89–96.

Herbawi, W., Weber, M., 2012b. The ridematching problem with time windows in dynamic ridesharing: A model and a genetic algorithm. In: 2012 IEEE Congress on Evolutionary Computation. IEEE, pp. 1–8.

Hou, Y., Li, X., Qiao, C., 2012. TicTac: From transfer-incapable carpooling to transfer-allowed carpooling. In: 2012 IEEE Global Communications Conference. GLOBECOM, IEEE, pp. 268–273.

Hryhoryeva, M., Leclercq, L., 2024. Regulating Ride-Hailing service operations: Impact of idle distance charge and demand distribution on network performance. Available At SSRN 4760780.

Huang, H., Bucher, D., Kissling, J., Weibel, R., Raubal, M., 2018. Multimodal route planning with public transport and carpooling. IEEE Trans. Intell. Transp. Syst. 20 (9), 3513–3525.

Jung, J., Jayakrishnan, R., Park, J.Y., 2016. Dynamic shared-taxi dispatch algorithm with hybrid-simulated annealing. Comput.- Aided Civ. Infrastruct. Eng. 31 (4), 275–291.

Kirkpatrick, S., Gelatt, Jr., C.D., Vecchi, M.P., 1983. Optimization by simulated annealing. Sci. 220 (4598), 671–680.

Kong, H., Zhang, X., Zhao, J., 2020. How does ridesourcing substitute for public transit? A geospatial perspective in Chengdu, China. J. Transp. Geogr. 86, 102769.

Kumar, P., Khani, A., 2021. An algorithm for integrating peer-to-peer ridesharing and schedule-based transit system for first mile/last mile access. Transp. Res. C: Emerg. Technol. 122, 102891.

Lee, E., Slama, R., Leclercq, L., 2024. Assessing how Ride-hailing rebalancing strategies improve the resilience of Multi-modal transportation systems. arXiv preprint arXiv:2412.00276.

Li, J., Embry, P., Mattingly, S.P., Sadabadi, K.F., Rasmidatta, I., Burris, M.W., 2007. Who chooses to carpool and why? Examination of Texas carpoolers. Transp. Res. Rec. 2021 (1), 110–117.

Li, Y., Li, X., Zhang, S., 2021. Optimal pricing of customized bus services and ride-sharing based on a competitive game model. Omega 103, 102413.

Li, M., Xu, J., Chen, C., Zhang, W., Long, X., Guo, S., 2024. HiCrowd: Hierarchical Multi-hop crowdsourcing delivery by leveraging the crowd of taxi. In: 2024 IEEE Smart World Congress. SWC, IEEE, pp. 306–312.

Liu, Y., Han, K., Yang, Z., Yu, Y., Ji, W., 2024. Integrating optimal ridesharing matching into multimodal traffic model: Implications for policy and sustainable transport system. arXiv preprint arXiv:2411.15427.

Lu, W., Liu, L., Wang, F., Zhou, X., Hu, G., 2020. Two-phase optimization model for ride-sharing with transfers in short-notice evacuations. Transp. Res. C: Emerg. Technol. 114, 272–296.

Ma, T.-Y., Rasulkhani, S., Chow, J.Y., Klein, S., 2019. A dynamic ridesharing dispatch and idle vehicle repositioning strategy with integrated transit transfers. Transp. Res. E: Logist. Transp. Rev. 128, 417–442.

Masoud, N., Jayakrishnan, R., 2017a. A decomposition algorithm to solve the multi-hop peer-to-peer ride-matching problem. Transp. Res. B: Methodol. 99, 1–29.

Masoud, N., Jayakrishnan, R., 2017b. A real-time algorithm to solve the peer-to-peer ride-matching problem in a flexible ridesharing system. Transp. Res. B: Methodol. 106, 218–236.
Masoud, N., Nam, D., Yu, J., Jayakrishnan, R., 2017. Promoting peer-to-peer ridesharing services as transit system feeders. Transp. Res. Rec. 2650 (1), 74–83.
Olszewski, P., Krukowski, P., 2012. Quantitative assessment of public transport interchanges. In: 40th European Transport Conference, Glasgow.
Özkan, E., Ward, A.R., 2020. Dynamic matching for real-time ride sharing. Stoch. Syst. 10 (1), 29–70.
Pincus, M., 1970. A Monte Carlo method for the approximate solution of certain types of constrained optimization problems. Oper. Res. 18 (6), 1225–1228.
Raviv, T., Tenzer, E.Z., 2018. Crowd-Shipping of Small Parcels in a Physical Internet. Workingpaper, Tel Aviv University.
Sasaki, G.H., Hajek, B., 1988. The time complexity of maximum matching by simulated annealing. J. ACM 35 (2), 387–403.
Schaller, B., 2021. Can sharing a ride make for less traffic? Evidence from Uber and Lyft and implications for cities. Transp. Policy 102, 1–10.
Shaheen, S., Cohen, A., 2019. Shared ride services in North America: definitions, impacts, and the future of pooling. Transp. Rev. 39 (4), 427–442.
Sil, A., Chaniotakis, E., Roy, U.K., Antoniou, C., 2022. Exploring satisfaction for transfers at intermodal interchanges: A comparison of Germany and India. J. Public Transp. 24, 100005.
Small, K.A., 1982. The scheduling of consumer activities: work trips. Am. Econ. Rev. 72 (3), 467–479.
Stiglic, M., Agatz, N., Savelsbergh, M., Gradisar, M., 2018. Enhancing urban mobility: Integrating ride-sharing and public transit. Comput. Oper. Res. 90, 12–21.
Stokkink, P., Cordeau, J.-F., Geroliminis, N., 2024. A column and row generation approach to the crowd-shipping problem with transfers. Omega 103134.
Tang, S., Zheng, F., Zheng, N., Liu, X., 2024. An efficient multi-modal urban transportation network partitioning approach for three-dimensional macroscopic fundamental diagram. Phys. A 637, 129487.
Teal, R.F., 1987. Carpooling: Who, how and why. Transp. Res. A: Gen. 21 (3), 203–214.
U.S. Department of Energy (DOE), 2022. Oak Ridge National Lab (2022) Transportation Energy Data Book Edition 40..
Vickrey, W.S., 1963. Pricing in urban and suburban transport. Am. Econ. Rev. 53 (2), 452–465.
Wang, X., Agatz, N., Erera, A., 2018. Stable matching for dynamic ride-sharing systems. Transp. Sci. 52 (4), 850–867.
Xu, M., Di, Y., Zhu, Z., Yang, H., Chen, X., 2024. Design and analysis of ride-sourcing services with auxiliary autonomous vehicles for transportation hubs in multi-modal transportation systems. Transp. B: Transp. Dyn. 12 (1), 2333869.
Xu, Y., Kulik, L., Borovica-Gajic, R., Aldwyish, A., Qi, J., 2020. Highly efficient and scalable multi-hop ride-sharing. In: Proceedings of the 28th International Conference on Advances in Geographic Information Systems. pp. 215–226.
Yan, P., Lee, C.-Y., Chu, C., Chen, C., Luo, Z., 2021. Matching and pricing in ride-sharing: Optimality, stability, and financial sustainability. Omega 102, 102351.
Yan, H., Liao, Y., Ma, Z., Ma, X., 2024. Improving multi-modal transportation recommendation systems through contrastive De-biased heterogenous graph neural networks. Transp. Res. C: Emerg. Technol. 164, 104689.
Yu, V.F., Purwanti, S.S., Redi, A.P., Lu, C.-C., Suprayogi, S., Jewpanya, P., 2018. Simulated annealing heuristic for the general share-a-ride problem. Eng. Optim. 50 (7), 1178–1197.