A ride-sharing problem with meeting points and return restrictions

W. Chen, M.R.K. Mes, J.MJ. Schutten, J. Quint

A ride-sharing problem with meeting points and return restrictions

Wenyi Chen*, Martijn Mes, Marco Schutten, Job Quint

Department of Industrial Engineering and Business Information System, University of Twente, the Netherlands

w.chen-5@utwente.nl,   m.r.k.mes@utwente.nl,   m.schutten@utwente.nl

Ride-sharing has been widely acknowledged as an effective solution for reducing travel costs, congestion and pollution. This paper considers the ride-sharing problem of the scheduled commuter and business traffic within a closed community of companies that agree to share the calendars of their employees. We propose a general ILP formulation for the aforementioned ride-sharing problem, which incorporates return restrictions in order to satisfy the business needs, as well as meeting points and the option for riders to transfer between drivers. All the instances with 40 and 60 participants, and most of the instances with 80 participants can be solved to optimality within a time limit of 2 hours. Using instances of up to 100 participants, the ILP can be solved with a gap of no more than 1.5% within the time limit. Due to the high computational complexity, we develop a constructive heuristic that is based on the saving concepts. This heuristic is also able to combine ride-sharing with the use of an external mobility service provider. Our numerical study shows that ride-sharing can be an effective way of reducing the number of trips and vehicle-miles. Particularly, ride-sharing creates more benefits when the participation is high, and when the origins and the destinations of the trips are more spatially concentrated. The results show that ride-sharing can create up to 35.2% mileage savings and up to 23.3% reduction in the number of cars needed to fulfill employees' travel schedules. We also illustrate our model using a real-life ride-sharing problem of a Dutch consultancy and research firm.

Key words: ride-sharing; carpooling; sustainable transport; transfers; return restrictions
History:

1.   Introduction

Rising motorization and increasing traffic density has intensified the problem of GHG emissions worldwide (Kuntzky et al. 2013). The transport sector accounts for 24.4% of the EU's total GHG emissions in 2013; passenger cars contribute almost 45% of the transport sector's emissions. Although the fuel efficiency and emission characteristics of passenger cars improved steadily, rapid growth in car ownership and in distances traveled offset any potential improvements in terms of environmental impacts (EEA 2013). Even with increasing environmental awareness and concern, many road users are still car-dependent, either by choice or constrained by circumstances (Stradling 2007). Given the constantly low car occupancy rate, e.g., Lisbon on average 1.2 persons per car, Sydney 1.3, London 1.5, Singapore 1.6 (ITF 2015), ride-sharing plays an increasingly important role in providing mobility, and reducing CO_2 emissions, traffic congestion and parking problems. Assuming that one person was added to each commute, IEA (2005) estimated that carpooling can reduce the number of kilometers traveled by 12.5%, which will lead to 7.7% reduction in fuel use. Delhomme and Gheorghiu (2016) provide an example: the people who carpool for a distance of 48 kilometers could save up to 33% of the monthly costs of commuting compared to those who choose to drive alone. Additionally, ride-sharing may save time because commuters are able to use high-occupancy vehicle lanes reserved for the vehicles with two or more occupants, reduce driver fatigue (Stiglic et al. 2015), and lead to social benefits by enlarging carpoolers' social networks (Agatz et al. 2012).

Ride-sharing services currently on the market range from simple online bulletin boards to sophisticated systems that offer real-time on-demand matching, routing and payment service (Stiglic

---

* corresponding author

et al. 2015, Furuhata et al. 2013). Ride-sharing providers like Uber, Lyft, and Flinc show that innovative use of technology can revolutionize personal mobility (Savelsbergh and Van Woensel 2016). Due to the challenge of coordinating itineraries and schedules between participants, ride-sharing coordination is mostly an informal and disorganized activity and only in certain cases travelers can make use of ride-sharing as a regular transportation alternative (Furuhata et al. 2013). In this paper, we focus on systems that offer automated matching of scheduled commuter/business traffic within a closed community of companies that agree to share their calendars. An example of a provider offering such a service is Zimride (`http://zimride.com/`). The service provider connects inter-city drivers and passengers through social networking. One of its signature service is to provide a private ride-sharing network within a corporate or a university (e.g., Harvard University). People within the same network register their trip, typically on a daily basis. Each trip contains detailed information, such as origin and destination of each trip, earliest departure and latest arrival times, the acceptable maximal detour, and the capacity of the car, if it is available. In addition, it also contains registration information of the person such as age, gender, educational level, special interests, etc., which may also be obtained from his/her social network. The system provides suggestions to individuals to share a car based on the details of their trips, as well as their registration information.

Companies also have a financial motivation to support ride-sharing among employees, especially companies that offer individually owned company cars to their employees. Since employees are not directly confronted with the marginal costs of using the cars, access to company cars leads to higher car use (van Dender and Clever 2013), and higher costs for companies to provide mobility service with these individually owned company cars. It is also socially undesirable to have cars with low occupancy on the road. Therefore, it is for a company's benefit to reduce the number of single-person trips by car. A recent development is to replace the individually owned companies cars with accessibility to mobility in general, through the use of public transport and shared cars. Ride-sharing among colleagues can also be a promising solution. In doing so, the company can reduce the total mobility cost, by reducing the total vehicle-miles driven by all employees as well as the total number of vehicles needed to fulfill their mobility demands. Furthermore, the outcomes of ride-sharing are also aligned with societal objectives for reducing emissions and traffic congestion, which should also be of companies' interest concerning the corporate social responsibility.

The goal of this paper is to provide the means for a closed corporate community to facilitate ride-sharing by matching the employees. Such a community can be a company, or a consortium of companies that are willing to share their calendars in some way. It can also be initiated by a car leasing company that aims to provide mobility services instead of only leased cars. In general, employees' agendas are well planned in advance. Thus, we limit our attention to the offline ride-share matching problem: given the detailed information about people's trips within a given time period, as well as the information of the users, such as maximum acceptable detour and the capacity of the car, find the optimal matching to minimize the overall cost of the commuter traffic. In order to increase the chance of matching, we incorporate the features of meeting points and multiple hops in the model. Furthermore, we provide a heuristic for solving non-trivial problem instances of the considered NP-hard optimization problem.

From a company's perspective, the ride-sharing service should not affect the employees' work related mobility. Thus, one of the main contributions of our work is the incorporation of return restrictions. That is, a match is possible only if (i) the rider can reach all his destinations during the planning horizon and (ii) he can return to the meeting point where his car is parked to drive to his final destination. As a result, the role of a person being a driver or a passenger is not fixed. Even more challengingly, it can change during the planning horizon.

The remainder of the paper is organized as follows. In the next section, we position our research in the context of the relevant literature. After introducing the ride-sharing model with meeting

points and return restriction (RS-M&R) in Section 3, the mixed integer programming formulation is presented in Section 4. We propose a heuristic for solving the RS-M&R in Section 5. In Section 6, the proposed methodology is illustrated using a real-life ride-sharing problem of Significant BV, a Dutch consultancy and research firm. Section 7 provides the experimental settings. The numerical results are presented in Section 8. Section 9 concludes the paper with key findings and directions for future research.

2. Literature review

In recent years, ride-sharing has received a growing interest both in academia and industry. Using criteria such as degree of automation, demand types and matching search criteria, Furuhata et al. (2013) present a classification of existing ride-sharing systems. The popularity of dynamic ride-sharing providers has stimulated a growing body of research on optimization technology in this topic, e.g., Winter and Nittel (2006), Agatz et al. (2011), Stiglic et al. (2015), and Lee and Savelsbergh (2015). Agatz et al. (2012) provide an overview of dynamic ride-sharing and the relevant algorithmic approaches for matching drivers and riders in real-time. They point out that transfers are not considered in the literature yet due to the increasing computational burden resulting from the increased number of drivers and riders that are involved in a matching in a multi-hop setting.

One way to increase a rider's chance of finding a match is to allow him to transfer between different drivers to reach his destination. Gruebele (2008) describes such a multi-hop and multi-rider routing system in detail, without providing a solution methodology. Herbawi and Weber (2011) consider a single rider version of the multi-hop ride-sharing problem, where drivers do not deviate from their routes and schedules. An evolutionary multi-objective route planning algorithm is used to obtain good quality solutions in reasonable runtime. Herbawi and Weber (2012) extend the previous work to match multiple riders with multiple drivers having time windows and allowing a possible detour from their routes. They propose a genetic algorithm and show that it can be used to solve the model in reasonable time. Drews and Luxen (2013) show that the problem studied by Herbawi and Weber (2012) can also be solved by exploiting time-expanded graphs representing the drivers' offers.

In the traditional recurring ride-sharing problem, a match is possible only if a driver is able to pick up the rider at his starting location and drop him off at his ending location. A more recent development to facilitate ride-sharing is to consider a setting in which commuters travel to and from meeting points. We are aware of two papers that consider meeting points in the context of ride-sharing. Assuming the matching between a driver and a rider is done, Aissat and Oulamara (2014) focus on finding the start and end meeting points so as to minimize the total travel distance of all the drivers. No restriction is placed on the locations of the starting and end meeting points relative to the rider's origin and destination, respectively. If a rider's travel effort from his origin to the starting meeting point and from the end meeting point to his destination is considered, the optimal solution may be very different. Closely related research that also consider ride-sharing with meeting points can be found in Stiglic et al. (2015). This work represents the single driver multiple rider setting. The drivers are allowed to make only one pickup and one drop-off. The potential locations of the pickup and drop-off points are confined in a certain walking distance from the riders' origin or destination. Our work is different in several ways. First, the participants who have a car are also allowed to ride with others. Thus, potential meeting points are defined differently, and the participants with a car can have flexible roles. Second, transfers are allowed. Third, return restrictions are incorporated. Fourth, a participant can have more than two trips during the planning horizon. As a result of the listed differences, the bipartite matching structure and the solution method proposed in Stiglic et al. (2015) are not applicable to our problem.

Agatz et al. (2011) also consider return trip matches and flexible roles, but focus on the single driver, single rider ride-sharing problem. When instances contain large numbers of participants with flexible roles, they find it difficult to solve the offline problems to optimality.

In this paper, we consider a multi-hop ride-sharing problem with the incorporation of meeting points and return restrictions. The contribution of this paper is threefold. First, we design an integer linear program (ILP) specifically tailored to the essentials of a closed corporate ride-sharing environment, which is characterized by the foreseeability of the participants' schedules yet limited flexibility in their schedules and itineraries. Accordingly, our model incorporates the following design features: (i) meeting points, (ii) return restrictions, (iii) transfers, (iv) flexible roles of being a driver or a rider, if possible, and (v) the possibility of switching roles during the planning horizon. Second, we show that the ILP model provides high quality solutions to ride-sharing problem instances up to 100 participants, and also provides useful insights into the potential benefits of implementing ride-sharing. Third, considering the difficulty in finding a good feasible solution for such a complex problem, we propose a constructive heuristic, which can serve as a good starting point for improvement heuristics for solving large-scale problem instances. The proposed heuristic is also able to consider an extended problem where an external mobility service provider can be used as a backup option.

3.   Problem description

The problem under consideration is defined on a directed graph $G = (\mathcal{N}, \mathcal{A})$, where $\mathcal{N}$ is the set of nodes representing the possible locations for departure, arrival, or transfer, and $\mathcal{A}$ is the set of arcs that directly connect two aforementioned locations, i.e., represents the road network. With each arc $(i, j) \in \mathcal{A}$, a travel time t_{ij} is associated based on the travel distance.

We are given a set of persons $\mathcal{P}$, with or without a car, who need to travel from one location to another. During the planning horizon, say a day, a person $p \in \mathcal{P}$ might have multiple trips, each of which is considered as a ride $r \in \mathcal{R}$. On ride r, p will travel from his origin o_{pr} to his destination d_{pr}, and SP_{pr} represents the set of arcs belonging to his shortest path from o_{pr} to d_{pr}. An earliest time e_{pr} at which he can depart from his origin o_{pr} and a latest time l_{pr} at which he has to arrive at his destination d_{pr} are also associated with person p. For those participants who have a car, the choice of being a driver or a passenger is flexible. A driver may take a single passenger or multiple passengers (sequentially or simultaneously) along the journey, as long as the capacity of his vehicle v_p is not exceeded. Similarly, a passenger may ride with a single driver from his origin to his destination or may transfer from one driver to another en route to his destination. It is possible that a participant p needs to wait during the transfers, the total waiting time is capped by the participant's maximum waiting time m_{pr}.

Our objective is to develop a mechanism for ride-sharing within a closed corporate community in order to reduce the overall cost of operating the commuter and business traffic, which consists of the costs associated with (i) vehicle miles such as fuel and tolls, and (ii) the inconvenience and efficiency loss due to the ride-sharing including the transfers and time losses. One of the key characteristics of the commuter/business traffic is the limited flexibility in employees, itineraries and schedules. To overcome such a challenge, the ride-sharing system has to be well designed to minimize the effort and inconvenience for the participants.

That being said, three important features are considered in our ride-sharing problem. First, the participants will be matched only when limited detour is required and a pre-processing procedure will be introduced to determine the limited allowable detour in Section 6. Second, meeting points are introduced to take advantage of any flexibility of every participant in terms of time and mobility, and construct routes with smaller detours. Most probably, people have multiple trips during a day. The simplest example is a return trip between home and the workplace. Thus, the third feature is to impose return restrictions, such that (i) the participants who leave their car and start sharing rides are able to arrive at their destinations throughout the day, and (ii) the participants and their cars are able to return to the home locations at the end of the day.

Table 1   **Parameters and decision variables for the RS-M&R model.**

Sets	
$\mathcal{P}$	Set of persons, containing drivers and passengers (indexed by q, p)
$\mathcal{N}$	Set of nodes (indexed by i, j, k)
$\mathcal{A}$	Set of arcs, each arc is indicated by (i,j), with $i,j \in \mathcal{N}$
$\mathcal{R}_p$	Set of rides of participant p (indexed by r, w)
Decision variables	
Z_{prij}	Binary variable equal to 1 if person p drives directly from node i to node j on ride r; 0 otherwise
Y_{pqrwij}	Binary variable equal to 1 if passenger p rides with driver q on p's ride r and q's ride w from node i to node j; 0 otherwise
D_{pri}	Departure time of person p at node i on ride r
Dependent variables	
S_{pqri}	Binary variable equal to 1 if persons p and q start carpooling at node i on p's ride r; 0 otherwise
C_{pri}	Binary variable equal to 1 if person p's car is parked at node i at the end of ride r; 0 otherwise
A_{pri}	Arrival time of person p at node i on ride r
Parameters	
o_{pr}	person p's origin on ride r
d_{pr}	person p's destination on ride r
e_{pr}	Earliest departure time of person p on ride r
l_{pr}	Latest arrival time of person p on ride r
SP_{pr}	A set of arcs that make up the shortest path of person p on ride r
x_{prij}	Binary parameter equal to 1 if arc (i,j) belongs to the shortest path of person p i.e., $(i,j) \in SP_{pr}$; 0 otherwise
v_p	Number of passenger seats available in person p's vehicle
m_{pr}	Maximum waiting time during person p's ride r
t_{ij}	Travel time from node i to node j, $\forall i,j \in \mathcal{N}$
α_1	Weight coefficient with respect to total vehicle miles
α_2	Weight coefficient with respect to penalising the arrivals that are earlier than the latest arrival time
α_3	Weight coefficient with respect to penalising the waiting time during transfers
α_4	Weight coefficient with respect to number of transfers
M	Large number

Let Z_{prij} be a binary variable that denotes whether person p travels as a driver on arc (i,j) on his ride r. Let Y_{pqrwij} be a binary variable that represents whether passenger p rides with driver q on p's ride r and q's ride w on arc (i,j). Let D_{pri} and A_{pri} denote person p's departure and arrival times at node i on his ride r. A well designed ride-sharing plan requires a seamless coordination among drivers and passengers, including the determination of (i) the ride-sharing plan Z_{prij} and Y_{pqrwij}, and (ii) the associated departure times D_{pri} and the arrival times A_{pri} at each node. A summary of the notations used in formulating the problem can be found in Table 1.

4.   Mathematical formulation

In this section, we present a mixed-integer program for the ride-sharing model with meeting points and return restriction (RS-M&R). Given the complexity of the problem, we start illustrating the ride-sharing mechanism with meeting points only (denoted by RS-M) in Section 4.1, without considering the return restrictions. Section 4.2 lays the groundwork for formulating the RS-M&R by incorporating the return restrictions into the RS-M. In Section 4.3, the proposed mixed-integer program for the RS-M&R is presented.

4.1. Ride-sharing with meeting points (RS-M)

In this subsection, we propose the ride-sharing model with meeting points (RS-M) for the single-trip problem. In order to present ideas in a transparent manner, we simplify the notations by removing the subscript r. By using the RS-M, the company can determine (i) the optimal single-trip ride-sharing arrangement among the employees, and (ii) the corresponding time schedule. The objective is to minimize the cost of the commuter and business traffic of a company, which consists of the cost incurred from vehicle miles and the costs of penalising the efficiency losses. The efficiency losses are related to (i) arriving too early for appointments, (ii) the waiting time for transfers, and (iii) the inconvenience and potential risks associated with the number of transfers. Accordingly, the objective function in our formulation of the RS-M is given by (1). Each of the four terms has a weight attached.

$$
\begin{aligned}
\min \quad & \alpha_1 \sum_p \sum_{(i,j) \in SP_p} t_{ij} Z_{pij} + \alpha_2 \sum_p (l_p - A_{p,d_p}) \\
& + \alpha_3 \sum_p \left((A_{p,d_p} - D_{p,o_p}) - \sum_{i,j} t_{ij} x_{pij} \right) \\
& + \alpha_4 \sum_p \sum_q \sum_i S_{pqi}
\end{aligned}
\tag{1}
$$

The RS-M is confined by two sets of constraints: (i) spatial constraints and (ii) capacity and time constraints.

Spatial constraints

$$
\begin{aligned}
& Y_{ppij} = 0 & & \forall p \in P, i, j \in \mathcal{N} & (2) \\
& \sum_q Y_{pqij} + Z_{pij} = x_{pij} & & \forall p \in P, i, j \in \mathcal{N} & (3) \\
& \sum_q \sum_i Y_{pqij} + \sum_k Z_{pjk} \leq 1 & & \forall p \in \mathcal{P}, j \in \mathcal{N} & (4) \\
& Y_{pqij} \leq Z_{qij} & & \forall p, q \in \mathcal{P}, i, j \in \mathcal{N} & (5) \\
& S_{pqj} \geq \sum_k Y_{pqjk} - \sum_i Y_{pqij} & & \forall p, q \in \mathcal{P}, j \in \mathcal{N} & (6) \\
& Z_{pij}, Y_{pqij}, S_{pqi} \in \{0, 1\} & & \forall p, q \in \mathcal{P}, i, j \in \mathcal{N} & (7)
\end{aligned}
$$

Constraints (2)-(7) are imposed to find feasible matches among drivers based on the spatial information (i.e., origins and destinations). Constraints (2) exclude the possibility of persons carpooling with themselves. Constraints (3) ensure that a person can either carpool with a driver or drive on his own on each arc (i,j) he has to travel. Constraints (4) ensure that person p cannot drive anymore for the rest of the trip, after he leaves his car and rides with someone else. Constraints (5) ensure that person p can ride with person q from node i to node j only if q drives from i to j. Constraints (6) keep track of the nodes where the carpools start. Constraints (7) are domain constraints.

Capacity and time constraints

$$
\begin{aligned}
& \sum_p Y_{pqij} \leq v_q & & \forall q \in \mathcal{P}, i, j \in \mathcal{N} & (8) \\
& A_{pj} = D_{pi} + t_{ij} & & \forall p \in \mathcal{P}, (i,j) \in SP_p & (9) \\
& A_{p,d_p} \leq l_p & & \forall p \in \mathcal{P} & (10) \\
& D_{p,o_p} \geq e_p & & \forall p \in \mathcal{P} & (11)
\end{aligned}
$$

$$
\begin{align}
& D_{pi} \geq A_{pi} && \forall p \in \mathcal{P}, i \in \mathcal{N} && (12) \\
& D_{pi} - D_{qi} \leq M(1 - Y_{pqij}) && \forall p, q \in \mathcal{P}, i, j \in \mathcal{N} && (13) \\
& D_{pi} - D_{qi} \geq -M(1 - Y_{pqij}) && \forall p, q \in \mathcal{P}, i, j \in \mathcal{N} && (14) \\
& (A_{p,d_p} - D_{p,o_p}) - \sum_{(i,j) \in SP_p} t_{ij} x_{pij} \leq m_p && \forall p \in \mathcal{P} && (15) \\
& D_{pi}, A_{pi} \geq 0 && \forall p \in \mathcal{P}, i \in \mathcal{N} && (16)
\end{align}
$$

Constraints (8)-(16) concern the capacity and time related issues. Constraints (8) are capacity constraints for all the drivers. Constraints (9) calculate the arrival times of persons based on the associated departure times. Constraints (10) and (11) ensure that each person departs after the corresponding earliest departure time and arrives before the corresponding latest arrival time. Clearly, the departure time cannot be earlier than the arrival time at the same station, which is considered by Constraints (12). Constraints (13) and (14) ensure that the departure time of a driver equals the departure time of the passenger whom the driver shares a ride with. Constraints (15) prevent a person's waiting time during the trip being greater than his maximum waiting time. Constraints (16) are non-negativity constraints.

4.2. Incorporating return restrictions

In this subsection, we extend the basic model to consider return restrictions. This extension does not lead to an additional objective. However, the return restriction imposes two extra conditions on the model, i.e., a person can leave his car at a certain node and ride with someone else, if and only if (i) he will pass through this node on a later ride, and (ii) he is able to return via ride-sharing. To this end, we introduce a new set $\mathcal{R}_p = \{1, 2, ..., n_p\}$, the elements of which are used as an indicator of the current ride of a person p. This enables us to cope with multiple rides per person in a certain period of time. We also introduce a new (dependent) binary variable C_{pri} to keep track of the location where a person's car is parked at the end of each ride. The constraints concerning the return restriction are written as follows.

$$
\begin{align}
& C_{pri} \leq \sum_{w=r+1}^{n_p} \sum_j Z_{pwij} && \forall p \in \mathcal{P}, r \in \mathcal{R}_p, i \in \mathcal{N} && (17) \\
& C_{pri} \leq 1 - \sum_j Z_{prij} && \forall p \in \mathcal{P}, r \in \mathcal{R}_p, i \in \mathcal{N} && (18) \\
& C_{pri} \geq \sum_j Z_{prji} - \sum_k Z_{prik} && \forall p \in \mathcal{P}, r \in \mathcal{R}_p, i \in \mathcal{N} && (19) \\
& C_{pri} \geq C_{p,r-1,i} - \sum_j Z_{prij} && \forall p \in \mathcal{P}, 2 \leq r \leq n_p, i \in \mathcal{N} && (20) \\
& C_{pri} \leq \sum_j Z_{prji} + C_{p,r-1,i} && \forall p \in \mathcal{P}, 2 \leq r \leq n_p, i \in \mathcal{N} && (21) \\
& C_{pri} \in \{0, 1\} && \forall p \in \mathcal{P}, r \in \mathcal{R}_p, i, j \in \mathcal{N} && (22)
\end{align}
$$

Constraints (17) construct the first condition of the return restriction mentioned at the beginning of Section 4.2, that is, if person p does not drive from node i to any other node on a later ride ($w > r$), then he cannot park his car at i on ride r. Constraints (18) state that if person p drives away from node i, then his car cannot be parked there. In contrast, if p drives to node i on ride r, but does not drive out of i on the same ride, then the car is parked at i at the end of ride r, which is ensured by Constraints (19). Constraints (20) synchronize the location of the car at the end of different rides. That is, the car is still parked at node i, if person p parks his car at i on ride r and carpools on the entire ride $r + 1$. Once the car is picked up from this node, Constraints (18) set the corresponding variable back to 0. Constraints (21) prevent a car being parked at node i at the end

of ride r, if the car was not parked there at the end of ride $r-1$ nor being driven to i on ride r. Constraints (19)-(21) determine where a car is parked at the end of a ride. With this information, we can model the return restriction condition (ii) that a person is able to pick up the parked car, otherwise the car cannot be parked. Constraints (22) are domain constraints for the newly defined variables.

4.3. Ride-sharing model with meeting points and return restriction (RS-M&R)

Based on the RS-M, and the constraints we added to address return restrictions, we provide the complete mathematical formulation for the RS-M&R in this subsection. As addressed later on, the complete formulation requires more changes besides adding the return constraints to the RS-M formulation.

$$
\begin{aligned}
\min \quad & \alpha_1 \sum_p \sum_r \sum_{(i,j)\in SP_{pr}} t_{ij} Z_{prij} + \alpha_2 \sum_p \sum_r (l_p - D_{p,r,d_{pr}}) \\
& + \alpha_3 \sum_p \sum_r \left((A_{p,r,d_{pr}} - D_{p,r,o_{pr}}) - \sum_{i,j} t_{ij} x_{prij} \right) \\
& + \alpha_4 \sum_p \sum_q \sum_r \sum_i S_{pqri}
\end{aligned} \tag{1'}
$$

s.t.

$$Y_{pprwij} = 0 \qquad \forall p \in \mathcal{P}; r, w \in \mathcal{R}_p; i, j \in \mathcal{N} \tag{2'}$$

$$\sum_q \sum_w Y_{pqrwij} + Z_{prij} = x_{prij} \qquad \forall p, r, i, j \tag{3'}$$

$$\sum_q \sum_w \sum_i Y_{pqrwij} + \sum_k Z_{prjk} \leq 1 + C_{p,r-1,j} \qquad \forall p \in \mathcal{P}; 2 \leq r \leq n_p; i, j \tag{4'}$$

$$\sum_q \sum_w \sum_i Y_{pqrwij} + \sum_k Z_{p,r+1,j,k} \leq 1 + C_{prj} \qquad j = d_{pr}, \forall p; 1 \leq r \leq n_p - 1; i, j \tag{23}$$

$$Y_{pqrwij} \leq Z_{qwij} \qquad \forall p, q \in \mathcal{P}; r \in \mathcal{R}_p; w \in \mathcal{R}_q; i, j \tag{5'}$$

$$S_{pqrj} \geq \sum_w \sum_i Y_{pqrwij} - \sum_w \sum_k Y_{pqrwjk} \qquad \forall p, q \in \mathcal{P}; r, j \tag{6'}$$

$$\sum_p \sum_r Y_{pqrwij} \leq v_q \qquad \forall q \in \mathcal{P}, w \in \mathcal{R}_q, i, j \tag{8'}$$

$$A_{prj} = D_{pri} + t_{ij} \qquad \forall p, r, i, j \tag{9'}$$

$$A_{pri} \leq l_{pr} \qquad i = d_{pr}, \forall p, r \tag{10'}$$

$$D_{pri} \geq e_{pr} \qquad i = o_{pr}, \forall p, r \tag{11'}$$

$$D_{pri} \geq A_{pri} \qquad \forall p, r, i \tag{12'}$$

$$D_{pri} - D_{qwi} \leq M(1 - Y_{pqrwij}) \qquad \forall p, q \in \mathcal{P}; r \in \mathcal{R}_p; w \in \mathcal{R}_q; i, j \tag{13'}$$

$$D_{pri} - D_{qwi} \geq -M(1 - Y_{pqrwij}) \qquad \forall p, q \in \mathcal{P}; r \in \mathcal{R}_p; w \in \mathcal{R}_q; i, j \tag{14'}$$

$$(A_{p,r,d_{pr}} - D_{p,r,o_{pr}}) - \sum_{(i,j)} t_{ij} x_{prij} \leq m_{pr} \qquad \forall p, r, (i,j) \in SP_{pr} \tag{15'}$$

$$C_{pri} \leq \sum_{w=r+1}^{n_p} \sum_j Z_{pwij} \qquad \forall p \in \mathcal{P}, 1 \leq r \leq n_p - 1, i \in \mathcal{N} \tag{17}$$

$$C_{pri} \leq 1 - \sum_j Z_{prij} \qquad \forall p, r, i \tag{18}$$

$$C_{pri} \geq \sum_j Z_{prji} - \sum_k Z_{prik} \qquad \forall p, r, i \tag{19}$$

$$C_{pri} \geq C_{p,r-1,i} - \sum_j Z_{prij} \qquad \forall p, r, i \tag{20}$$

$$C_{pri} \leq \sum_j Z_{prji} + C_{p,r-1,i} \qquad\qquad \forall p, 2 \leq r \leq n_p, i \qquad (21)$$

$$Z_{prij}, Y_{pqrwij}, S_{pqri}, C_{pri} \in \{0,1\} \qquad\qquad \forall p,q \in \mathcal{P}; r \in \mathcal{R}_p; w \in \mathcal{R}_q; i,j \in \mathcal{N} \qquad (24)$$

$$D_{pri}, A_{pri} \geq 0 \qquad\qquad \forall p \in \mathcal{P}, r \in \mathcal{R}_p, i \in \mathcal{N} \qquad (25)$$

The labels of the constraints link the constraints from RS-M&R with the constraints from RS-M, which we have already explained in Sections 4.1-4.2. Most of these constraints are simply modified by adding an additional dimension of r in the notation. Here, we draw the reader's attention to Constraints (4') and (23). Constraints (4') are modified based on Constraints (4) to cope with return restrictions. As described in Section 4.1, Constraints (4) are to prevent a person from driving once he leaves his car to join a carpool. With the return restrictions, however, it is possible for him to return to the location where he leaves his car and drive again. This situation is considered in Constraints (4'). Constraints (23) consider the boundary situation at the destinations. These constraints ensure that person p cannot carpool to his destination on ride r and leaves this node with his car on ride $r+1$, unless his car has been parked at this location.

5. Heuristic approach

When the number of persons and transfer points are large, the RS-M&R can become computationally prohibitive to solve. Moreover, when considering the return restrictions, it is not even an easy task to find a good solution to start with. In this section, we propose a constructive heuristic, which is based on the savings concept, to serve as a good starting point for improvement heuristics that can be used to solve large-scale problem instances of the RS-M&R that we may face in practice. For expository purposes, we assume that each of the participants has a car. However, this assumption can be easily relaxed to cover a more general setting that some of the participants do not have a car, see Agatz et al. (2011). Considering the flexibility that each participant has as being a driver or a rider, we make the following design choices to limit the number of role changes.

- Once a person is assigned to provide a ride to others, he will always be a driver.
- A person can leave his car at most once. In other words, once a person leaves his car (either at home or at a car parking point), he will be a rider until he picks up his car at the car parking point where he left his car. Then, he will become a driver and drive himself all the way to his home address.

The basic idea of the heuristic is to ensure the return restrictions of the persons who have the most saving potentials, by greedily assigning others who can share rides with them. The additional complexity from the return restrictions motivates us to consider the use of an external mobility service provider (denoted by EMS), such as taxi, as a backup option for taking the unmatched riders who have left their cars. Although this is an extension of the RS-M&R problem presented before, we can still compare the performance of the heuristic with the ILP from Section 4.3 by setting the cost of EMS sufficiently high. Sufficiently high in this setting means that a driver will use his own car when the ride-sharing plan includes the use of EMS. Hence, the resulting optimal solution to the relaxed problem is the same as the optimal solution to the RS-M&R.

The remainder of this section is organized as follows. In Section 5.1, we introduce the observations that are used to determine candidates for meeting points and car parking points. In Section 5.2, we propose the solution approach, the core of which is an extension of the set coverage problem.

5.1. Determining candidates for meeting points and car parking points

Meeting points are the nodes where two or more persons start a shared ride. We use the concept of the time interval at a node as the possible time frame of being present at the node, given the earliest departure time, the latest arrival time, and the assignment of the shared rides. Our approach for determining potential meeting points depends on the following observations.
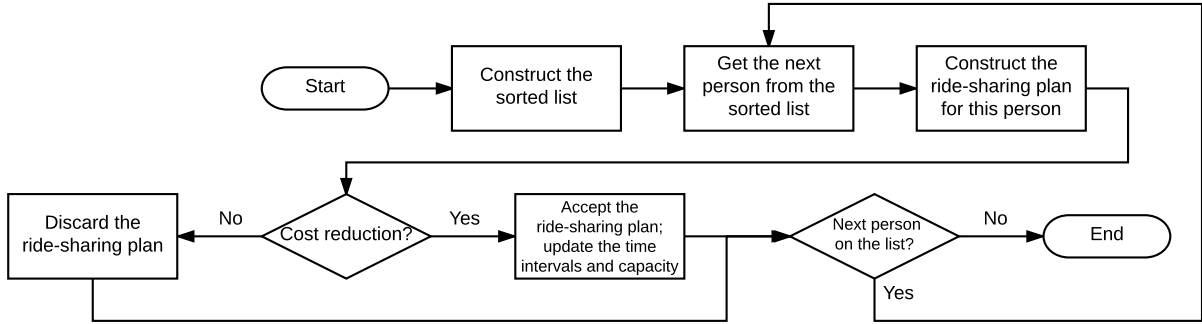
Figure 1   Overall structure of the heuristic

OBSERVATION 1. A node i can be a *potential* meeting point between persons p and q if and only if (i) the intersection of the time intervals of p and q at i is non-empty, and (ii) i is the starting point of a common arc in SP_{pr} and SP_{qw}.

On top of Observation 1, we add Observation 2 that extends the concept of meeting points to the multiple-rider matchings.

OBSERVATION 2. A node i can be a potential meeting point among a set of persons $Q \subseteq P$ if and only if i is a potential meeting point among subset of persons $Q' \subseteq Q$ for all $Q' \subseteq Q$.

Our approach for determining potential car parking points depends on the following observation that comes on top of Observations 1 and 2.

OBSERVATION 3. A potential meeting point i can be a *potential* car parking point for person p with r trips if and only if i is visited at least twice by person p.

According to Observation 3, a car parking point is a special type of meeting point. It only occurs when the role of a person is changed from a driver to a rider. Note that home addresses are also potential car parking points according to Observation 3. It is also important to point out that Observation 3 also holds for the relaxed problem if we consider EMS as special participant of the system.

5.2.   General procedure
The constructive heuristic is an iterative procedure. We construct a sorted list of participants according to their ride-sharing potential. In each iteration, we select a participant as a potential rider from the list, and determine the ride-sharing plan for this person. The resulting ride-sharing plan for a participant is accepted only if it leads to a reduction in total mobility cost. The overall structure of the heuristic is summarized in Figure 1. We will now further elaborate on the procedures of (i) constructing the sorted list, and (ii) iteratively constructing the ride-sharing plan for each participant.

5.3.   Constructing the sorted list
The purpose of the list is to select the most promising participant who might be able to satisfy the return restrictions only by riding with others. The procedure involves two steps. First, for each participant p, we calculate the initial uncovered distance of p, when he shares rides as a rider with each other participant q that acts as a driver. Second, we sort the participants by the *initial* minimum uncovered distance in non-descending order. Note that the list is sorted only once according to the initial minimum uncovered distance.

5.4.   Constructing the ride-sharing plan

As described in Figure 1, we construct the ride-sharing plan using an iterative procedure. In each iteration, we determine the ride-sharing plan for one participant using a two-phase procedure. In Phase 1, we solve a variant of the set cover problem, and obtain a feasible ride-sharing plan by using EMS to cover the uncovered distance in the solution. In Phase 2, we improve the ride-sharing plan using a subroutine, the goal of which is to reduce the use of EMS in the plan.

5.4.1.   Phase 1: dynamic greedy cover algorithm We start with a formal definition of the set cover problem, and explain how to model the ride-sharing problem as a variant of it. In particular, we address the main differences between them.

The *set cover problem* is defined as follows. Let $\mathcal{U}$ be a set of n elements $x_1, x_2, ..., x_n$, and $\mathcal{S} = \{S_1, S_2, ..., S_m\}$ be a collection of m subsets of $\mathcal{U}$ such that $\bigcup_i S_i = \mathcal{U}$. The goal is to select as few subsets as possible from $\mathcal{S}$ such that their union covers $\mathcal{U}$. The literature proposes the greedy cover algorithm that chooses sets according to one rule: at each iteration, choose the set that contains the largest number of uncovered elements. It has been shown that the greedy cover algorithm is the best-possible polynomial time approximation algorithm for the set cover problems (Feige 1998).

In our problem context, the set of arcs that p can ride with participant q (concerning spatiality, time, and capacity) is defined as a subset S_i, and the universe set for p refers to the set of all arcs that can be covered by his peers, which is a subset of $\bigcup_{r \in \mathcal{R}_p} SP_{pr}$. It is important to note that a full cover might not be possible even by selecting all the subsets. This is due to the fact that the subsets in our problem, i.e., the arcs where p can ride with q for different q are not independent. The selection of one subset may lead to infeasibility of other subsets. Thus, the goal of our problem is to select as few participants as possible to be drivers to cover the longest possible distance for all $r \in \mathcal{R}_p$. Having said that, the greedy cover algorithm, that is to pick the subset that covers the maximum uncovered distance until no further improvement can be achieved, is still applicable to our problem. The second difference is a continuation of the first one. Since the time interval of a participant on an arc can be influenced by the previous assignments, our problem is a dynamic set cover problem in the sense that the set of feasible arcs of a participant that can be used to cover p evolves over the greedy iterations. Consequently, we need to update the subsets in each iteration. The third difference is that an additional step (line 6 of Algorithm 1) is required to construct a feasible ride-sharing plan. To summarize, the following procedure allows us to determine a feasible ride-sharing plan for the current participant as a potential rider, which is used as the input of the EMS reduction subroutine in Phase 2.

Algorithm 1 The dynamic greedy cover algorithm

1: **while** uncovered distance that can be covered by participants > 0 **do**
2:     Pick the set that covers the maximum uncovered distance
3:     Mark elements in the chosen set as covered
4:     Update the time intervals and capacity of the participants
5: **end while**
6: Use EMS to cover all the uncovered arcs

5.4.2.   Phase 2: EMS reduction subroutine It is unlikely that the potential riders can be fully covered by their peers. Thus, the solution constructed by the dynamic greedy cover algorithm could include the use of EMS. Therefore, we use a subroutine to minimize the use of EMS in the ride-sharing plan of the current potential rider by self-driving between a potential car parking point and his home address. The key idea is to determine the actual car parking point so as to

---

Algorithm 2 The EMS reduction subroutine

 1: Find the starting node v_s of the covered arcs by participates
 2: Find the ending node v_e of the covered arcs by participates
 3: **if** both v_s and v_e are potential car parking points **then**
 4:     **if** $v_s = v_e$ **then**
 5:         Set it as the actual car parking point of the current potential rider
 6:     **else if** $v_s \neq v_e$ **then**
 7:         Remove the covered arc(s) between them from the ride-sharing plan
 8:         Set the node that is further away from home address as the actual car parking point
 9:     **end if**
10: **else**
11:     Find the closest potential car parking point
12:     Set it as the actual car parking point of the current potential rider
13: **end if**
14: update the ride-sharing plan with self-driving between home address and the car parking point

---

minimize the total distance from this point to where he starts/stops riding with his peer. Algorithm 2 outlines the procedure of Phase 2.

Although the proposed solution approach evaluates the potential matches only based on mobility costs, other objectives are considered implicitly. As we greedily assign drivers to the potential rider based on the maximum distance each driver can cover, it potentially helps to reduce the number of transfers of the potential rider. To reduce the total deviation from the latest arrival time at the destinations, we set the departure time of each person at his latest possible departure time within the time interval of a feasible match. Although we are not considering it in our experiment, post-processing may be used to improve the objective values of the total waiting time during the transfer and the total deviation from the latest arrival time at the destinations by re-optimizing the time schedule of the given ride-sharing plan.

6. An illustrative case: Significant BV

In this section, we apply the RS-M&R model on a case based on the commuter/business traffic of Significant BV, a consulting company that initiated this research. To this end, we use the rides driven by Significant personnel to evaluate the ride-sharing potential for this company. In Section 6.1, we present the case and discuss the characteristics of the data. Section 6.2 reports the results of applying the RS-M&R model to the rides driven by Significant personnel in March 2014 (weekends exclusive).

6.1. Case description

Significant BV is a Dutch consultancy and research firm located in Barneveld, advising issues of organizational, operational and procurement in the public domain. Currently, the company employs 42 consultants and researchers, all of whom are equipped with a company car. Due to the characteristics of the consultancy work, they are quite flexible in where and when they work. As a result, almost all employees travel by car. To reduce the costs and greenhouse gas emissions of the commuter/business traffic, Significant considers to reduce the number of single-person trips by car via ride-sharing among the colleagues. However, employees lack insights into their colleagues' whereabouts, and thus ride-sharing only happens accidentally for joint meetings. Our proposed automated matching methodology could offer a solution for this.

As input data, we collect data of Significant-related rides of all the employees in a period of four weeks via some prepared forms filled by the employees in order to keep track of their rides. The

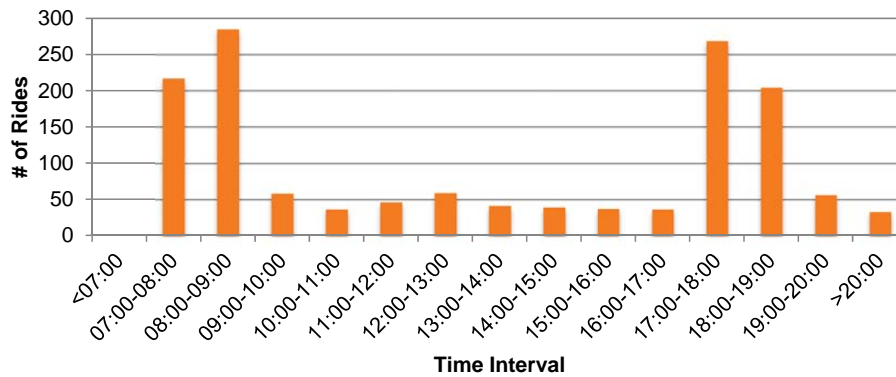**Figure 2**   **Origins and destinations of observed trips**

Figure 3   **The distribution of the persons over the time intervals**

rides are defined on a city level; time schedules of the rides are defined on an hourly basis. The rides within the same city are excluded from the set of rides we use for the experiments. During these four weeks, 1416 rides are considered. 87% of the these rides are from/to a home address; the other 13% consists of the rides between the office in Barneveld and a client, and those between two clients. 78 locations are visited in total. Figure 2 presents an overview of these locations; the size of a circle indicates the frequency of visits of this location. For instance, Barneveld is visited 430 times, Utrecht is visited 242 times, and Amsterdam is visited 97 times during the four weeks. Figure 3 shows the number of rides in different time intervals. Most of the rides are driven in early mornings or late afternoons, which is consistent with the high percentage of rides from/to home addresses. The main purpose of this illustrative case is to show the ride-sharing potential, even on such a small scale with only 42 participants, and 71 trips per day on average.

6.1.1.   Network and shortest paths The network used in this case study consists of the 78 cities shown in Figure 2 and the 358 carpool parking lots in the Netherlands[1]. The arcs between each node pair represent the travel route chosen by Google Maps under the criteria of shortest

---

[1] We retrieve the names and the coordinates of these parking lots from `carpoolplein.nl`

Table 2   Parameters of the Significant case

Parameter	Value
Max. waiting time at a transfer point	10 min
Max. time deviation from the latest arrival time	60 min
Capacity of passenger seats	3
Weight of the vehicle-miles α_1	0.5
Weight of the time deviation α_2	0.01
Weight of the waiting time α_3	1
Weight of the number of transfers α_4	0.000001
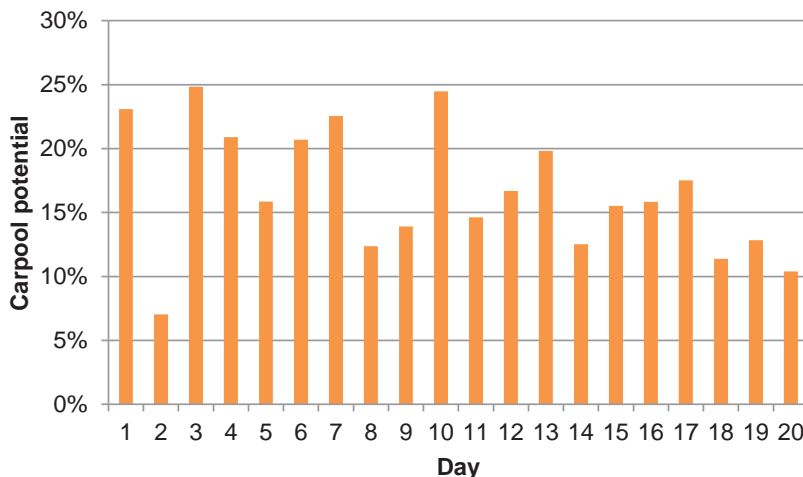
Figure 4   Percentage saved of the total driving distance by ride-sharing

driving time. To collect the travel time between each node pair, Google Maps API and Bing Maps API are used. The Google Maps API is used to retrieve accurate coordinates of a location. Bing Maps API is used to retrieve the travel time information for the 189,660 arcs.

Although most of the carpool parking lots are easily accessible from highways, a limited detour is still necessary. Therefore, every carpool parking lot that results in less than 5 minutes of additional driving distance if it is inserted into the route is assumed to be located on the shortest path. Similarly, we also assume that every city that leads to less than 10 minutes of additional driving distance if it is inserted into the route is located on the shortest path. The visiting order of the nodes along the shortest path is determined in a non-decreasing order according to the travel time between the origin and each node. This automated process of assigning nodes to the shortest paths may lead to undesirable outcomes that need to be adjusted (see Appendix for detailed illustrations).

6.2. Experiments

In this subsection, we show the results of applying the RS-M&R model to the rides driven by Significant personnel in a period of four weeks. The value of the parameters are selected based on interviews with employees at Significant, which are summarized in Table 2.

Figure 4 shows that by sharing the rides, Significant employees could have saved 7-25% of the distances driven. The optimal solutions also show that among the 1416 collected rides, 511 can be saved by ride-sharing, of which 506 rides are shared between 2 persons (i.e., one driver and one passenger), 2 rides are shared between 3 persons, and 3 rides are shared between 4 persons. In other words, two-person trips account for the vast majority of the shared rides. In addition, no one needs to wait during the trip according to the optimal ride-sharing plan.

Although the user group is very small, and the origins and destinations are rather scattered across the Netherlands, the results show that Significant can still benefit from ride-sharing by reducing the distance traveled by 7-25%.

In the next section, the RS-M&R model is tested through extensive computational experiments on virtual networks. We aim at providing valuable insights into successfully implementing a ride-sharing service with the consideration of meeting points and return restrictions.

7. Experimental settings

There are three factors that affect the possibility of ride-sharing: the spatial network, the hotspot density, and the participation density. Therefore, we design our problem instances to reflect these factors in three different settings, including the basic setting presented in Section 7.1, the high transfer point density setting (denoted by HTP setting) described in Section 7.2, and the high business traffic density setting (denoted by HBT setting) described in Section 7.3. Section 7.4 explains the method we use to compare the algorithmic performance, followed by the experimental setting for investigating the benefits of using EMS in Section 7.5. The performance indicators are described in Section 7.6.

7.1. Basic setting

Two different spatial distributed sets with 51 nodes generated from Solomon's 50 customers problems instances are considered: the scattered set, and the clustered set. In particular, R101 is used to generate the scattered set, while C101 is used to generate the clustered set. Assuming that we are facing a single-workplace problem, we retain the coordinates of the depot, which becomes the workplace, and of the customers, which are considered as locations for the cities and carpool places. Note that only cities are considered as potential home addresses and origins and destinations for business trips, while carpool places are only used for transfers. We multiply the coordinates of these nodes by 3, resulting in an area of 210×210 kilometers (roughly the size of the Netherlands). The networks are generated based on the minimum spanning tree concept. In order to better reflect the connectivity of the Dutch intercity road network, we directly connect those two nodes whose ratio between their Euclidean distance and the shortest distance on the current network is above a predetermined threshold (0.3 for the scattered set and 0.5 for the clustered set). The resulting graphs are depicted in Figure 5, where the arcs in red constitute the minimum spanning tree and the arcs in blue are added for better connectivity. The Euclidean distance is used to calculate the distances between the connected nodes. The average speed of the drivers is assumed to be 60km/h.

We consider three types of city nodes, the main difference of which is the associated probability of being an origin or a destination of a trip. Since all the trips belong to the commuter/business traffic of a company, the probability of the workplace being either the origin or the destination of these trips should be higher. We set it to 0.3 based on the data from Significant. The spatial distribution of the origins and destinations of individual trips depends on the socio-economic status of the cities. For instance, when there are no significant demographic and socio-economic differences within the region, it is possible to assume that the probability does not vary with the remaining 50 nodes. On the other hand, if such differences exist, some of the nodes may become hotspots for residential and business purposes, such as Utrecht, Amsterdam and Den Haag in the Significant case. It is natural to assign a higher probability to these city nodes. The number of hot spots are assumed to be 1, 2, or 3 and the aggregate probability of any of the hot spots being the origin or the destination of a trip is 0.15. That is, with one hot spot, the probability of it being chosen is 0.15; with two hot spots, the probability of either being chosen is 0.075; with three hot spots, the probability of any being chosen is 0.05. These observations lead to the following spatial distributions of trips:
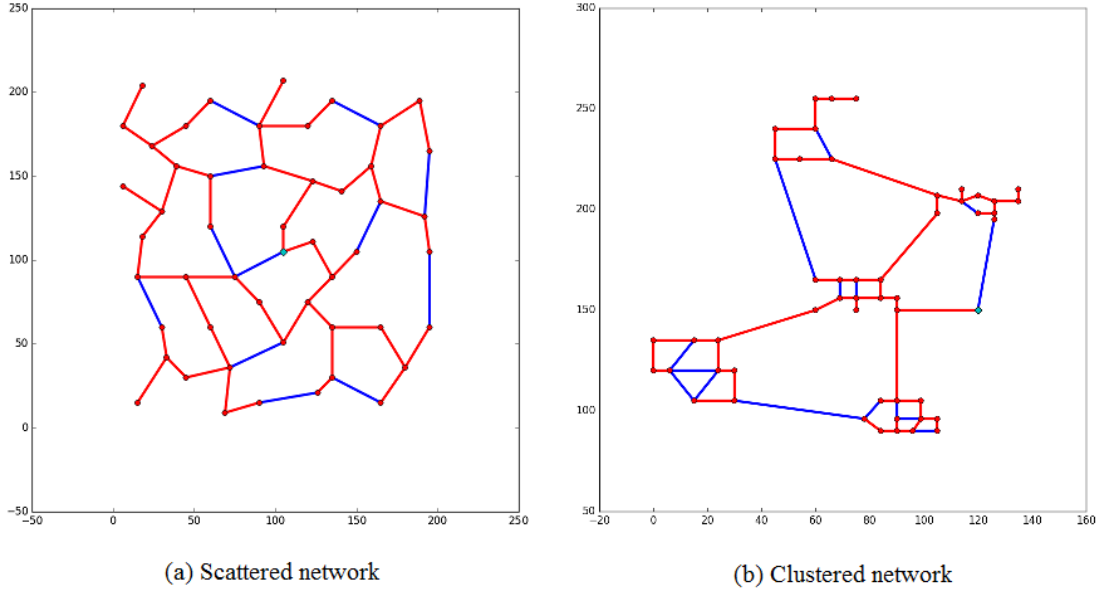
(a) Scattered network   (b) Clustered network

Figure 5   Networks

- SU: Scatteredly distributed city/transfer nodes, uniformly distributed origins and destinations. We use the following probability function for this setting.

$$\Pr(i = n_{pr}) = \begin{cases} 0.3 & i \text{ is the workplace} \\ \frac{0.7}{50} & \text{otherwise} \end{cases}$$

where n_{pr} represents an end point of participant p's ride r, either the origin o_{pr} or the destination d_{pr}.

- SH: Scatteredly distributed city/transfer nodes with origin and destination hot spots. We use the following probability function for this setting.

$$\Pr(i = n_{pr}) = \begin{cases} 0.3 & i \text{ is the workplace} \\ \frac{0.15}{h} & i \text{ is a hot spot} \\ \frac{0.55}{50-h} & \text{otherwise} \end{cases}$$

where h denotes the number of hotspots in an instance. We consider the settings with 1, 2 and 3 hotspots, which are denoted by SH1, SH2 and SH3, respectively.

- CU: Clusteredly distributed city/transfer nodes, uniformly distributed origins and destinations). We use the following probability function for this setting.

In total, we have 200 instances (20 scenarios, each including 10 instances) in the basic setting. Each scenario has a certain number of users, and a certain spatial distribution for the trips.

For each person, the number of trips was set to 2, 3 or 4 with probability 0.7, 0.2 and 0.1. The first and the last rides of each person belong to the commuter traffic. The origins and destinations of the trips, except for the destination of the last trip, are generated based on the aforementioned probabilities in different settings. The destination of the last trip is the same as the origin of the first trip, representing the user's home location. Latest arrival times l_{pr} for $r = 1$ are uniformly distributed between 480 and 540 (representing the time window between 8am and 9am); earliest departure times e_{pr} for $r = r_{max}$ are uniformly distributed between 1020 and 1140 (representing the time window between 5pm and 7pm). Latest arrival times of the remaining trips are uniformly distributed between 540 and 960 (representing the time window between 9am and 4pm). Earliest

arrival times of these trips are computed by subtracting a time slack from the corresponding latest arrival times. The time slack is set to 30. We assume that each person has a company car with capacity of 3 passenger seats. The weights α_1, α_2, α_3, and α_4 are the same as in the Significant case.

7.2. Effect of transfer point density
In order to study the effect of transfer point density, we compare the basic setting with the HTP setting. The HTP setting contains half as many cities but the same number of transfer points as in the basic setting. The cities in the HTP setting form a subset of the transfer points. Only cities can be considered as locations for the origin or the destination of a trip, but both cities and transfer points can be used for transfers. Thus, a smaller number of cities represents an environment in which more trips may share the same origin or (and) destination. To allow for a fair comparison, the relative probability of a trip starting from (ending at) a hotspot as compared to a regular city is the same as their counterpart in the basic setting. The set of 25 cities are randomly selected from the 50 nodes in each problem instance.

7.3. Effect of business traffic density
To access the impact of the business traffic density on the performance of the ride-sharing system, we compare the basic setting with the HBT setting. In the HBT setting, the number of trips of a participant is randomly generated among 2, 3 or 4 with equal probabilities. Since two trips are always required for the commuter traffic, it implies that the number of business trips is 0, 1, or 2.

7.4. Algorithmic comparison
We compare the computational performance of the proposed heuristic with different trip patterns and different numbers of participants. To show the efficiency of the proposed heuristic, we also consider a *naive* heuristic for comparison. This naive heuristic selects the current potential rider, and assigns drivers to cover the potential riders in a random manner, followed by Phase 2 of the proposed heuristic. To allow for a fair comparison with the ILP, we set the cost of EMS prohibitively high ($c_e = 1000$) to avoid any use of it in the heuristics.

7.5. Benefits of the potential use of EMS
The return restriction may significantly reduce the success rate of ride-sharing. It is likely that a participant has to travel with his own car because only a small part of his entire journey cannot be covered by any other driver. In this case, using EMS for the small part of his journey may substantially reduce the total cost of the system. Thus, we want to investigate the potential benefit of allowing the utilization of EMS as a backup option for taking the unmatched riders who have left their cars. To this end, we evaluate the mobility cost $c_e = 2$, which resembles the cost of using a taxi in the Netherlands, and $c_e = 0.5$, which resembles a situation when the community hires a mobility service provider to handle all the mobility demands. We also consider an intermediate cost setting $c_e = 1.25$, which aims at representing the cost of using a contracted EMS as an emergency backup. We only use the heuristic to solve these problem instances.

7.6. Performance indicators
We evaluate and compare the solutions using the following metrics: (1) the percentage mileage savings, (2) the travel time increase, i.e., the average waiting time per person during the entire day, and (3) the time deviation from the latest arrival time, i.e., the average time difference (per person per ride) between a person's arrival time and the latest arrival time at the destination. In addition, we evaluate the efficiency of ride-sharing using the following indicators: (4) the percentage car savings, i.e., the number of saved cars as a fraction of the total number of participants, (5) the matching rate, i.e., the vehicle-miles of the shared rides as a fraction of vehicle-miles when

ride-sharing is implemented, (6) the average number of riders in a shared ride, and (7) the average number of transfers needed for a matched rider to reach his destination. Note that we do not follow the commonly used definition of the matching rate, namely, the fraction of participants that are matched. This is because with the permission of transfers, a higher number of shared rides does not necessarily indicate a better performance of the ride-sharing system. It is important to note that, for the matching rate, the mileage savings, and the average number of riders in a shared ride, a higher value is preferred. In contrast, for the time deviation and the average number of transfers, a lower value is preferred.

To conclude, Table 3 provides a summary of the three settings that are described in this section. Given these settings, we perform a number of experiments as shown in Table 4.

Table 3   **Instance settings**

Basic setting	
#cities:	50
probability density of #trips:	$\Pr(r=2)=0.7$, $\Pr(r=3)=0.2$, $\Pr(r=4)=0.1$
trip distribution:	SU, SH3, SH2, SH1, CU
#participants:	40, 60, 80, 100
HTP setting	
#cities:	25
probability density of #trips:	$\Pr(r=2)=0.7$, $\Pr(r=3)=0.2$, $\Pr(r=4)=0.1$
trip pattern:	SU, SH3, SH2, SH1, CU
#participants:	40, 60, 80, 100
HBP setting	
#cities:	50
probability density of #trips:	$\Pr(r=2)=\Pr(r=3)=\Pr(r=4)=\frac{1}{3}$
trip distribution:	SU, SH3, SH2, SH1, CU
#participants:	40, 60, 80, 100

Table 4   **Characteristics of instances in each experiment**

Experiment	Setting	c_e	Solution method
Effect of trip patterns and #participants	Basic	1000	ILP
Effect of transfer point density	HTP versus Basic	1000	ILP
Effect of business traffic density	HBT versus Basic	1000	ILP
Heuristic performance	HTP and Basic	1000	ILP, Greedy, Naive
Benefit of using EMS	Basic	2, 1.25, 0.5	Greedy

8.   Numerical results

Test instances are performed on an Intel Core i5-5200U 2.20GHz, 8 GB RAM computer. The standard CPLEX 12.40 MIP solver in AIMMS is used for the ILP model, and the heuristic is implemented in Java. In our results, we report averages over 10 instances. The corresponding ILPs run with a time limit of 2 hours using AIMMS's default parameter settings. In most cases, the ILP finds the optimal solution; only in a few cases the solver reached the time limit and terminated with a gap of at most 1.5%.

Section 8.1 studies the features of the RS-M&R, based on the optimal solutions obtained by the ILP in the basic setting. Sections 8.2 and 8.3 report the results of the HTP setting and the HBT

setting, respectively, as compared to the results of the basic setting. In Section 8.4, we report the performance of the proposed heuristics. In Section 8.5, we report the benefits of introducing an EMS into the ride-sharing system.

8.1. Effects of trip patterns and participation density

In this section, we study the effect of the number of participants in the system and their trip patterns on the system performance. Our experiments show that the average waiting time per person is negligible, less than one minute in all the tested scenarios. Thus, only the results of the other six indicators are presented.

Table 5 shows the percentage mileage savings (denoted by MileSav), the time deviation from the latest arrival time (denoted by TimeDev), and the percentage car savings (denoted by CarSav) in the basic setting. As expected, the MileSav is higher when more people participate. Comparing the cases (i.e., SU, SH3, SH2, and SH1) in the scattered network, the MileSav is higher when the origins and the destinations of the trips are more concentrated. The reason is that it is more likely for two or even more persons to share a ride if they have the same origin or destination. The CU case offers the highest MileSav across all cases. Although the origins and destinations are evenly distributed across the cities, the clusters in the clustered network provide a natural concentration of them. Since the road connections between clusters are limited, it results in an increased number of shared rides among the trips between different clusters. We also see that ride-sharing creates a promising opportunity to reduce the number of cars needed to satisfy the mobility demand (positive values of CarSav). This opportunity increases when more people participate.

Table 5   Basic setting: results of the ILP with varying number of participants and trip patterns

	MileSav (%)				TimeDev (minutes)				CarSav (%)			
$\rho =$	40	60	80	100	40	60	80	100	40	60	80	100
SU	8.88	11.92	15.33	18.38	2.61	3.45	4.45	5.12	6.75	10.33	13.63	14.40
SH3	8.76	14.33	16.48	19.99	2.46	3.86	4.63	5.43	5.00	10.33	12.63	15.30
SH2	9.55	14.37	17.67	20.62	2.87	4.20	4.74	5.87	7.25	10.67	14.13	16.80
SH1	12.41	16.28	21.39	23.00	3.38	4.11	5.45	5.49	9.50	11.00	16.00	19.40
CU	16.46	20.09	23.76	26.88	4.16	4.88	5.93	6.35	7.50	6.67	13.50	25.80
Average	11.21	15.40	18.93	21.77	3.10	4.10	5.04	5.65	7.20	9.80	13.98	18.34

Note: ρ represents the number of participants.

We observe that the TimeDev fluctuates in different trip patterns, because it is more closely related to the earliest departure times and the latest arrival times of the persons who are assigned to share a ride. However, we see ascending TimeDev with increasing number of participants. Such an increase in TimeDev results from an increased dependency of the participants. A feasible matching requires the time coordination of all the corresponding participants, and thus some matched participants might need to execute their trips earlier so as to accommodate others.

Table 6 presents the results of the matching rate, the average number of riders and the average number of transfers. In general, these indicators increase as the number of participants increase across all trip patterns. The performance of the SU, SH3 and SH2 cases are similar to each other, the relative order of which fluctuates slightly. As compared to the other three trip patterns in the scattered network, SH1 results in a higher value of the matching rate and the average number of riders, but no clear difference in the number of transfers. Given the similarity and differences between CU and SH1, our discussion focuses on the comparison between them in terms of these three indicators.

Table 6   **Basic setting: ride-sharing efficiency with varying number of participants and trip patterns**

	Matching rate (%)				Avg #riders				Avg #transfers			
$\rho =$	40	60	80	100	40	60	80	100	40	60	80	100
SU	9.81	13.63	18.23	22.68	1.27	1.40	1.54	1.53	0.21	0.28	0.37	0.39
SH3	9.75	16.92	19.88	25.22	1.28	1.39	1.48	1.52	0.22	0.31	0.36	0.44
SH2	10.61	16.88	21.65	26.25	1.32	1.44	1.51	1.57	0.27	0.29	0.33	0.39
SH1	14.33	19.63	27.55	30.23	1.48	1.60	1.67	1.65	0.22	0.34	0.36	0.39
CU	19.88	25.16	31.18	36.85	1.39	1.54	1.65	1.69	0.35	0.36	0.49	0.47

We find that a higher matching rate is obtained when the concentration of the trips is higher. For a given number of participants, the matching rate in CU is higher than in SH1. On the other hand, SH1 results in a higher number of riders in a shared ride, especially when the number of participants is not too high (i.e., 40, 60, and 80). As compared to the scattered network, the clustered network requires a higher number of transfers.

8.2.   Effect of transfer point density

In this section, we report the impact of the number of transfer points on the system performance. Table 7 presents the results of MileSav, TimeDev, and CarSav, measured as an absolute difference from the basic setting where all the transfer points are considered as cities. We find that the benefit of introducing more transfer points, from the perspective of MileSav, increases in the number of participants. We also find that the HTP setting results in at most 1.05 minutes increase in TD. In general, the HTP setting also leads to larger reductions in CarSav.

Table 7   **Effect of transfer point density regarding MileSav, TimeDev and CarSav.**

	MileSav (%)				TimeDev (minutes)				CarSav (%)			
$\rho =$	40	60	80	100	40	60	80	100	40	60	80	100
SU	2.35	5.09	5.90	5.90	0.26	0.98	0.77	1.05	2.50	1.67	3.50	5.50
SH3	3.42	3.83	5.77	5.51	0.76	0.41	0.22	0.75	3.00	4.00	7.00	7.80
SH2	5.77	5.89	8.04	6.20	0.29	0.29	0.8	0.53	4.50	5.83	6.50	6.40
SH1	4.41	6.55	3.42	7.28	0.07	0.68	-0.11	1.27	8.00	11.83	10.25	9.40
CU	-0.22	5.39	4.73	8.28	-0.22	1.01	0.10	1.03	3.75	7.67	3.75	-4.50

Note: The results are measured as an absolute difference from the base case scenarios.

The results on the matching rate, the average number of riders, and the average number of transfers are presented in Table 8, measured as the absolute difference from the basic setting. We find that the matching rate is more sensitive to the change of the number of participants in the high transfer point density scenarios. This is shown by the fact that the difference in matching rate increases in the number of participants. We also find that the average number of riders in a shared ride is more sensitive to both the number of participants and different trip patterns. Unlike the basic setting, the performance of SU, SH3 and SH2 are much more distinguishable when the density of the transfer points is higher. More surprising is the fact that, on average, slightly less transfers are needed in the high transfer point density environment. This demonstrates the fact that the average number of transfers highly depends on the network (nodes and arcs), and less on the origin and the destination of the trips. The results point to the fact that the increase in the matching rate is mainly due to the increasing number of shared rides with multiple riders. In other words, more concentrated origins and destinations create more shared rides with multiple riders.

Table 8   **Effect of transfer point density regarding matching rate, average #riders and average #transfers.**

	Matching rate (%)				Avg #riders				Avg #transfers			
	$\rho =$40	60	80	100	40	60	80	100	40	60	80	100
SU	2.98	7.14	8.99	9.72	-0.02	0.03	0.05	0.13	-0.08	-0.05	-0.04	0.01
SH3	4.35	5.54	8.94	9.02	0.06	0.12	0.13	0.11	-0.01	-0.02	-0.05	-0.10
SH2	7.73	8.72	13.18	10.58	0.14	0.21	0.16	0.21	-0.05	0.00	0.03	0.02
SH1	6.27	10.49	5.71	13.57	-0.04	0.18	0.17	0.22	-0.07	-0.12	-0.08	-0.07
CU	-0.01	9.70	9.07	12.43	0.06	0.07	0.10	0.17	-0.05	0.08	-0.04	0.04

Table 9   **Effects of business traffic density**

	SU			SH1			CU		
	Base	High	% diff	Base	High	% diff	Base	High	% diff
System									
Mileage savings(%)	11.92	7.50	-37.08	16.28	10.06	-38.20	20.09	16.26	-19.06
Matching rate(%)	13.63	8.15	-40.21	19.63	11.24	-44.78	25.16	19.53	-22.38
Avg #riders	1.40	1.20	-14.29	1.60	1.35	-15.63	1.54	1.39	-9.74
Time deviation	3.45	2.18	-36.81	4.11	2.86	-30.41	4.88	3.88	-20.49
Avg #transfers	0.28	0.18	-35.71	0.34	0.20	-41.18	0.36	0.26	-27.78
Commuter									
Mileage savings(%)	13.70	8.58	-40.36	18.36	11.22	-38.89	23.27	18.43	-20.78
Matching rate(%)	15.97	9.45	-37.37	22.75	12.70	-44.18	30.36	22.83	-24.80
Avg #riders	1.11	0.76	-31.53	1.20	0.75	-37.50	1.20	0.80	-33.33
Time deviation	3.95	2.26	-42.78	4.65	2.99	-35.70	5.48	3.93	-28.28
Avg #transfers	0.29	0.23	-20.69	0.36	0.25	-30.56	0.38	0.29	-23.68
Business									
Mileage savings(%)	2.53	5.42	114.23	4.40	7.45	69.32	3.83	11.97	212.53
Matching rate(%)	2.62	5.76	119.85	4.81	8.25	71.52	4.05	13.69	238.02
Avg #riders	0.15	0.25	66.67	0.28	0.40	42.86	0.22	0.35	59.09
Time deviation	0.97	2.03	109.28	1.31	2.54	93.89	1.83	3.73	103.83
Avg #transfers	0.05	0.08	60.00	–	0.08	–	0.08	0.16	100.00

Note: The results are based on the scenarios of 60 participants in the base case.

8.3.   Effect of business traffic density

In this section, we report on the effect of the business traffic density on the performance of the system. We evaluate the indicators (i) from the perspectives of the system, (ii) for the commuter trips only, and (iii) for the business trips only. Only the results for the scenario of 60 participants in SU, SH1 and CU are presented, because the results of SU and SH1 provide the lower and upper bounds of the results of SH2 and SH3, respectively. As explained in Section 8.1.1, this is due to the degree of concentration in the origins and the destinations of the trips in the scattered network. The results are shown in Table 9. Since the total number of business trips are 150% higher in the HBT setting, the number of saved trips is divided by the total number of rides to get a fair comparison.

As expected, the business traffic density has a positive impact on the mileage savings, the matching rate, and the average number of riders for the business traffic. More surprising is the fact that the increase of business trip density has a negative impact on the mileage savings, the matching rate and the car occupancy in the commuter traffic. To some extent, this result may be a consequence of the choice of objective hierarchy: minimizing the waiting time is twice as important as minimizing the vehicle-miles. In the basic setting, due to the scarceness of the business trips, we see a larger percentage of mixture between commuter traffic and business traffic within a shared ride, especially between the early return commuter traffic and the late business traffic. The potential drawbacks of these mixed matchings are the increase in waiting time and TimeDev. When

Table 10   **Summary of the run time**

| Trip patterns | Number of participants | Run time (sec) | | | | | |
| | | 50 cities | | | 25 cities | | |
		ILP	Greedy	Naive	ILP	Greedy	Naive
SU	40	9.42	0.27	0.18	11.57	0.25	0.16
	60	20.12	0.35	0.25	23.52	0.32	0.28
	80	38.49	0.42	0.32	189.80	0.86	0.34
	100	636.13	0.52	0.36	2044.89	1.07	0.38
SH3	40	14.02	0.25	0.18	9.33	0.27	0.18
	60	18.20	0.33	0.27	28.59	0.29	0.27
	80	52.41	0.42	0.32	241.66	0.35	0.40
	100	360.63	0.45	0.41	2460.12	0.44	0.38
SH2	40	13.44	0.26	0.18	9.54	0.22	0.20
	60	17.36	0.33	0.26	26.08	0.34	0.33
	80	45.77	0.43	0.30	929.99	0.38	0.34
	100	301.19	0.47	0.41	4178.31	0.49	0.39
SH1	40	11.42	0.26	0.20	10.59	0.46	0.24
	60	23.34	0.34	0.29	61.20	0.45	0.32
	80	249.93	0.40	0.30	2617.77	0.74	0.42
	100	3450.78	0.47	0.39	7200.00	0.54	0.36
CU	40	11.35	0.23	0.18	13.02	0.20	0.18
	60	55.62	0.32	0.28	148.10	0.24	0.26
	80	944.83	0.36	0.31	3882.78	0.36	0.33
	100	4004.24	0.44	0.41	7200.00	0.48	0.38

the number of business trips increases, it is more likely to be better off by sharing rides within the business traffic. Implicitly, it results in a reduction of the participants that the commuter traffic can share a ride with, and thus leads to reductions in the MileSav, the matching rate and the car occupancy. Given that the commuter traffic is dominant over the business traffic in the basic setting, the diminishing commuter traffic density in the HBT setting also worsens these indicators from the system's standpoint. This is because the improvement gained from the business traffic is not enough to offset the rollback from the commuter traffic from the vehicle-miles perspective.

These findings stress the importance of having a high participation from the viewpoint of the two sub-systems, i.e., the commuter traffic and the business traffic. Similar conclusions can be observed in the scenarios with 40, 80, and 100 participants.

8.4. Heuristics performance

In this subsection, we report the algorithmic performance with different trip patterns and different number of participants. Within a time limit of 2 hours using AIMMS's default parameter settings, all the instances with 40 and 60 persons can be solved to optimality. However, in 9 instances across the scenarios with 80 participants and 35 instances across the scenarios with 100 participants, the solver reached the time limit and terminated with a gap of no more than 1.5%. The run time of the ILP and the heuristics in solving the different scenarios is summarized in Table 10.

We see that the run time in solving the ILP increases super-linearly with respect to the number of participants. The scenarios with a higher concentration of origins and destinations, i.e., the HTP settings and CU and SH1 in general, are much more difficult to solve. The run time of these scenarios are also more sensitive to an increase in the number of participants. In contrast, the run time of the greedy heuristic increases linearly with respect to the number of persons. We also find that the run time is insensitive to the change of trip patterns. As expected, the naive heuristic is faster, as compared to the greedy heuristic.
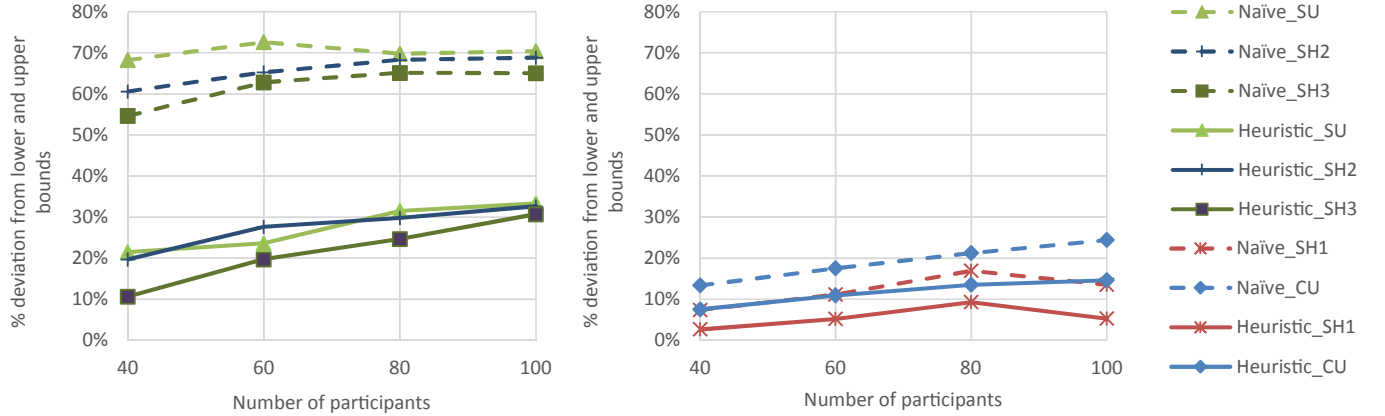
Figure 6   **Performance of the proposed heuristic**

Figure 6 presents the performance of the greedy heuristic and the naive heuristic for the basic setting, as compared to the solutions obtained by the ILP and the benchmark solutions obtained by the situation where ride-sharing is not implemented. We assign a value 0 to the ILP solutions and a value 1 to the solutions with no ride-sharing. The solutions obtained by the two heuristics are scaled between 0 and 1. These values are calculated proportionally based on their relative distance away from the corresponding optimal and benchmark solutions. Therefore, the closer to zero, the better the performance. Overall, the greedy heuristic outperforms the naive heuristic consistently. We observe that the performance in the SH1 and CU cases is better than in the other trip patterns. This can be explained by the fact that the higher concentration of origins and destinations in SH1 and CU provides more ride-sharing alternatives, the performance of which is not too far away from the optimum. Thus, it is more likely for the proposed greedy heuristic to find one of these reasonably good solutions. When the trip pattern is more scattered (i.e., SU, SH3, SH2), both the greedy heuristic and the naive heuristic perform less, but the performance of the greedy heuristic is much more robust as compared to the naive heuristic. As compared to the ILP solutions, the proposed heuristics perform less with increasing number of participants. This is because the potential for system optimization grows as the number of participants increases. Even so, the greedy heuristic can be used to obtain a good feasible ride-sharing plan that provides a much better result as compared to the situation where ride-sharing is not implemented.

8.5. Benefits of using EMS

In this subsection, we report the potential benefit of allowing the utilization of EMS as a backup option for taking the unmatched riders who have left their cars. In principle, solutions obtained by the heuristics in the previous experiments in Section 8.4 could have used EMS. However, we used a prohibitively high EMS cost to avoid any use of it. Table 11 presents the results of the total cost savings and the MileSav. These results are obtained by the greedy heuristic, measured as percentage deviation from the results with $c_e = 1000$. As expected, the benefit of EMS is negatively correlated with its cost. We also see that the benefit of EMS is highest for the CU case and smallest for the SU case. This result is explained in the next paragraph by carefully examining the composition of the mobility cost.

Figure 7 describes the composition of the mobility cost in more detail. The total costs are measured as a percentage of the results with $c_e = 1000$. In the scattered network (i.e., SU and SH1), the fraction of total external mobility cost increases significantly as the unit cost decreases. That is, the effect of an increase in the vehicle-miles provided by EMS is larger than the impact of a decrease in its unit cost. In contrast, the fraction of total external mobility cost remains approximately the same in CU. The results show that the use of EMS indeed reduce the total

Table 11   Benefits of using EMS, measured as a percentage deviation from the results of c_e=1000.

	Mobility cost savings			Mileage savings		
	$c_e = 2$	$c_e = 1.25$	$c_e = 0.5$	$c_e = 2$	$c_e = 1.25$	$c_e = 0.5$
SU	0.59%	2.28%	8.45%	3.16%	5.70%	8.45%
SH1	1.33%	2.96%	9.55%	3.68%	7.20%	9.55%
CU	4.40%	6.77%	12.02%	9.19%	11.28%	12.02%

Note: The results are based on the scenarios of 60 participants in the base case.
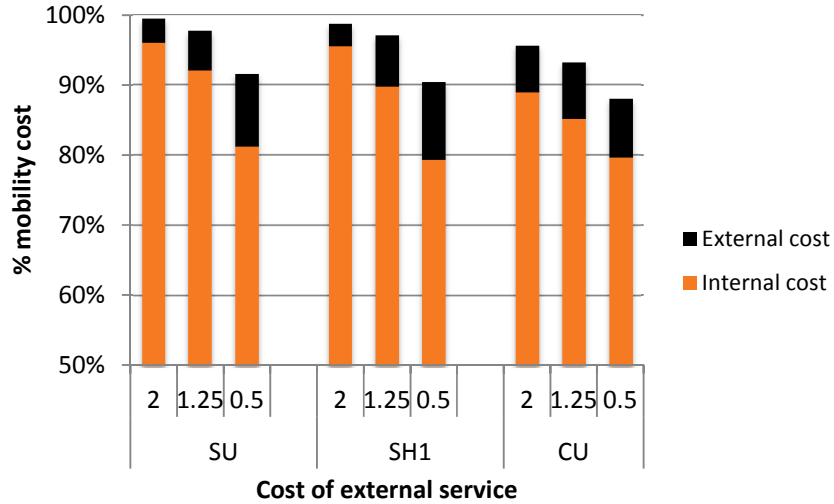
Figure 7   Division of the internal and external mobility costs.

mobility costs. In particular, EMS is most valuable in the clustered network, which points to the fact that more participants in the CU fail to share rides due to a short disconnection along the journey.

9.   Conclusion

In this paper, we consider the ride-sharing problem within a closed community of companies that agree to share the calendars of their employees. In particular, we introduce return restrictions in the ride-sharing problem to meet companies' need for their commuter/business traffic. We provide a general ILP formulation for the ride-sharing problem with meeting points and return restrictions. The model incorporates meeting points, and the option for a rider to transfer between drivers.

Due to the high computational complexity of the problem, we propose a greedy heuristic. The basic idea of this heuristic is to guarantee the return restrictions of the participants who have the most ride-sharing potential, by greedily assigning other participants who can share rides with them taking the spatial, time, and capacity constraints into account. This heuristic can also be used to solve an extended problem that allows the utilization of an external mobility service (EMS) as a backup option.

The proposed ILP model is illustrated using a real-life ride-sharing problem of Significant BV, a Dutch consultancy and research firm. Although the user group is small, and the origins and destinations of the trips are rather scattered across the Netherlands, the results show that Significant can still benefit from ride-sharing by a reduction of vehicle miles of 7-25%.

The performance of the ILP and the heuristic are also tested on a large number of virtual problem instances. The numerical results show that ride-sharing by using the proposed ILP can substantially improve a number of critical performance indicators, i.e., the mileage savings (thus the

mobility cost), the number of cars needed, the matching rate, and the occupancy rate of the cars. On average, the price that has to paid to achieve these performance increases is minor: participants will (i) spend no more than one extra minute on a ride, (ii) arrive at their destination no more than 10 minutes before the corresponding latest arrival time of a trip, and (iii) have to transfer 0.5 times to reach their destination. Overall, our results demonstrate that ride-sharing creates more benefits when the origins and the destinations of the trips are more concentrated. Such a concentration relates to participation density, city density, hotspot density, and city clusters. The results also suggest that it may be desirable to consider the use of EMS. The total mobility cost may further reduce if the companies can make long-term agreements with EMS based on observed trip patterns. Finally, we show that the heuristic is able to provide good feasible solutions quickly, which offers opportunities for further development of improvement heuristics.

Future research can be done along three lines: (i) extending the mathematical model to incorporate the option of using EMS, which allows us to investigate the true benefit of EMS; (ii) developing improvement heuristics to improve the performance in larger-scale problem instance; and (iii) extending the model by introducing an objective related to the resilience of the system, given that ride-sharing systems are much more integrated with the introduction of the return restriction.

Acknowledgments

Appendix A: Shortest path computation and manual adjustment

We assume that every carpool parking lot that results in less than 5 minutes of additional driving distance if it is inserted into the route is located on the shortest path. Similarly, we also assume that every city that leads to less than 10 minutes of additional driving distance if it is inserted into the route is located on the shortest path. We use an example of a ride between Utrecht and Barneveld to illustrate the process of defining the shortest path between two cities in our case. Figure 8 shows the 8 nodes (3 cities and 5 carpool parking lots) located near the route from Utrecht to Barneveld, which satisfy the abovementioned shortest path assumptions. The visiting order of the nodes along the shortest path is determined in a non-decreasing order according to the travel time between the origin and each node.

Figure 8   Nodes of shortest path between Utrecht and Barneveld (Google)

Figure 9   Unrealistic visiting sequence

Figure 10   Connecting nodes on alternative routes

The abovementioned process leads to three undesirable results. First, the travel time between the origin and the destination is longer than that in reality. For instance, the sum of the travel times of the arcs between Utrecht and Barneveld is 63 minutes, while the direct travel time between these two cities is 35 minutes. However, we believe that the consistent input generated by the shortest path generation method will result in representative ride-sharing results. Second, sorting the nodes on travel time between the origin and the intermediate nodes does not always represent a logical visit order, see Figure 9. Third, when alternative routes with comparable travel times exist for a ride between two nodes, the shortest path can consist of nodes along different routes. This leads to arcs that connect nodes along one route with nodes along an alternative route (see Figure 10). Due to the use of distance matrices, the latter two consequences require manual adjustment for the unrealistic shortest paths in our experiments for the purpose of maintaining the original shortest path between any two nodes. In practice, companies usually have access to geographic information system, and thus these issues can be fixed automatically.

References

Agatz, N., A Erera, M Savelsbergh, X Wang. 2011. Dynamic ride-sharing: a simulation study in metro atlanta. *Procedia - Social and Behavioral Sciences* **17**(1) 1450–1464.

Agatz, N, A Erera, M Savelsbergh, X Wang. 2012. Optimization for dynamic ride-sharing: A review. *European Journal of Operational Research* **223**(2) 295–303.

Aissat, K, A Oulamara. 2014. A priori approach of real-time ridesharing problem with intermediate meeting locations. *Journal of Artificial Intelligence and Soft Computing Research* **4**(4) 287–299.

Delhomme, P, A Gheorghiu. 2016. Comparing french carpoolers and non-carpoolers: Which factors contribute the most to carpooling? *Transportation Research Part D: Transport and Environment* **42** 1–15.

Drews, F, D Luxen. 2013. Multi-hop ride sharing. *Sixth Annual Symposium on Combinatorial Search*.

EEA. 2013. Evaluating 15 years of transport and environmental policy integration term 2015: Transport indicators tracking progress towards environmental targets in europe. URL `http://www.eea.europa.eu/publications/term-report-2015`.

Feige, U. 1998. A threshold of ln n for approximating set cover. *Journal of the ACM (JACM)* **45**(4) 634–652.

Furuhata, M, M Dessouky, F Ordóñez, M Brunet, X Wang, S Koenig. 2013. Ridesharing: The state-of-the-art and future directions. *Transportation Research Part B: Methodological* **57** 28–46.

Gruebele, P. 2008. Interactive system for real time dynamic multi-hop carpooling. *Global Transport Knowledge Partnership* .

Herbawi, W, M Weber. 2011. Evolutionary multiobjective route planning in dynamic multi-hop ridesharing. *Evolutionary Computation in Combinatorial Optimization*. Springer, 84–95.

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

IEA. 2005. Saving oil in a hurry: measures for rapid demand restraint in transport. URL `https://www.iea.org/publications/freepublications/publication/savingoil.pdf`.

ITF. 2015. A new paradigm for urban mobility. URL `http://www.itf-oecd.org/sites/default/files/docs/cop-pdf-03.pdf`.

Kuntzky, K, S Wittke, C Herrmann. 2013. Car and ride sharing concept as a product service system simulation as a tool to reduce environmental impacts. *The Philosopher's Stone for Sustainability*. Springer Berlin Heidelberg, 381–386.

Lee, Alan, Martin Savelsbergh. 2015. Dynamic ridesharing: Is there a role for dedicated drivers? *Transportation Research Part B: Methodological* **81** 483–497.

Savelsbergh, M, T Van Woensel. 2016. 50th anniversary invited articlecity logistics: Challenges and opportunities. *Transportation Science* **50**(2) 579–590.

Stiglic, M, N Agatz, M Savelsbergh, M Gradisar. 2015. The benefits of meeting points in ride-sharing systems. *Transportation Research Part B: Methodological* **82** 36–53.

Stradling, S. 2007. *Determinants of Car Dependence*, chap. 10. 187–204.

van Dender, K, M Clever. 2013. Recent trends in car usage in advanced economies–slower growth ahead?: Summary and conclusions. *International Transport Forum Discussion Papers*. 2013/9, OECD Publishing.

Winter, S, S Nittel. 2006. Ad hoc shared-ride trip planning by mobile geosensor networks. *International Journal of Geographical Information Science* **20**(8) 899–916.