

Online Ridesharing with Meeting Points

Jiachuan Wang
The Hong Kong University of Science
and Technology
Hong Kong, China
jwangey@cse.ust.hk

Peng Cheng
East China Normal University
Shanghai, China
pcheng@sei.ecnu.edu.cn

Libin Zheng
Sun Yat-sen University
Guangzhou, China
zhenglb6@mail.sysu.edu.cn

Lei Chen
The Hong Kong University of Science
and Technology
Hong Kong, China
leichen@cse.ust.hk

Wenjie Zhang
The University of New South Wales
Australia
wenjie.zhang@unsw.edu.au

ABSTRACT

Nowadays, ridesharing becomes a popular commuting mode. Dynamically arriving riders post their origins and destinations, then the platform assigns drivers to serve them. In ridesharing, different groups of riders can be served by one driver if their trips can share common routes. Recently, many ridesharing companies (e.g., Didi and Uber) further propose a new mode, namely “ridesharing with meeting points”. Specifically, with a short walking distance but less payment, riders can be picked up and dropped off *around* their origins and destinations, respectively. In addition, *meeting points* enables more flexible routing for drivers, which can potentially improve the global profit of the system. In this paper, we first formally define the Meeting-Point-based Online Ridesharing Problem (MORP). We prove that MORP is NP-hard and there is no polynomial-time deterministic algorithm with a constant competitive ratio for it. We notice that a structure of vertex set, k -skip cover, fits well to the MORP. k -skip cover tends to find the vertices (meeting points) that are convenient for riders and drivers to come and go. With meeting points, MORP tends to serve more riders with these convenient vertices. Based on the idea, we introduce a convenience-based meeting point candidates selection algorithm. We further propose a hierarchical meeting-point oriented graph (HMPO graph), which ranks vertices for assignment effectiveness and constructs k -skip cover to accelerate the whole assignment process. Finally, we utilize the merits of k -skip cover points for ridesharing and propose a novel algorithm, namely SMDB, to solve MORP. Extensive experiments on real and synthetic datasets validate the effectiveness and efficiency of our algorithms.

PVLDB Reference Format:

Jiachuan Wang, Peng Cheng, Libin Zheng, Lei Chen, and Wenjie Zhang. Online Ridesharing with Meeting Points. PVLDB, 15(13): 3963–3975, 2022. doi:10.14778/3565838.3565849

PVLDB Artifact Availability:

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 15, No. 13 ISSN 2150-8097.
doi:10.14778/3565838.3565849




Figure 1: An Example of Meeting Points

The source code, data, and/or other artifacts have been made available at <https://github.com/dominatorX/open>.

1 INTRODUCTION

Nowadays, on-demand ridesharing becomes important in civil commuting services. Together with online platforms (e.g., DiDi [4]), ridesharing surpasses traditional taxi services with more saved energy, less air pollution, and lower cost [32].

In *online ridesharing*, riders arrive dynamically. Platforms need to deal with them immediately for different objectives, including maximizing the number of served riders [12, 14, 25, 30, 40], minimizing the total travel distance [8, 20, 21, 23, 26, 28–30, 34], or maximizing the unified revenue [9, 10, 36].

Ridesharing allows one driver to serve more than one group of riders simultaneously. The route of a driver is a sequence of pick-up/drop-off points. Given a set of drivers and riders, *route planning* is to design and update routes every time a rider arrives. A key operation, called *insertion*, shows great effectiveness and efficiency for solving online ridesharing problem [11, 12, 23, 26, 28, 30, 34, 36, 40]. It tries to insert a newly coming rider’s origin and destination into a driver’s route without changing the order of his/her current sequence of pick-up/drop-off points.

However, due to the complex topology of the city road network, some locations are spatially close to each other but hard to access for vehicles. Especially, if two locations are only connected by a Pedestrian Street, where vehicles cannot go through, a short walking could greatly reduce the travel cost of the assigned vehicle. To deal with the case, *meeting points* (MP for short) are introduced as alternative locations for pick-up/drop-off locations of riders [31]. As shown in Figure 1, a rider r at location A wants to go to location C . The nearby roads are directed roads. A driver w at location B is assigned to serve r . Then, the shortest route for w is represented in

blue dashed arrow lines. If r can move a short distance, for example, to location \mathbf{D} (i.e., an MP), w can serve r through a much shorter route displayed in the green line.

In a recent work [41], the authors utilize meeting points to improve the results of offline ridesharing problems, whose method however is slow and only can handle up to 40 riders/vehicles in real time. Thus, it is not practical for online applications (e.g., Uber and DiDi) with hundreds of riders/vehicles every several seconds. Existing studies also investigate the strategies to properly select MPs with online surveys [13, 15]. In industry, Uber recently offers Express POOL to encourage riders to walk to Express spots (meeting points) for efficient routing [1]. But Uber Express Pool only schedules the route for each vehicle when there are shareable ride-requests to group with. In addition, the MPs in Uber Express Pool are similar to the stops of buses for nearby riders to come together and thus not flexible [2]. For the example in Figure 1, Uber Express POOL will not assign driver w to pick up rider r until another rider r' appears close to point D (i.e., the selected pick-up stop). In summary, to the best of our knowledge, there is no existing solution for *online* ride-sharing services boosted with flexible MPs.

With MPs, online ridesharing is more flexible but challenging. To solve it, we first define *Meeting-Point-based Online Ridesharing Problem* (MORP) mathematically. Based on existing studies [9, 36], we prove that the MORP problem is NP-hard and has no deterministic algorithms with a constant competitive ratio, thus intractable.

In the traffic network, some vertices are more convenient to come and go and thus “popular” during assignments, such as those close to highways. Flexible MPs makes it possible to serve more riders at or near those vertices, which makes them even more frequently used. This motivates us to take the advantage of k -skip cover V^* [33], which is a subset of vertices to be the skeleton of a graph G .

To solve MORP, we prepare *Meeting point candidates* for each vertex offline. On the other hand, we propose a hierarchical meeting-point oriented (HMPO) graph, which further filters MPs for effectiveness and accelerate shortest path queries during insertion. Based on the k -skip cover in HMPO graph, we propose a meeting-point-based insertion operator, named SMDB, which can solve MORP effectively and efficiently. We make the following contributions:

- We formulate the online route planning problem with MPs mathematically, namely MORP. We prove that it is NP-hard and has no algorithm with constant competitive ratio in Section 3.
- With observations and analyses, we propose a heuristic algorithm to select MP candidates for riders in Section 5, which is based on a unified cost function considering the travel cost from additional walking. We propose a novel hierarchical structure of the road network, namely hierarchical meeting-point oriented (HMPO) graph, to fasten the solution for MORP in Section 6.
- With HMPO graph, we propose an effective and efficient inserter, namely SMDB, to handle the requests in MORP in Section 7.
- Extensive experiments on synthetic and real data sets show the efficiency and effectiveness of SMDB in Section 8.

2 BACKGROUND AND RELATED WORKS

2.1 Online Ridesharing

Route planning for ridesharing, which has been widely studied in recent years, is a variant of the dial-a-ride problem (DARP) proposed

in 1975 [38, 39]. Traditional DARP problems usually have additional restrictions, such as limiting the drivers to start from/return to depot(s) and serve all the requests [19, 22]. These settings lead to small scale datasets with near-optimal solutions. In comparison, route planning for ridesharing is more applicable in the real world, which applies to hundreds of thousands of requests and tens of thousands of drivers with locations distributed over large scale road network [9, 10, 23, 36]. Realistic revenue and serving cost can be designed as objectives to meet the requirement of ridesharing platforms [9, 10, 36, 42]. A common setting for the serving cost is a unified score based on distance/time cost of driving and penalty of rejecting riders. One can further extend the unified cost to an application-specific one, such as maximizing the score combined with complicated social utilities from both workers and requests [11, 16]. Using meeting point results in additional costs such as walking, which is handled with a unified cost function.

2.2 Insertion

Real-world ridesharing services require solutions for online instead of off-line mode. Efficient heuristic methods are developed for route planning without information of future workers and requests in advance [9, 10, 12, 23, 26, 28, 34, 40]. With large scale dataset and requirement for real-time response, a commonly used operator called insertion shows good performance for route planning [11, 12, 23, 24, 26–30, 34, 36]. Insertion greatly reduces the search space of possible new routes to serve each rider from $O(N!)$ to $O(N^2)$. Tong *et al.* further reduce its time complexity to linear time using dynamic programming [35, 36]. We adapt the linear insertion [36] for our MORP problem as baseline and further propose a more effective inserter based on a new graph structure.

2.3 Ridesharing with meeting points

As an effective way to improve ridesharing experience, meeting points (MPs) are used in online hailing companies, such as Didi and Uber. Stiglic *et al.* [31] first introduce the concept of “meeting points” to give alternatives to pick up and drop off riders. They devise a heuristic algorithm for meeting-point-based offline ridesharing problem. Zhao *et al.* [41] develop the mathematical model for the offline ridesharing problem with flexible pickup and delivery locations and propose an integer linear programming model to solve it. In recent years, Uber had proposed Express Pool as an online ridesharing service, in which riders need to walk a little and may have a longer waiting time, but get a discount. On the other hand, Uber prefers to group passengers together with the same MPs, then pick up and drop off them like a bus with selectable stations, which has less flexibility [2]. In this paper, we focus on the online ridesharing problem with MPs, which needs to respond to requests within a very short time (e.g., within 5 seconds).

2.4 k -skip cover

Tao *et al.* [33] first propose k -skip cover: given a graph $G(V, E)$, we call a set $V^* \subseteq V$ a k -skip cover if for any shortest path SP on G with exactly k vertices, there is at least one vertex $u \in SP$ satisfying $u \in V^*$. In general, for any shortest path SP in G , vertices of $SP \cap V^*$ succinctly describes SP by sampling the vertices in SP with a rate of at least $\frac{1}{k}$. Such a sub-path out of the whole

path is called a k -skip shortest path. In many applications, such as electronic map presentation, given all vertices are unnecessary and the k -skip shortest path gives a good skeleton of it. The study further shows that answering k -skip queries is significantly faster than finding the original shortest paths. Funke et al. [17] further generalize the work of [33] by constructing k -skip path cover for all paths instead of only shortest paths. Besides, they devise a new way of constructing a smaller size of k -skip path cover. To construct k -skip cover, k -skip neighbor is defined [33]. If vertices u and v are both in a k -skip cover and u is a k -skip neighbor of v , the shortest path from u to v does not pass any other vertices in the k -skip cover.

3 PROBLEM DEFINITION

3.1 Basic Notations

We use graph $G_c = \langle V_c, E_c \rangle$ to represent a road network for cars, where V_c and E_c indicate a set of vertices and a set of edges, respectively. Each edge, $(u, v) \in E_c (u, v \in V_c)$, is associated with a weight $t_c(u, v)$ indicating travel time for driving from vertex u to v through it. Similarly, graph $G_p = \langle V_p, E_p \rangle$ is used to represent a road network for passengers. Each edge $(u, v) \in E_p (u, v \in V_p)$ is weighed by $t_p(u, v)$ as its travel time for walking. For the two graphs, we denote the union of vertices as $V = V_c \cup V_p$. In the city network, passengers are more flexible. We set all the edges in V_p undirected according to the network of OSM [6]. In addition, usually for any edge (u, v) , walking is slower than driving (e.g., $t_c(u, v) < t_p(u, v)$). We denote *path* as a sequence of vertices $\{v_1, v_2, \dots, v_k\}$ with travel time $\sum_{i=1}^{k-1} t(v_i, v_{i+1})$. For each pair of vertices (u, v) , we represent the time cost of its shortest path for cars and passengers as $SP_c(u, v)$ and $SP_p(u, v)$, respectively.

Definition 3.1 (Drivers). Let $W = \{w_1, w_2, \dots, w_n\}$ be a set of n drivers that can provide transportation services. Each driver w_i is defined as a tuple $w_i = \langle l_i, a_i \rangle$ with a current location l_i and a capacity limitation a_i .

At any time, the number of riders in a taxi of driver w_i must not exceed its capacity a_i .

Definition 3.2 (Requests). Let $R = \{r_1, r_2, \dots, r_m\}$ be a set of m requests. Each request $r_j = \langle s_j, e_j, tr_j, tp_j, td_j, p_j, a_j, pi_j, de_j, wp_j, wd_j \rangle$ is denoted with its source location s_j , destination location e_j , release time tr_j , latest pick-up time tp_j , deadline td_j , rejection penalty p_j , and a capacity a_j . Once it is assigned, two vertices as pick-up point pi_j and drop-off point de_j will be recorded. The shortest time for a request to walk from source to pick-up point is represented as $wp_j = SP_p(s_j, pi_j)$ and from drop-off point to destination is denoted as $wd_j = SP_p(de_j, e_j)$.

In practice, we do not ask riders to set all the parameters in Definition 3.2. Excluding s_j, e_j , and a_j , which are given by the rider, other parameters can be auto-filled by the platform to improve the user's experience, such as deadline td_j for reasonable serving time [23]. A request r_j can be served by driver w_i only if: (a) w_i can arrive at pi_j after tr_j ; (b) the remaining capacity of w_i is at least a_j when he/she arrives at pi_j ; and (c) w_i can pick r_j at pi_j no later than tp_j and deliver r_j at de_j no later than $td_j - wd_j$.

Note that in real-application, rejections are unavoidable for the "urgent" requests on a platform, especially at rush hours. The loss from rejecting r_j is denoted by penalty p_j . The penalty can be

application-specific. Furthermore, we denote all the requests that are served by driver w_i as R_{w_i} . Then, $\hat{R} = \cup_{w_i \in W} R_{w_i}$ and $\bar{R} = R \setminus \hat{R}$ refer to the total served and unserved requests, respectively. To simplify, we will use r_j to indicate a request or a rider of a request without differentiation.

Definition 3.3 (Meeting Points). For a request $r_j \in R_{w_i}$, the pick-up point $pi_j = u \in V_c \cap V_p$ denotes that driver w_i will pick up rider r_j at vertex u . The drop-off point $de_j = v \in V_c \cap V_p$ denotes that w_i will drop off r_j at vertex v . Pick-up and drop-off points are *meeting points* (MP for short).

With MPs, we allow the drivers to flexibly pick up and drop off passengers. Traditional online ridesharing solutions only assign drivers to pick up a rider r_j from its source s_j and drop off r_j to its destination e_j . With MPs, the rider r_j can move a short distance to location pi_j and be picked up there by a driver. After being dropped off at location de_j , r_j walks to his/her destination e_j .

Definition 3.4 (Route). The route of a driver w_i located at l_i is a sequence, $S_{w_i} = [l_i, l_{x_1}, l_{x_2}, \dots, l_{x_k}]$, where each l_{x_n} ($n \in [1, k]$) is a pick-up or drop-off point of a request $r_j \in R_{w_i}$ and the driver will reach these locations in the order from 1 to k .

We call the vertices of a route as *stations*. Drivers move on shortest paths between stations. A feasible route satisfies: (a) $\forall r_j \in R_{w_i}$, its drop-off time of de_j is no later than $td_j - wd_j$; (b) $\forall r_j \in R_{w_i}$, its pick-up point pi_j appears earlier than its drop-off point de_j in S_{w_i} ; (c) The total capacity of undropped riders is no larger than the driver's capacity a_i at any time. We denote $S = \{S_{w_i} | w_i \in W\}$ as all the route plans.

Here we define $D(S_{w_i})$ as the shortest time to finish S_{w_i} :

$$D(S_{w_i}) = SP_c(l_i, l_{x_1}) + \sum_{k=1}^{|S_{w_i}|-2} SP_c(l_{x_k}, l_{x_{k+1}})$$

3.2 Meeting-Point-based Online Ridesharing

Definition 3.5 (Meeting-Point-based Online Ridesharing Problem, MORP). Given transportation networks G_c for cars and G_p for passengers, a set of drivers W , a set of dynamically arriving requests R , a driving distance cost coefficient α , a walking distance cost coefficient β , MORP problem is to find a set of routes $S = \{S_{w_i} | w_i \in W\}$ for all the drivers with the minimal unified cost:

$$UC(W, R) = \alpha \sum_{w_i \in W} D(S_{w_i}) + \beta \sum_{r_j \in R} (wp_j + wd_j) + \sum_{r_j \in \bar{R}} p_j \quad (1)$$

which satisfies the following constraints: (i) Feasibility constraint: each driver is assigned with a feasible route; (ii) Non-undo constraint: if a request is assigned in a route, it cannot be canceled or assigned to another route; if it is rejected, it cannot be revoked.

3.3 Hardness Analysis

LEMMA 3.6. *The MORP problem is NP-hard.*

PROOF. For details, please refer to Appendix A in [37]. \square

The Competitive Ratio (CR) is commonly used to analyze the online problem. CR is defined as the ratio between the result achieved by a given algorithm and the optimal result for the corresponding offline scenario. It is proved no constant CR to maximize the total revenue for basic route planning for shareable mobility problems with neither deterministic nor randomized algorithm [9, 36].

LEMMA 3.7. *There is no randomized or deterministic algorithm guaranteeing constant CP for the MORP problem.*

PROOF. For details, please refer to Appendix B in [37]. \square

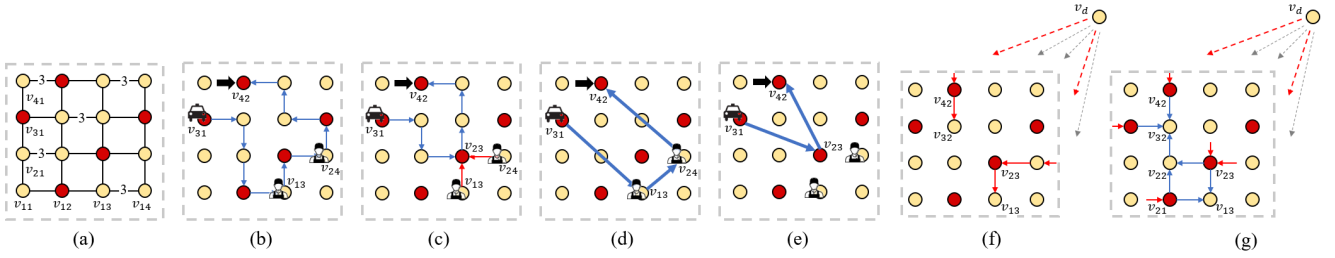


Figure 2: An Example for the fitness of meeting points and k -skip cover. (a) undirected graph G has edges with weight of 1 except for those marked as 3. Red vertices form a 2-skip cover, which are also good MP candidates. (b) Two requests at v_{13}, v_{24} are heading to v_{42} . Without meeting points, driver at v_{31} need to serve them with route $v_{31} \rightarrow v_{13} \rightarrow v_{24} \rightarrow v_{42}$. (c) Serving them with MP v_{23} is much more effective, traversing along the popular vertices $v_{31} \rightarrow v_{23} \rightarrow v_{42}$. (d) Traditional ridesharing computes 3 shortest paths $SP(v_{31}, v_{13})$, $SP(v_{13}, v_{24})$, and $SP(v_{24}, v_{42})$. (e) With MPs, it only computes 2 shortest paths $SP(v_{31}, v_{23})$ and $SP(v_{23}, v_{42})$ beginning and ending with vertices in V^* , which can be computed efficiently with V^* . (f) v_d is far away from the subgraph, and we want to compare 2 shortest path queries from v_d to v_{32} and to v_{13} . Before reaching them, the two paths must reach their surrounding MPs first (e.g., v_{42} and v_{23}). (g) k -skip cover “cut off” shortest paths, so that any shortest paths towards v_{32} and v_{13} must reach v_{32}, v_{42}, v_{21} , and v_{23} first. The distance relationships inside the “cut” can help us to bound the distance difference between expensive queries from v_d .

4 OVERVIEW OF THE FRAMEWORK

In this section, we first introduce the k -skip cover and how it coincides with the demand of MORP [33]. Then we show the detail of our framework, which makes full use of the k -skip cover.

A k -skip cover V^* is vertex set on graph G . By definition, any shortest path of length k has at least 1 vertex that is $\in V^*$ [33]. A good V^* has small size, such that each of its vertices is frequently passed for transportation. We claim that k -skip cover suits MORP problem well for 2 reasons. Figure 2 is shown as an example:

- For a road network, vertices that are convenient to come and go are good candidates for both a k -skip cover V^* and MPs, as they are usually components of many short paths and thus vital for transportation. These convenient vertices are fast for computing and rider-concentrated, thus “popular” during assignments. MPs enable drivers to serve more requests through them, which make them more popular. After constructing the cover V^* , shortest path queries on points of V^* can be computed quickly. This motivates us to build a hierarchical meeting-point oriented graph with k -skip cover in Section 6, which encourages more requests to be served effectively through V^* and boost the overall query.
- During assignment, we try to insert nearby MP candidates into each worker, where many queries are from same source to different MPs. We claim that k -skip cover has underexplored merits to bound the differences between these queries, shown in Section 7. If one of them is infeasible to insert, the bound makes it possible to prune other MPs, which greatly improve the efficiency. We explore this attribute and devise a new insertion algorithm SMDB in Section 7.2.

Taking the advantage of k -skip cover, we construct our framework to solve MORP problem, shown in Figure 3.

During the online assignment, requests arrive and are assigned one-by-one. Given a new request r_j , we first select meeting points (MP) according to its source and destination locations. Then we iteratively insert each pair of MPs into each driver w_i . In our work, we adapt the insertion algorithm with time complexity $O(n)$ for MP insertion [36]. If there exists valid insertion(s), we choose the one with the minimal unified cost; otherwise, we reject r_j .

Compared with traditional ridesharing problem, we need to select MPs and accelerate related computations. Our work conducts off-line pre-processing to improve the efficiency and effectiveness

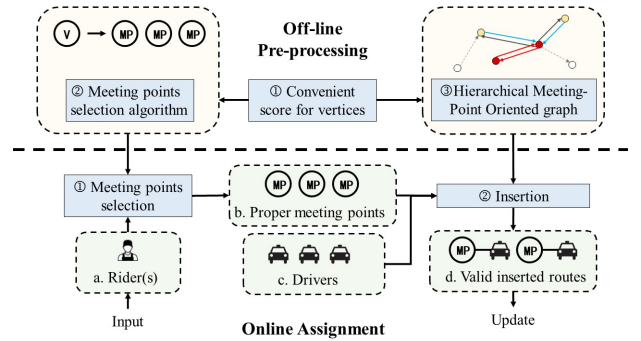


Figure 3: Assignment framework overview.

of online assignment. we first propose a method to evaluate the convenience of each vertex, which bases on statistics of shortest path queries to meet the demand of MPs and k -skip cover. MP candidates are selected for each vertex in Section 5, which greatly shrink search space with $O(1)$ time complexity during online assignment. In addition, we design a structure, namely hierarchical meeting-point oriented graph (HMPO graph), to rank vertices for effective assignments in Section 6. k -skip cover is embedded for efficiency. Based on it, we further devise a new insertion algorithm in Section 7, namely SMDB, which prunes candidate MPs and drivers during the insertion phase.

5 SELECT MEETING POINT CANDIDATES

After Stiglic et al. [31] introduced the concept of “meeting points” (MP) to provide flexible pick-up and drop-off points for riders, many researchers aim to find an effective solution for ridesharing with MPs [13, 31]. In this paper, we pre-select a set of vertices as candidates to serve their nearby vertices.

In this section, we first introduce the motivation of selecting MP candidates. Then, we propose a heuristic algorithm, Local-Flexibility-Filter, to select them.

5.1 Meeting Point Candidates

To insert one rider into a route, traditional ridesharing only inserts 1 pair of pick-up and drop-off points. Assume that on average, one vertex has K nearby vertices, which are within the acceptable distance for rider to walk to. Enumerating K pick-up points and K drop-off points as MPs increases the time cost by a factor of K^2 ,

which is unacceptable. Here, we pre-select *Meeting Point Candidates* for each vertex. To insert a pair of origin and destination, we can directly get their MPs, instead of searching all their neighbors.

Definition 5.1. (Meeting Point Candidates) Given the road network for cars $G_c = \langle V_c, E_c \rangle$ and passengers $G_p = \langle V_p, E_p \rangle$, MP Candidates MC is a dictionary, which maps each vertex $u \in V_p$ to a vertex set $MC(u) = \{v_1, v_2, \dots\} \subseteq V_c \cap V_p$. For the MORP problem, we only select MPs for a vertex from its MP candidates, that is, $\forall r_j \in R, pi_j \in MC(s_j)$ and $de_j \in MC(e_j)$.

5.2 Meeting Point Candidate Selection

MP candidates should easily get to and conveniently reach other vertices. We introduce our Local-Flexibility-Filter algorithm to find the candidate sets $MC(\cdot)$ in two phases.

Vertices convenient for drivers. The first phase aims to find the vertices which are convenient for drivers, thus boosting transportation efficiency. As the example shown in Figure 2, MPs and k -skip cover have similar preferences. We quantify the convenience from the statistic of shortest path queries, named *equivalent in/out cost* ECI/ECO for each vertex. If a vertex u is inserted into a route between v_1 and v_2 , the $ECI(u)$ and $ECO(u)$ indicates the average cost from v_1 to u and from u to v_2 , respectively.

As riders are usually assigned to nearby drivers, shortest path queries between a vertex and its surrounding vertices will indicate its convenience. For each source vertex u , we directly select its n_r nearest vertices on the car graph G_c as *reference vertices* $n_o(u)$. Intuitively, we define the equivalent out cost of u as the average distance towards its reference vertices:

$$ECO(u) = \frac{\sum_{v \in n_o(u)} SP_c(u, v)}{n_r}$$

Similarly, for $ECI(u)$, we reverse the graph and select its n_r nearest inward neighbors $n_i(u)$. $ECI(u) = \frac{\sum_{v \in n_i(u)} SP_c(v, u)}{n_r}$ indicates the average cost of reaching u from other vertices. Here, n_r depends on the density of a road network and the speed of drivers.

Vertices convenient for riders. The second phase takes the walking convenience into account. For each vertex u , we select vertices $\{v_1, v_2, \dots\}$ no farther than a maximum walking distance d_m . For each reachable vertex v_i , we calculate a serving-cost score $SCS(u, v_i)$ combining both walking distance and the equivalent in/out costs as $SCS(u, v_i) = \beta \cdot SP_p(u, v_i) + \alpha (ECI(v_i) + ECO(v_i))$, where α and β are the weight factors for driving and walking cost defined in MORP. Especially, each vertex u has the SCS score for itself: $SCS(u, u) = \beta \cdot SP_p(u, u) + \alpha (ECI(u) + ECO(u)) = \alpha (ECI(u) + ECO(u))$.

Filtering of MP candidates. We want a small candidate set for pruning effectiveness, while only good MPs are retained. A vertex with a low average cost to serve a rider (low SCS) does not need alternatives. But a vertex with high SCS needs more choices to find a good MP.

Here, for a vertex u , we prune all its candidates with a score higher than $SCS(u, u) + thr_{CS}$. thr_{CS} is a user-specified threshold. We also set an upper bound nc_m for the number of candidates of each vertex.




Figure 4: Original Graph for car and passenger

Now, we have the MP candidate set for each vertex. Once a request arrives, we find the candidate MPs only in the MP candidate sets of its source and destination.

We show an $O(|V|)$ selection algorithm with detailed analysis in Appendix C in [37].

6 HIERARCHICAL MEETING-POINT ORIENTED GRAPH

With MPs, assigned routes can be concentrated on the convenient vertices, where inconvenient ones can be replaced by nearby MPs. In this section, we first validate this assumption by analyzing the real-world data. Then we rank the vertices and design a Hierarchical Meeting-Point Oriented graph (*HMPO graph*). To be more specific, we 1) find defective vertices and guide us to assign riders through convenient vertices for effectiveness; and 2) further define core vertices, which form a k -skip cover and reduce additional computing costs.

6.1 Graph Analysis for Meeting Point

Take the road network of New York City on Open Street Map (OSM) [6] as an example, which contains 58189 vertices and 122337 edges. In this real-world road network, some vertices are more “convenient” than the others, which have larger traffic flows and lower transition costs. Based on this intuition, we evaluate the original road graph and vertex-flexibility with methods in Section 5. Based on the result in Section 5, vertices can be ranked by $ECI(\cdot) + ECO(\cdot)$ as an indication of convenience. We find that for more than 70% vertices, each of their MP candidates has at least one vertex among the 20% most convenience vertices. According to this observation, convenient vertices would be more frequently used for ridesharing with MPs compared without MPs. On the other hand, some vertices are inconvenient to drive in and out for drivers, while MPs can be their convenient alternatives. This motivates us to build a hierarchical graph, which gives an indication for effective assignment and boosts the queries on convenient vertices.

Formally, we introduce Hierarchical Meeting-Point Oriented Graph (HMPO Graph), which gives the hierarchical order over the vertex set V and has 3 levels of vertices:

- **Core vertices** V_{co} . They are used as MPs frequently.
- **Defective vertices** V_{de} . They are inconvenient to access.
- **Sub-level vertices** V_{su} . The remaining vertices are classified as sub-level vertices.

The three sets of vertices form a partition of all vertices in G_c and G_p , that is, $V_{co} \cup V_{su} \cup V_{de} = V$, $V_{co} \cap V_{su} = \emptyset$, $V_{co} \cap V_{de} = \emptyset$, and $V_{su} \cap V_{de} = \emptyset$. We introduce an example below, which will be used to show the main steps of our algorithms to build a HMPO graph in this section.

Table 1: ECO and ECI with 3 nearby vertices, and $MC(\cdot)$

ID	A	B	C	D	E	F
$ECO(\cdot)$	5.33	4.67	5.33	5	6.67	∞
$ECI(\cdot)$	5.33	3.67	5.67	∞	6.33	6.33
$MC(\cdot)$	{A, B}	{B}	{B, C}	{A, D, E}	{E}	{C, E, F}

Example 6.1. There are 6 vertices A to F in Figure 4. The graph for car is directed in Figure 4(a). The graph for passenger is undirected shown in Figure 4(b). Considering $n_r = 3$ nearby vertices, we derive $ECO(\cdot)$ and $ECI(\cdot)$ according to Section 5 in Table 1. With $\alpha = \beta = 1$, $nc_m = 3$, $d_m = 30$, and $thr_{CS} = 15$, we derive the MP candidates $MC(\cdot)$ shown in Table 1.

6.2 Defective Vertices

Defective vertices are inconvenient vertices for vehicles to access and will be eliminated. These defective vertices are not welcomed in traditional ridesharing either, but only with meeting points, removing them is feasible as we can serve riders with MPs. However, vertex removal also induces additional travel cost. We use a heuristic method to select and remove these vertices and improve the assignment effectiveness.

Vertex removing cost. Removing a vertex u from a traditional graph leads to 2 kinds of additional costs for transportation: (i) *the detour cost*. If a query finds a path containing u , removing u means that we need to find a new path without passing u . The new path is usually longer than the original one and results in a detour; (ii) *the inaccessibility cost*, which is from queries with u as origins or destinations. No path will exist after removing u .

With prepared MP candidates for each vertex in G_p , the walking cost is checked directly without query on G_p . Thus, for assignment, shortest path queries are only on the graph G_c , which is affected by vertex removing.

With MPs, the inaccessibility cost of removing vertices can decrease from ∞ to a limited mixture of walking and driving costs. We propose a heuristic algorithm called Defective Vertices Selection Algorithm (*DVS algorithm* for short) to remove the defective vertices without error from the detour cost and maintain its accessibility in the meantime. As the vertices with larger equivalent in/out cost ECI and ECO are harder to reach and leave, we try to remove defective vertices in decreasing order of $ECI + ECO$. For each vertex, we first check its MP candidates to ensure that at least one candidate is not defective and the vertex is accessible. Then, we search the shortest path queries passing it with a bounded search space. If there is no detour cost, we mark it as defective vertex.

We show the detail of DVS algorithm with time complexity $O(N \log N)$ in Appendix D and an example in Appendix E in [37]. In addition, we propose two lemmas about the removing cost of defective vertices.

LEMMA 6.2. *Removing all vertices selected by the DVS algorithm from G_c with their edges leads to no detour cost.*

PROOF. For details, please refer to Appendix F in [37]. \square

LEMMA 6.3. $\forall u \in V$ is accessible after removing vertices selected by the DVS algorithm from G_c with MPs.

PROOF. For details, please refer to Appendix G in [37]. \square

6.3 Core Vertices

Some convenient vertices take over the major traffic flow of a city network, such as the vertices along the highways. People in nearby places usually drive to the highway and go along it, then finally turn to small road for the destination. By assigning more route segments along highways, drivers can finish requests fast along highways in most time. With MPs, we can efficiently organize more routes along highways. Thus, these fast and request-concentrated vertices are more frequently used in solving the MORP problem. Finding such a group of convenient vertices and optimizing their related operations are pretty beneficial.

In this subsection, we select *core vertices* V_{co} as a backbone of the whole graph. As shown in Figure 2, k -skip cover [33] is a skeleton vertex set, which coincides with our demand on finding convenient vertices and accelerating their related computations. Thus, we define the Core Vertices Selection Problem to find core vertices, which have a good coverage of MP with embedded k -skip cover. We first formulate the coverage of MP as an integer linear program, which can be solved with constant approximation ratio. Then, using an evaluational cost from the first step, we further complement its output and make it a k -skip cover as well.

First, we define the Core Vertices Selection Problem:

Definition 6.4. Given the union of vertices $V = V_c \cup V_p$ and car graph G_c with its selected defective vertices V_{de} , the *Core Vertices Selection Problem* is to find a set of core vertices $V_{co} \subseteq V - V_{de}$ with the minimum size $|V_{co}|$, such that

- (i) V_{co} is a k -skip cover of the updated graph with V_{de} and adjacent edges removed from G_c .
- (ii) Proportion factor ϵ of vertices in passenger vertex set V_p has at least one vertex $u \in V_{co}$ as its MP candidate.

For attribute (i), the concept of the **k -skip cover** is introduced in Section 2. Both of [17, 33] tend to keep the vertices that are components of more shortest paths and remove the ones that belong to fewer shortest paths. They simply rank their vertices according to the degree, which is pretty coarse as road network has very low vertex degree. Instead, we use a cost which is evaluated from the internal connections between vertices for ranking, aiming at comparing vertex importances more quantitatively. Attributes (ii) assures that most of the requests can be assigned with core vertices as MPs. So the optimization for core vertices can benefit more assignments. If a set of core vertices is valid to represent the graph skeleton, the smaller its size is, the more efficient its inner operations are after further preprocessing. So, we want to find a V_{co} which satisfies attribute (i) (ii) with minimal size.

Now, to satisfy the attribute (ii), we can formulate an equivalent interger linear program. The detailed mapping steps include: (i) constructing candidate serving set MS to indicate the set of servable vertices for each vertex $u \in V - V_{de}$, that is, for any pair of vertex $u, v \in V$, $u \in MC(v)$ iff $v \in MS(u)$. As an inverse-function relationship of MC , we can convert the MP candidate set by scanning each vertex $u \in V$ and for each $v \in MC(u)$, add u to $MS(v)$. This process is linear; (ii) the universe \mathcal{U} is defined to be the set of all the passenger-accessible-vertices v_p ; (iii) find the smallest sub-collection of vertices $\{u_1, u_2, \dots\} \subseteq V - V_{de}$ that the union of their serving sets covers at least ϵ of the universe, that is, $|\cup_k MS(u_k)| \geq \epsilon \cdot |V_p|$.

Then, we have the following integer linear program (ILP):

$$\begin{aligned}
\min \quad & \sum_{u \in V - V_{de}} \delta_u \\
\text{s.t.} \quad & \phi_v + \sum_{u: v \in MS(u)} \delta_u \geq 1 \quad \forall v \in V_p \\
& \sum_{v \in V_p} \phi_v \leq (1 - \epsilon) |V_p| \\
& \delta_u \in \{0, 1\} \quad \forall u \in V - V_{de}, \\
& \phi_v \in \{0, 1\} \quad \forall v \in V_p,
\end{aligned} \tag{2}$$

where $\delta_u = 1$ iff vertex u is chosen to serve its candidate serving set $MS(u)$. The binary variable $\phi_v = 1$ iff vertex v is not served by any selected vertex.

The partial set cover problem is an NPC problem. Luckily, in our setting, each element (vertex) is in at most nc_m candidate serving sets as the times of adding a vertex u to some $MS(\cdot)$ equals to the number of u 's MP candidates $|MC(u)| \leq nc_m$. Such a case is called a low-frequency system. There exists a solution to approximate the optimum within a factor nc_m using LP relaxation in polynomial time [18]. In detail, the corresponding LP relaxation can be derived by substituting the constraints $\delta_u \in \{0, 1\}$ by $\delta_u \geq 0, \forall u \in V - V_{de}$ and $\phi_v \in \{0, 1\}$ by $\phi_v \geq 0, \forall v \in V_p$. Then the problem is transferred to a linear Program LP . After deriving the dual LP of it, [18] iteratively chooses each set to be the highest cost set and obtains a feasible cover with the primal-dual stage. This step generates a cost $cost_u$ for each set. Intuitively, in our setting, this cost $cost_u$ can be treated as how bad the “quantity” each vertex u with its serving set $MS(u)$ to be included in the final output. The higher its value, the worse a vertex is. Finally, it chooses the solution with minimum cost.

However, to guarantee attribute (i) simultaneously, we cannot apply this algorithm directly to our problem. Both of the two attributes require the set of chosen vertices to have a good “coverage” of other vertices, that is, each of them can serve many nearby vertices. Motivated by this observation, we first use our equivalent in/out cost ECI/ECO as a coarse estimation in the solution of [18]. Then we use the $cost_u$ of [18] as the ranking for k -skip cover completion. For the detail of our CVS algorithm and complexity analysis, please refer to Appendix H in [37].

We show that the size of V_{de} is bounded with the following lemma and present an example below.

LEMMA 6.5. *Assume that we have N vertices in total, with M set as optimal solution for the attribute (ii), the upper bound of the size of core vertex set is $\sigma(k) = \max(\frac{N}{k} \log \frac{N}{k}, nc_m \cdot M)$.*

PROOF. For details, please refer to Appendix I in [37]. \square

Example 6.6. Let us continue using the setting in Example 6.1. After selecting F and D to V_{de} , we select the core vertices among the rest of vertices. We set $\epsilon = 80\%$ to guarantee that no fewer than $\epsilon |V| = 4.8$ vertices are covered. $MS(\cdot)$ and the result of partial set cover are shown in Table 2. The final cover is $V'_{co} = \{B, E\}$ with cost 2. To maintain a 2-skip cover, we initialize $V_{co} = \{A, B, C, E\}$ and iteratively check each vertex $u \in V - V_{de} - V'_{co} = \{A, C\}$. As none of these removals violates the 2-skip cover, they are removed from V_{co} . Removing B or E violates attribute (ii) and we output $V_{co} = \{B, E\}$.

Finally, as V_{co} , V_{su} , and V_{de} is a partition of V , $V_{su} = V - V_{co} - V_{de}$.

Table 2: Candidate serving set and partial set covers.

ID	A	B	C	E
$MS(\cdot)$	$\{A, D\}$	$\{A, B, C\}$	$\{C, F\}$	$\{D, E, F\}$
Partial set cover	None	None	$\{C, A, B\}$	$\{E, B\}$
Cost	∞	∞	3	2

6.4 Construction of HMPO colored Graph and Fast Query

After obtaining the three levels of vertices, we use them to construct the hierarchical graph $G_h(V, E_h)$. In this subsection, we first concentrate on the formulation of edge E_h . Then, we show that how to compute shortest path queries fast on G_h .

HMOG graph construction. First, for the defective vertices, we “discard” them so there is no edge for drivers to come nor leave them. We get an updated car graph G_c by removing V_{de} and its adjacent vertices from G_c .

Recall that V_{co} is a k -skip cover of G_c . Based on that, all the shortest distance queries can be answered by G^* efficiently.

Core vertices serve as the skeleton of the original road network. Borrowing the definition in the previous work [33], for each $u \in V_{co}$, one can find its k -skip neighbors and build *super-edges*, leading to a new graph for fast query. To be more specific, one can construct a graph $G^* = (V_{co}, E_{cc})$, which has 2 vertex sets V_{co} and V_{su} and 3 edge sets E_{cs} , E_{sc} , and E_{ss} . Any query result between these core vertices on the new graph is the same as that on the original graph. Query start from or aim at sub-level vertices can be answered by temporally extending the graph with super-edges between sub-level vertex and its nearby core vertices.

Here we summarize the construction steps. (1) Instead of the k -skip neighbors of u , $N_k(u)$, we can find a superset $M_k(u)$ of $N_k(u)$ efficiently according to [33]. (2) We build super-edge, which is weighted by the shortest path distance between two vertices, from $u \in V_{co}$ to each $v \in M_k(u)$. Add these super-edges between core vertices to E_{cc} . (3) For $u \in V_{su}$, we find its $M_k(u)$ on $G - V_{de}$ and $M'_k(u)$ on the reversed graph. The super-edges built to $M_k(u)$ are stored into edge set E_{sc} , each represents a path from a sub-level vertex to a core vertex. Similarly, the super-edges built to $M'_k(u)$ are stored into edge set E_{cs} for core-to-sub-level vertex pairs. (4) In the middle of finding $M_k(u)$ for each $u \in V_{su}$, every $v \in V_{su}$ sharing a shortest path with u without core vertex generates a sub-to-sub level super-edge, added to E_{ss} .

Fast queries. We can calculate the exact k -shortest path using the three edge sets [33]. Consider the following cases: (i) $u, v \in V_{co}$. We simply return the query result on graph $G^* = (V_{co}, E_{cc})$; (ii) $u, v \in V_{sub}$. First we check the super-edges of u in E_{ss} . If u can reach v directly, return the weight of super-edge. Otherwise, follow the general cases in (iii, iv); (iii) $u \in V_{co}$. We need an additional step to add u with its super-edges $(u, \cdot) \in E_{sc}$ into G^* ; (iv) $v \in V_{co}$. Add v with its super-edges $(\cdot, v) \in E_{cs}$ into G^* and return the query result. Its correctness has been proved. For details, please refer to [33].

In summary, our final output for HMPO graph $G_h(V, E_h)$ is consists of: V_{co} , which forms the graph G^* for query together with super-edges E_{cc} ; V_{su} , of which each vertex is added to G^* with edge sets E_{sc} , E_{cs} , E_{ss} temporally for query; V_{de} , which is only called at the route planning stage with MC . $E_h = E_{cc} \cup E_{cs} \cup E_{sc} \cup E_{ss}$ and $V = V_{co} \cup V_{su} \cup V_{de}$.

Example 6.7. Following the setting in Example 6.1, we construct the G_h in Figure 5. $V_{de} = \{D, F\}$ are marked in white and their edges are removed. $V_{co} = \{B, E\}$ are marked in dark red. Sub-level vertices $\{A, C\}$ are marked in light yellow. Two core-to-core edges between them are added to E_{cc} and marked in red. Two edges from B to A and C are added to E_{cs} and marked in grey. Two edges from A and C to B are added to E_{sc} and marked in blue.

7 HMPO GRAPH BASED INSERTOR

With the HMPO Graph, we introduce a new algorithm to boost the insertion phase for the solution of MORP problem.

Requests are served one-by-one, where the MP candidates of their origins and destinations are inserted into drivers' routes. With limited walking distance, the MP candidates $MC(u)$ for each $u \in V_p$ are close to each other. One interesting problem is, if we fail to insert a candidate $v \in MC(u)$, do the rest $MC(u) - \{v\}$ help? To answer it, we define a new distance correlation, which bounds the time saving of switching to any vertex in $MC(u) - \{v\}$. If deducting the saving still cannot meet the time limitation, we can safely prune the whole set. However, to derive it on the traditional graph, we need all the distances from $|V|$ sources even though vertices in $MC(u)$ are close to each other. Fortunately, we find that based on the k -skip structure in our hierarchical graph, an upper bound can be derived effectively. As shown in Figure 2, k -skip cover "cuts off" the shortest paths using the core vertices, which can be used as anchors to compare the difference between the paths from the same source to a pair of vertices. Most of the related works only use k -skip cover for faster query and vertex clustering [11]. To the best of our knowledge, this is the first paper to explore this property of k -skip cover graph for task assignment.

We demonstrate an example through this section in Figure 6. The network and edge weights are consistent with Figure 2(a), where red vertices belong to core vertices and yellow ones belong to sub-level vertices. In Figure 6(b), a rider is waiting to be picked up at origin v_{22} and a driver is at v_d . The MP candidates for v_{22} are $\{v_{22}, v_{32}\}$. We want to know that if driver cannot reach MP v_{22}/v_{23} in time, can we prune the other MP.

In the following subsections, we first extract the important distance correlations in subsection 7.1, then devise an effective algorithm *SMDB* for MORP problem in subsection 7.2.

7.1 Maximum Difference for Reaching Distances

Insertion for MPs induces many shortest path queries, which are from the *same* source to closely-located MPs. Here, we bound the distance differences between these queries. We first give related definitions for the bound. Given a vertex u and its MP candidates $MC(u)$, we have (i) one checker vertex $Ch(u) \in MC(u)$, and (ii) the distance bound, named *set maximum difference* ($SMD(u)$). We show that after computing the query of one long path from a vertex v to checker $Ch(u)$, (i.e., $SP_c(v, Ch(u))$), the distance from v to any other candidate vertices is no shorter than $SP_c(v, Ch(u)) - SMD(u)$.

We first introduce the above concepts. Then, as any shortest path is split by k -skip cover V^* , we can quickly derive the bound locally, where only the subgraph formed by surrounding V^* is needed.




Figure 5: HMPO Graph

First, we define *maximum difference* $MD(\cdot, \cdot)$ between two vertices as a general distance correlation as follows:

Definition 7.1. (Maximum Difference) Given a graph $G(V, E)$, for each pair of vertices $v_1, v_2 \in V$, the maximum difference for them is:

$$MD(v_1, v_2) = \max_{v_3 \in V} (SP_c(v_3, v_1) - SP_c(v_3, v_2))$$

In general, starting from any source $v \in V$, the largest difference between distances to reach two vertices is denoted as their MD .

When we insert a set of MP candidates of source or destination into a route, it is more reasonable to use a set-based relationship for maximum difference rather than pairwise MD . To be more specific, for the MP candidates $MC(u)$ of vertex u , by defining a checker vertex $Ch(u) \in MC(u)$, we want to know the Maximum Differences between checker $Ch(u)$ and all the vertices in the candidate set. The upper bound of these MD s is defined as *set maximum difference* (SMD). That is, $SMD(u) = \max_{v \in MC(u)} MD(Ch(u), v)$.

In the following part, we show that for a modified SMD with narrower source vertices, which are not sub-level vertices that have super-edges to $v \in V_M$, the k -skip structure can help us get SMD in $O(|V|)$ time. Mathematically, we define

$$MD^*(u_1, u_2, V_M) = \max_{v: (v, u) \notin E_{sc} \cup E_{ss}, u \in V_M} (SP_c(v, u_1) - SP_c(v, u_2)),$$

and

$$SMD(u) = \max_{v \in MC(u)} MD^*(Ch(u), v, MC(u)), \quad (3)$$

where E_{sc} and E_{ss} are super-edge sets in Section 6.4.

With a k -skip cover V^* , any shortest paths of length $> k$ ended at a vertex u must reach its surrounding V^* . Our core vertices V_{co} , also a k -skip cover, further builds super edges between u and these surrounding V_{co} . As shown in Figure 6(c), any paths ended at v_{22} and v_{32} need reach v_{31}, v_{42}, v_{21} , and v_{23} first. They are also the only 4 vertices having super edges with v_{22} and v_{32} in the HMPO graph, e.g., Figure 6(d). We check super edges and get the surrounding V_{co} for each vertex in an MP candidate set $MC(u)$ and define their union as $VC(u)$. In Figure 6, $VC(v_{22}) = \{v_{31}, v_{42}, v_{21}, v_{23}\}$. As vertices in $MC(\cdot)$ are close to each other, $VC(\cdot)$ is a small set.

Note that the distance from any $vc \in VC(u)$ to any $v \in MC(u)$ is already recorded in HMPO graph, which does not need to compute. We denote its cost as $CC(vc, v)$. By fixing a vertex $w \in MC(u)$ as checker vertex, we can compute the following Local Maximum Difference LMD quickly:

$$LMD(w) = \max_{v \in MC(u), vc \in VC(u)} (CC(vc, w) - CC(vc, v)),$$

LEMMA 7.2. *If we use $w \in MC(u)$ as the checker vertex, the Local Maximum Difference $LMD(w)$ is a valid Set Maximum Difference $SMD(u)$. Mathematically, if a vertex $l_c \in V$ has no super-edge $(l_c, v) \in E_{ss} \cup E_{sc}$ towards any $v \in MC(u)$, then $SP_c(l_c, w) - SP_c(l_c, v) \leq LMD(w)$.*

PROOF. We prove it by contradiction. Given a vertex $u \in V$ and its $MC(u)$ with outputs $Ch(u) = w$ and $SMD(u)$, assume that $\exists l_c \in V$ has no super-edges in $E_{ss} \cup E_{sc}$ to $MC(u)$, $\exists v \in MC(u)$ which satisfies




Figure 6: An Example for the Set Maximum Difference for Meeting Points

$SP_c(l_c, w) - SP_c(l_c, v) > SMD(u)$. We denote the last core vertex in the path from l_c to v as vc . Here we have

$$\begin{aligned} LMD(w) &< SP_c(l_c, w) - SP_c(l_c, v) \\ &\leq (SP_c(l_c, vc) + SP_c(vc, w)) - (SP_c(l_c, vc) + SP_c(vc, v)) \\ &= SP_c(vc, w) - SP_c(vc, v) \end{aligned}$$

As vc is the last core vertex along the path, vc must be a k -skip neighbors of $v \in MC(u)$. Thus, we have:

$$SP_c(vc, w) - SP_c(vc, v) = CC(vc, w) - CC(vc, v) \leq LMD(w)$$

Contradiction. The original lemma is proved. \square

For example, in Figure 6(f), if we choose v_{32} and check vertex, we can calculate $CC(v_{31}, v_{32}) - CC(v_{31}, v_{22}) = CC(v_{42}, v_{32}) - CC(v_{42}, v_{22}) = -1$ and $CC(v_{21}, v_{32}) - CC(v_{21}, v_{22}) = CC(v_{23}, v_{32}) - CC(v_{23}, v_{22}) = 1$. So $LMD(v_{32}) = 1$ when $vc = v_{21}$ or v_{23} . After we compute the cost from v_d to v_{32} as $SP_c(v_d, v_{32})$, we want to bound the cost from v_d to v_{22} (i.e., $SP_c(v_d, v_{22})$). In Figure 6(g), its last passed core vertex is $v_{23} \in VC$. we have $SP_c(v_d, v_{22}) = SP_c(v_d, v_{23}) + SP_c(v_{23}, v_{22}) \geq SP_c(v_d, v_{23}) + [SP_c(v_{23}, v_{32}) - LMD(v_{32})] \geq SP_c(v_d, v_{32}) - LMD(v_{32})$.

Now we can efficiently find SMD given a checker. Here, we show the detail in our HMPO Graph-based Maximum Difference Generator (HMDG), which finds the $Ch(\cdot)$ for each $MC(\cdot)$ that minimize $SMD(\cdot)$ when calculating LMD .

Algorithm sketch Lines 3-8 collect all the cost from vertices $vc \in VC$ to vertices $v \in MC(u)$. As $VC \subseteq V_{co}$, if $v \in V_{co}$, we check the its super edges from core to core vertices. The costs are stored in the form of dictionary $CC[vc][v]$; if $v \in V_{su}$, we check its super edges from sub to core vertices instead.

Secondly, we initialize $LMD(v) \rightarrow SMD$ as a dictionary, which can be further used to choose the checker with minimal SMD . Instead of fixing checker and calculating corresponding SMD , for each $vc \in VC$, we find the vertex v^- which has minimal $CC[vc][v^-]$ in line 11, that is, among vertices in $MC(u)$, v^- is the closest destination for vc . If we use v as checker, the SMD is the maximal of $CC[\cdot][v] - CC[\cdot][v^-]$, which should be minimized. Lines 12-14 enumerate MC and update the $LMD[v]$ if a higher maximum difference is found, that is, $LMD[v] < CC[vc][v] - CC[vc][v^-]$. Finally, $LMD[\cdot]$ saves the required SMD for each checker. We find the minimal value of LMD with its key v and return $Ch(u) = v$ and $SMD = LMD[v]$.

So our algorithm finds $Ch(u)$ with $SMD(u)$ such that: (i) for any source $vc \in VC$ from nearby core vertices, $\forall v \in MC(u)$, $SP_c(vc, Ch(u)) - SP_c(vc, v) \leq SMD(u)$; (ii) $Ch(u)$ is chosen among $MC(u)$ to minimize $SMD(u)$.

Time Complexity. The time costs of Lines 2 and 3 are $O(nc_m)$. There are $O(nc_m)$ iterations in lines 4-8 and each iteration costs $O(\sigma^*)$, where σ^* is the number of super edges a vertex has on average. Line 9 costs $O(nc_m)$ time. Line 11 costs nc_m to find the minimum value and lines 12-14 form $O(nc_m)$ iterations taking $O(1)$ time in each iteration. For the size of VC , we borrow the definition $\bar{\sigma}_k$ from [33], where $\bar{\sigma}_k$ is the average number of k -hop neighbors

Algorithm 1: HMPO graph-based Max Difference Generator

Input: HMPO graph $G_h = (V, E_h)$, MP candidate sets MC .

Output: Checker Ch and Set Max Diff SMD for each MP candidate set

```

1 foreach  $u \in V$  do
2   Build set  $VC = \{vc | (vc, v) \in E_{cs}, v \in MC(u)\}$ 
3   Initialize dictionary  $CC$  for costs from  $VC$  to  $MC(u)$ 
4   foreach  $v \in MC(u)$  do
5     if  $v \in V_{co}$  then
6       Check the super edges of  $v$ . For each edge  $\in E_{cc}$  from
7          $vc$ , record the costs into  $CC[vc][v]$ 
8     if  $v \in V_{su}$  then
9       Check the super edges of  $v$ . For each edge  $\in E_{cs}$  from
10         $vc$ , record the costs into  $CC[vc][v]$ 
11   Initialize  $LMD[v] = 0$  for all  $v \in MC(u)$ 
12   foreach  $vc \in VC$  do
13     Find the minimal in  $CC[vc][\cdot]$ , denote the key as  $v^-$ 
14     foreach  $v \in MC(u)$  do
15       if  $LMD[v] < CC[vc][v] - CC[vc][v^-]$  then
16          $LMD[v] = CC[vc][v] - CC[vc][v^-]$ 
17   Find the minimal value of  $LMD[v]$ , return  $Ch(u) = v$  and
18      $SMD(u) = LMD[v]$ 
19 return  $Ch, SMD$ 

```

of the vertices in V . Thus, there are $O(|VC|) = O(\bar{\sigma}_k nc_m)$ iterations in lines 10-14. Their total time complexity is $O(\bar{\sigma}_k nc_m^2)$. Line 15 cost $O(nc_m)$ to find the minimum. So the total time complexity of the big loop in lines 1-15 is $O(|V| (nc_m \sigma^* + \bar{\sigma}_k nc_m^2))$ and grows linearly with $|V|$, where both σ^* and $\bar{\sigma}_k$ depend on the structure of the road network instead of the size.

Note that, though we only discuss the distance relationships from the same source to different destinations, similar property still hold for queries from different sources to the same destination. We further emphasize that such a relationship is not only applicable for meeting points, but a property of k -skip cover. It is useful for many crowdsourcing scenarios, which have shortest path queries with close sources and destinations.

7.2 SMD-Boost Algorithm

We use SMD to boost insertion in this subsection.

Recall that whenever we try to insert a request r_j into a route, the rider should be picked up no later than tp_j . Considering that we try to insert one pick-up into a fixed position of route. If we use $Ch(s_j)$ as the MP, it is not insertable if the time to pick up rider at $Ch(s_j)$, denoted as t_c here, is later than tp_j . The worker needs

Algorithm 2: SMDBoost

Input: a driver w_i with route S_{w_i} , request r_j , MP candidate set MC , set maximum difference SMD , checker set Ch , dead vertices DV

Output: a route S_w^* for the driver w and updated DV

- 1 **if** Driver's location $l_i \in DV$ **then**
- 2 Return S_{w_i} and DV without insertion
- 3 Generate arriving time $arv[\cdot]$ for S_{w_i}
- 4 Collect all sub-level vertices which have super-edges to vertices in $MC(s_j)$ into set Ne
- 5 The largest index to insert pick-up: $id^* = |S_{w_i}|$
- 6 **foreach** $v \in S_{w_i}$ **do**
- 7 **if** $v \in Ne$ **then**
- 8 Continue
- 9 **if** $arv[v] + SP_h(v, Ch(s_j)) - SMD(Ch(s_j)) \geq tp_j$ **then**
- 10 **if** $v = l_i$ **then**
- 11 Add l_i to DV . Insertion fails and returns Null
- 12 Record $id^* = idx(v) - 1$
- 13 Break
- 14 Insert r_j with adapted insertion algorithm where insertion indexes of pick-ups larger than id^* are pruned.
- 15 **return** S_w^*, DV

to pick up rider at least $t_c - tp_j$ earlier. In the last subsection, we find the maximum time saving for each checker, SMD . If SMD of $Ch(s_j)$ is smaller than $t_c - tp_j$, no matter which MP of s_j is used for insertion, we cannot pick up the request by tp_j . In such case, there is no need to try other MPs for insertion with the help of SMD .

We illustrate our algorithm *SMDBoost* for the insertion phase in Algorithm 2. Note that we add one more set for pruning, dead vertices DV . It means that no driver w_i with current location $l_i \in DV$ can serve this rider. We initialize $DV = \emptyset$ for each new request. Assume that we try to insert rider r_j into the route of driver w_i . First, if $l_i \in DV$, we can prune driver w_i . Otherwise, we derive arriving time $arv[\cdot]$ for each route vertex according to [36]. In line 4, all the sub-level vertices which have super-edges towards $MC(s_j)$ are collected into a set Ne , which covers vertices that cannot be pruned. This is because that distances between these vertices and an MP can be arbitrarily short through super-edges in $E_{sc} \cup E_{ss}$ and not follow the definition of SMD in Equation 3.

Pruning strategy in lines 5-12 finds the largest index id^* to insert any pick-up in $MC(s_j)$. We initialize $id^* = |S_{w_i}|$. Then we check each vertex $v \in S_{w_i}$ in order. As the distances from vertices in Ne to some MPs in $MC(s_j)$ can be arbitrarily short, if $v \in Ne$, it is possible to insert rider after v . In this case, we continue to check the next insertion position in line 8. For each vertex $v \notin Ne$, it can be viewed as a source for $MC(u)$ subject to $SMD(u)$. So if $arv[v] + SP_h(v, Ch(s_j)) - SMD(Ch(s_j)) \geq tp_j$, according to Lemma 7.2, inserting any MP after v misses the deadline tp_j . Insertion position can only be smaller than index of v denoted as $idx(v)$, that is, $id^* = idx(v) - 1$. A special case is that if inserting after the driver's current location l_i cannot catch t_j , driver w_i and all the other drivers at l_c currently can not serve r_j . So we add l_i into DV for future pruning.

Table 3: Setting of Dataset and Model

Parameters	Settings
Number of vertices of NYC	57030
Number of edges of NYC	122337
Number of valid requests of NYC	277410

After the checking phase, we insert r_j without checking indexes after id^* for the pick-up point. The base algorithm adapts the linear insertion algorithm [36] for MPs.

Time Complexity. Line 3 is a linear operation according to [36]. Line 4 costs $O(nc_m)$. There is an $O(|S_{w_i}|)$ loop in lines 5-12. All other lines from 1 to 12 cost $O(1)$. As the algorithm in line 13 is linear, the total time complexity is $O(|S_{w_i}|)$.

LEMMA 7.3. *Pruning Algorithm 2 has no performance loss.*

PROOF. For details, please refer to Appendix J in [37]. \square

For example, in Figure 6(g), assume that rider r_i needs to be picked up in 10 time steps. If the shortest path from v_d to v_{32} is 12 time steps, as SMD is 1, the shortest time to reach any other MP candidates is bounded by $12 - 1 = 11 > 10$. We can stop checking other MP candidate(s) (v_{22}). Any other worker at v_d currently cannot serve r_i . v_d is marked as a dead vertex and added to DV .

8 EXPERIMENTAL STUDY

8.1 Experimental Methodology

Data set. We use both real and synthetic data to test our HMPO Graph-Based solution. Specifically, for the real data, we use a public data set NYC [3]. It is collected from two types of taxis (yellow and green) in New York City, USA. We use all the request data on December 30th to simulate the ridesharing requests in our experiments. Each taxi request in NYC contains the latitudes/longitudes of its source/destination locations, its starting timestamp, and its capacity. We can generate a ridesharing request and initialize its locations, release time, and capacity correspondingly.

In addition, we derive the distribution of requests of all the NYC requests in December and generate 5 synthetic datasets (Syn) with the size of requests as 100k, 200k, 400k, 800k, and 1000k. We extract the road network of NYC from Geofabrik [5]. It includes the labels of roads for both driving and walking. We clean it into the directed car graph G_c and passenger graph G_p according to road labels. The road network is widely used in existing ridesharing studies [36].

The settings of our dataset are summarized in Table 3.

Implementation. We follow the common settings for simulating ridesharing applications in [9, 23, 36]. While building the graph for road network, the weights of edges are set to their time cost (divide the road length on Geofabrik by the velocity of its road type). The walking speed is set to 3.6km/h, while the driving speed varies according to the road type in OSM, i.e., 60% of the maximum legal speed limit in their cities [7]. For each request, we map its source and destination to the closest vertex in the road network. The initial location of each driver is randomly chosen.

8.2 Setting Default Values for Parameters

Table 4 summarizes the major parameter settings for our experiments, where default values are in bold font. Other hyperparameters are discussed in Appendix K in [37] with extensive experiments.

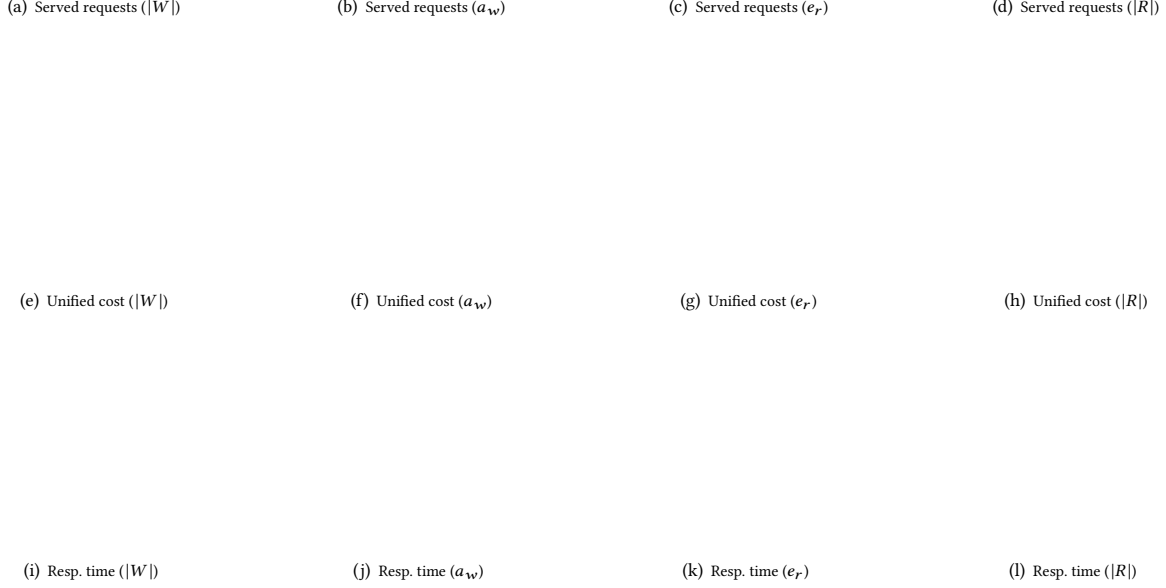


Figure 7: Performance of varying number of drivers $|W|$, capacity a_w , deadline coefficient e_r , and number of requests $|R|$

Table 4: Parameter Settings.

Parameters	Settings
Deadline Coefficient e_r	0.1, 0.2, 0.3 , 0.4, 0.5
Capacity a_w	2, 3 , 4, 7, 10
Driving Distance Weight α	1
Walking Distance Weight β	0.5, 1 , 1.5, 2
Penalty p_o	3, 5, 10, 15, 30
Number of drivers $ W $	5k, 10k, 20k , 30k, 40k
Number of requests $ R $	100k, 200k, 400k, 800k, 1000K

Note that for a mix mode where some riders are not willing to walk, our algorithm *SMDB* can directly handle it by manipulating its input. Instead of generating meeting point candidates according to Section 5, we restrict the candidates as source $\{s_j\}$ for pick-up point and destination $\{e_j\}$ for drop-off point, where the walking distance is 0. The candidates are fed into *SMDB*, which still works regardless of the number of candidates.

The experiments are conducted on a server with Intel(R) Xeon(R) E5 2.30GHz processors with hyper-threading enabled and 128 GB memory. The simulation implementation is single-threaded, and the total running time is limited to 14 hours for NYC. In reality, a real-time solution should stop before its time limit (24 hours for us) [23, 36]. All the algorithms are implemented in Java 11. We first preprocess our road network according to Section 6. The modified vertices and weighted edges can be loaded directly for route planning. We boost the shortest distance and path queries with an LRU cache according to the previous studies [23, 36].

Compared Algorithms. We compare *SMDB* with the state-of-the-art algorithms for route planning of ride-sharing.

- **GreedyDP** [36]. It uses a greedy strategy for route planning without MPs. Each request is assigned to the feasible new route with minimum increased cost.
- **Kinetic Tree (KT)**. For each driver, it saves all the possible routes for the assigned request using a structure, Kinetic Tree. The new request is inserted by traversing and updating the tree, which is more local optimal but slower compared with GreedyDP.
- **BasicMP**. It is an extension from GreedyDP by adapting MPs to solve the MORP problem.
- **First Serve (FS)**. A variant of BasicMP, where each request is directly assigned to the first driver who can serve it.
- **HSRP**. It uses the HMPO Graph to improve the effectiveness of BasicMP without pruning.

Metrics. All the algorithms are evaluated in terms of total unified cost, served requests $|R|$ and response time (average waiting time to arrange a request, *resp. time* for short), which are widely used as metrics in large-scale online ride-sharing proposals [23, 26, 36].

8.3 Experimental Results

In this subsection, we present the experimental results.

Impact of Number of Drivers $|W|$. The first column of Figure 7 presents the results with different numbers of drivers in NYC. Compared with GreedyDP, BasicMP outperforms it in terms of the number of served requests by 6.6% to 12.7% with the help of

MPs, while SMDB and HSRP outperforms it by 21.4% to 29.9%. More requests are served with more drivers, results in a decrease of unified costs and an increase in the served rates of all the algorithms. BasicMP decreases the cost by 2.7% to 13.2% and SMDB decreases it by 4.5% to 32.7%. GreedyDP and FS run the fastest, where GreedyDP does not use MPs and FS assigns each request to the first feasible worker suboptimally. SMDB runs faster than other local-optimal MP-based methods (HSRP and BasicMP) with a *resp. time* < 0.08s.

Impact of Capacity of Drivers a_w . The second column of Figure 7 presents the effect of the capacities of drivers. BasicMP serves 8.6% to 12.2% more requests than GreedyDP with 5.1% to 9.2% less cost. SMDB outperforms other algorithms on both serving rate, 27.3% to 33.1% higher than GreedyDP, and unified cost, 13.0% to 20.5% less than GreedyDP. With a larger capacity, all the algorithms serve more requests with lower costs. However, the improvement with capacity from 4 to 10 is not as significant as from 2 to 4. The reason is that the extra spaces are mostly wasted. GreedyDP and FS still runs faster but serve fewer requests. SMDB costs less time than BasicMP and HSRP. KT is much slower than other methods but serves the fewest requests.

Impact of Deadline Coefficient e_r . The third column of Figure 7 shows the results of varying the deadline coefficient e_r . With larger e_r , all the algorithms serve more requests with a lower unified cost. SMDB still serves more requests with lower cost, which outperforms GreedyDP by serving 26.6% to 30.3% more requests than GreedyDP and decreasing its cost by 9.8% to 18.7%. With meeting points, BasicMP increases the number of served requests by 8.0% to 12.6% and decreases the cost by 3.4% to 9.2% compared with GreedyDP. With a larger deadline coefficient e_r , they find more feasible routes and serve more requests. KT builds a tree to record all the possible routes for each driver, which even fails to finish with $e_r = 0.5$. Time cost increases with larger e_r as each request can find more feasible candidate routes. GreedyDP and FS still run the fastest and SMDB is faster than BasicMP and HSRP.

Impact of Number of Requests $|R|$. The fourth column of Figure 7 displays the results on different sizes of synthetic requests. The datasets are generated based on the distribution of all the NYC requests in December. We increase the number of requests from 10k to 100k to test the extreme case that drivers are fully busy, i.e., the number of served requests does not increase with 100k arrival requests compared with 80k requests for all the algorithms.

When drivers are not fully busy with requests fewer than 80k, all the algorithms serve more requests with a lower unified cost as the $|R|$ increases. Comparing with each other, SMDB serves 7.3% to 28.4% more requests than GreedyDP and decreasing its cost by 10.6% to 21.8%. BasicMP works weaker that increases the number of served requests by 32.5% to 10.8% and decreases the cost by 5.0% to 14.8% compared with GreedyDP. Response time of KT drops quickly with more requests while it rejects more requests. SMDB has a closer response time compared with GreedyDP and FS with increasing $|R|$, where HSRP and BasicMP are the slowest methods.

Summary of Results.

- Our SMDB algorithm can serve 7.3% to 33.1% more requests than the state-of-art algorithm [36]. The unified cost is decreased by 4.5% to 32.7%. These results validate the effectiveness of our algorithm in large scale datasets.

- With MPs, BasicMP, HSRP, and SMDB outperform GreedyDP and KT. Suboptimal MP-based First Serve method also has competitive performance compared to GreedyDP. SMDB potentially arranges more requests on the highway and prunes candidates, which uses less time to serve more requests than BasicMP and HSRP. On the real-world dataset, SMDB can save 11% to 24% time compared with HSRP. On the synthetic dataset where a large amount of requests is test, SMDB saves more time (up to 31% compared with HSRP) given more requests, which is important to handle heavy traffic cases. With a response time lower than 0.2 seconds, SMDB is acceptable to be used as a real-time solution for ridesharing tasks.

9 CONCLUSION

In this paper, we propose the MORP problem, which utilizes meeting points for better ride-sharing route planning. We formulate a modified objective function to cover the cost of walking. We prove that the MORP problem is NP-hard and there is no polynomial-time algorithm with a constant competitive ratio for it. To cut the search space of meeting points for fast assignments, we devise an algorithm to prepare meeting point candidates for each vertex. Besides, we construct an HMPO graph with hierarchical order on vertices for fast route planning, which takes the advantage of flexibility from meeting points to improve efficacy. Based on it, we propose SMDB algorithm to solve MORP problem effectively and efficiently. Extensive experiments on real and synthetic datasets show that our proposed solution outperforms baseline and the state-of-the-art algorithms for traditional ridesharing in effectiveness greatly without losing too much efficiency. Our paper is provided as a comprehensive theoretical reference for optimizing route planning with meeting points in ridesharing.

ACKNOWLEDGMENTS

Peng Cheng's work is partially supported by the National Natural Science Foundation of China under Grant No. 62102149, Shanghai Pujiang Program 19PJ1403300 and Open Foundation of Key Laboratory of Transport Industry of Big Data Application Technologies for Comprehensive Transport. Libin Zheng is supported by the National Natural Science Foundation of China No. 62102463, the Natural Science Foundation of Guangdong Province of China No. 2022A1515011135 and the Fundamental Research Funds for the Central Universities, Sun Yat-sen University No. 22qntd1101. Lei Chen's work is partially supported by National Key Research and Development Program of China Grant No. 2018AAA0101100, the Hong Kong RGC GRF Project 16209519, CRF Project C6030-18G, C1031-18G, C5026-18G, CRF C2004-21GF, AOE Project AoE/E603/18, RIF Project R6020-19, Theme-based project TRS T41-603/20R, China NSFC No. 61729201, Hong Kong ITC ITF grants ITS/044/18FX and ITS/470/18FX, Guangdong Basic and Applied Basic Research Foundation 2019B151530001, Microsoft Research Asia Collaborative Research Grant, HKUST-NAVER/LINE AI Lab, Didi-HKUST joint research lab, HKUST-Webank joint research lab grants and HKUST Global Strategic Partnership Fund (2021 SJTU-HKUST) Wenjie Zhang's work is partially supported by the ARC Future Fellowship FT210100303. Corresponding author: Peng Cheng.

REFERENCES

- [1] 2019. [Online] Express Pool. <https://www.uber.com/us/en/ride/express-pool/>.
- [2] 2019. [Online] Uber Express Just like a Bus. <https://gizmodo.com/i-tried-uber-s-new-pool-express-service-and-honestly-j-1823190462>.
- [3] 2021. [Online] TLC Trip Record Data. <https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page>.
- [4] 2022. [Online] Didi Chuxing. <http://www.didichuxing.com/>.
- [5] 2022. [Online] Geofabrik. <https://download.geofabrik.de/>.
- [6] 2022. [Online] Open Street Map. <https://www.openstreetmap.org/>.
- [7] 2022. [Online] Speed limits in the United States – Wikipedia. https://en.wikipedia.org/wiki/Speed_limits_in_the_United_States.
- [8] Javier Alonso-Mora, Samitha Samaranayake, Alex Wallar, Emilio Frazzoli, and Daniela Rus. 2017. On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences* 114, 3 (2017), 462–467.
- [9] Mohammad Asghari, Dingxiong Deng, Cyrus Shahabi, Ugur Demiryurek, and Yaguang Li. 2016. Price-aware real-time ride-sharing at scale: an auction-based approach. In *Proceedings of the 24th ACM SIGSPATIAL international conference on advances in geographic information systems*. 1–10.
- [10] Mohammad Asghari and Cyrus Shahabi. 2017. An on-line truthful and individually rational pricing mechanism for ride-sharing. In *Proceedings of the 25th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. 1–10.
- [11] Peng Cheng, Hao Xin, and Lei Chen. 2017. Utility-Aware Ridesharing on Road Networks. In *SIGMOD*. ACM.
- [12] Blerim Cici, Athina Markopoulou, and Nikolaos Laoutaris. 2015. Designing an on-line ride-sharing system. In *Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems*. 1–4.
- [13] Paul Cziotka, Dirk C Mattfeld, and Monika Sester. 2017. GIS-based identification and assessment of suitable meeting point locations for ride-sharing. *Transportation research procedia* 22 (2017), 314–324.
- [14] Pedro M d'Orey, Ricardo Fernandes, and Michel Ferreira. 2012. Empirical evaluation of a dynamic and distributed taxi-sharing system. In *2012 15th International IEEE Conference on Intelligent Transportation Systems*. IEEE, 140–146.
- [15] Elif Eser, Julien Monteil, and Andrea Simonetto. 2018. On the tracking of dynamical optimal meeting points. *IFAC-PapersOnLine* 51, 9 (2018), 434–439.
- [16] Esteban Feuerstein and Leen Stougie. 2001. On-line single-server dial-a-ride problems. *Theoretical Computer Science* 268, 1 (2001), 91–105.
- [17] Stefan Funke, André Nusser, and Sabine Storandt. 2014. On k-path covers and their applications. *Proceedings of the VLDB Endowment* 7, 10 (2014), 893–902.
- [18] Rajiv Gandhi, Samir Khuller, and Aravind Srinivasan. 2004. Approximation algorithms for partial covering problems. *Journal of Algorithms* 53, 1 (2004), 55–84.
- [19] Timo Gschwind and Stefan Irnich. 2015. Effective handling of dynamic time windows and its application to solving the dial-a-ride problem. *Transportation Science* 49, 2 (2015), 335–354.
- [20] Anupam Gupta, Mohammad Taghi Hajiaghayi, Viswanath Nagarajan, and R. Ravi. 2010. Dial a Ride from k -forest. *ACM Trans. Algorithms* (2010).
- [21] Wesam Mohamed Herbawi and Michael Weber. 2012. A genetic and insertion heuristic algorithm for solving the dynamic ride-matching problem with time windows. In *Proceedings of the 14th annual conference on Genetic and evolutionary computation*. 385–392.
- [22] Sin C Ho, Wai Yuen Szeto, Yong-Hong Kuo, Janny MY Leung, Matthew Petering, and Terence WH Tou. 2018. A survey of dial-a-ride problems: Literature review and recent developments. *Transportation Research Part B: Methodological* 111 (2018), 395–421.
- [23] Yan Huang, Favyen Bastani, Ruoming Jin, and Xiaoyang Sean Wang. 2014. Large Scale Real-time Ridesharing with Service Guarantee on Road Networks. *Proceedings of the VLDB Endowment* 7, 14 (2014).
- [24] Jang-Jei Jaw. 1984. *Solving large-scale dial-a-ride vehicle routing and scheduling problems*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [25] Alexander Kleiner, Bernhard Nebel, and Vittorio Amos Ziparo. 2011. A mechanism for dynamic ride sharing based on parallel auctions. In *Twenty-second international joint conference on artificial intelligence*.
- [26] Shuo Ma, Yu Zheng, and Ouri Wolfson. 2013. T-share: A large-scale dynamic taxi ridesharing service. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*. IEEE, 410–421.
- [27] Shuo Ma, Yu Zheng, and Ouri Wolfson. 2015. Real-Time City-Scale Taxi Ridesharing. *IEEE Transactions on Knowledge & Data Engineering* 27, 07 (2015), 1782–1795.
- [28] Masayo Ota, Huy Vo, Claudio Silva, and Juliana Freire. 2016. Stars: Simulating taxi ride sharing at scale. *IEEE Transactions on Big Data* 3, 3 (2016), 349–361.
- [29] Zachary B. Rubinstein, Stephen F. Smith, and Laura Barbulescu. 2012. Incremental Management of Oversubscribed Vehicle Schedules in Dynamic Dial-A-Ride Problems. In *AAAI*.
- [30] Douglas Oliveira Santos and Eduardo Candido Xavier. 2013. Dynamic taxi and ridesharing: A framework and heuristics for the optimization problem. In *Twenty-Third International Joint Conference on Artificial Intelligence*.
- [31] Mitja Stiglic, Niels Agatz, Martin Savelsbergh, and Mirko Gradisar. 2015. The benefits of meeting points in ride-sharing systems. *Transportation Research Part B: Methodological* 82 (2015), 36–53.
- [32] SUMC. 2018. What is shared-use mobility? <https://goo.gl/3Jw6z7>.
- [33] Yufei Tao, Cheng Sheng, and Jian Pei. 2011. On k-skip shortest paths. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*. 421–432.
- [34] Raja Subramaniam Thangaraj, Koyel Mukherjee, Gurulingesh Ravari, Asmita Metrewar, Narendra Annamaneni, and Koushik Chattopadhyay. 2017. Xhare-a-ride: A search optimized dynamic ride sharing system with approximation guarantee. In *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*. IEEE, 1117–1128.
- [35] Yongxin Tong, Libin Wang, Zhou Zimu, Bolin Ding, Lei Chen, Jieping Ye, and Ke Xu. 2017. Flexible online task assignment in real-time spatial data. *Proceedings of the VLDB Endowment* 10, 11 (2017), 1334–1345.
- [36] Yongxin Tong, Yuxiang Zeng, Zimu Zhou, Lei Chen, Jieping Ye, and Ke Xu. 2018. A Unified Approach to Route Planning for Shared Mobility. *PVLDB* (2018).
- [37] Jiachuan Wang, Peng Cheng, Libin Zheng, Lei Chen, and Wenjie Zhang. 2022. Online Ridesharing with Meeting Points [Technical Report]. *arXiv preprint arXiv:2209.10892* (2022). <https://arxiv.org/abs/2209.10892>
- [38] Nigel HM Wilson, RW Weissberg, BT Higonnet, and J Hauser. 1975. *Advanced dial-a-ride algorithms*. Technical Report.
- [39] Nigel Henry Moir Wilson, Richard Wayne Weissberg, and John Hauser. 1976. *Advanced dial-a-ride algorithms research project*. Technical Report.
- [40] San Yeung, Evan Miller, and Sanjay Madria. 2016. A flexible real-time ridesharing system considering current road conditions. In *2016 17th IEEE International Conference on Mobile Data Management (MDM)*, Vol. 1. IEEE, 186–191.
- [41] Meng Zhao, Jiateng Yin, Shi An, Jian Wang, and Dejian Feng. 2018. Ridesharing problem with flexible pickup and delivery locations for app-based transportation service: Mathematical modeling and decomposition methods. *Journal of Advanced Transportation* 2018 (2018).
- [42] Libin Zheng, Lei Chen, and Jieping Ye. 2018. Order dispatch in price-aware ridesharing. *Proceedings of the VLDB Endowment* 11, 8 (2018), 853–865.