Incentivizing Participation in Peer-to-Peer Ride-Sharing Platform

PI: Fei Fang (Institute for Software Research, School of Computer Science)
Co-PI: Alexandre Jacquillat (Heinz College of Information Systems and Public Policy)

FINAL RESEARCH REPORT

Contract # 69A3551747111

DISCLAIMER

PIs and Contributors

1. Project Investigator: Fei Fang (Carnegie Mellon University, ORCID: 0000-0003-2256-8329)

2. Co-Project Investigator: Alexandre Jacquillat (Carnegie Mellon University, ORCID: 0000-0002-2352-7839)

3. Contributors:

   (a) Hoon Oh (Carnegie Mellon University, ORCID: 0000-0002-0228-652X)

   (b) Yanhan Tang (Carnegie Mellon University, ORCID: 0000-0002-7372-9738)

   (c) Zong Zhang (Carnegie Mellon University, ORCID: 0000-0003-3672-2825)

1 Problem Description

By providing more flexible commuting options than public transportation and private vehicles, on-demand ride-hailing and ride-sharing platforms have become increasingly popular in urban areas. However, the availability and affordability of on-demand transportation remain much more limited in sub-urban and rural areas. In practice, an overwhelming majority of commuting trips rely on self-driving with private vehicles. An increasingly popular option to promote alternative forms of mobility in non-urban areas lies in peer-to-peer (P2P) ride-sharing: by bringing together commuters traveling along similar routes at similar times, P2P ride-sharing can enhance mobility while reducing the costs of transportation, traffic congestion, and greenhouse gas emissions. Moreover, P2P ride-sharing can also improve access to basic needs for disadvantaged populations with limited car ownership.

To be successful, P2P ride-sharing platforms require effective algorithms to match rider requests with available drivers. The topic of matching in ride-sharing platforms has attracted considerable research interest in recent years Ozkan and Ward [2017]; Bertsimas *et al.* [2019]; Santi *et al.* [2014]; Bei and Zhang [2018]; Alonso-Mora *et al.* [2017], building upon related problems such as the dial-a-ride problem (DARP) Cordeau [2006]; Parragh *et al.* [2010] and the vehicle routing problem with time windows (VRPTW) Cordeau *et al.* [2007], matching in spatial-temporal networks, and is sometimes studied jointly with the topic of pricing Bimpikis *et al.* [2016]. However, there is only limited research for P2P ride-sharing without direct payments from riders to drivers Masoud and Jayakrishnan [2017], which can provide viable solutions for a community, e.g., the community of residents from close-by regions and employees of the same company, featuring drivers who have their own travel plans and are willing to share part of their trips with riders. Moreover, most existing work in matching in P2P ride-sharing only considers the flexibility windows of the drivers and riders ("users", henceforth) as constraints and ignores users' preferences and incentives for participation. In addition, the predominant objective used in this setting is to minimize total costs (e.g., travel costs and inconvenience costs). However, such approaches do not capture the impact of matching decisions on individual users, including the fairness among users and whether or not the users will accept the matching outcome. In this work, we address these limitations.

In contrast to previous work, we focus on matching riders and drivers in P2P ride-sharing without payment and study the problem from a user-centric perspective, with the objective of balancing system-wide efficiency and user satisfaction. We elevate users' preferences as a first-order concern by considering the users' preferred travel times and modeling drivers' altruism—motivated by the community-based P2P context. Furthermore, we explicitly consider the fairness and stability of the matching outcome. Fairness is formalized by guaranteeing that every rider is matched with at least a threshold probability. Stability is formalized by guaranteeing that users have no incentive to reject the current matching and take an alternative transportation option on his own or engage in independently shared rides outside of the platform. We develop a multi-objective framework that quantifies the trade-offs between efficiency, fairness, and stability.

2 Methodology

First, to compute a matching between drivers and riders to maximize system-wise efficiency, we extend the Request-trip-vehicle (RTV) framework from Alonso-Mora *et al.* [2017] to incorporate users' preferred times in P2P ride-sharing; the resulting model is computationally complex, and we propose a pruning algorithm to enhance its performance. Second, we formalize the notions of fairness and stability in P2P ride-sharing building upon a user utility model that incorporates altruism. We also design algorithms for computing efficient solution given fairness and stability constraints. Here we provides a high-level description of our models and algorithms. More detailed description can be found in the corresponding publications.

In our basic model, each rider and driver is characterized by his origin, destination, a time window describing the earliest possible departure time and latest possible arrival time, and value of the trip. In addition, we consider his preferred departure time and his maximum acceptable detour time. The goal is to match the drivers to disjoint subsets of riders and determine the schedule for each driver (an ordered sequence of node-time pairs describing how a single driver travels to pick up and drop off each rider matched to him) so as to minimize the total travel cost for all the participants.

The RTV framework has three main parts: (a) construction of all feasible (d, S) pairs where d is a driver and S is a subset of riders; (b) computation of the minimum cost (c_{dS}) for each feasible pair (d, S); and (c) computation of the cost-minimizing matching. We adopt this framework to our problem. The main challenge in applying the framework to our problem lies in (b). Since we incorporate the novel aspect of users' preferred times, existing approaches cannot be directly applied and we propose a novel algorithm TripCost for computing c_{dS}.

More specifically, in part (a), for each driver d, we incrementally construct a list of subsets of riders that are feasible with d by gradually increasing the size of the subset. In each step, we add one rider to an existing feasible subset S of size $h - 1$, getting a subset S' of size h. If the subsets of S' with $h - 1$ riders are all feasible with d, i.e., already in the list, algorithm TripCost of part (b) is called to further verify the compatibility of S'. When no more subsets with size h can be constructed, we will construct subsets with size $h+1$. After constructing subsets of sizes $h \leq |\mathcal{R}|$ where $\mathcal{R}$ is the total number of riders, we obtain the set of feasible (d, S) pairs.

For part (b), we develop $\mathsf{TripCost}(d, S)$, a novel tree-search based algorithm to compute c_{dS}. Each node of the tree represents a partial route, i.e., an ordered list of stops that the driver needs to visit, with the root representing the origin of the driver. In each step in the tree-search, we compute a lower bound of the cost of a full route that extends the current partial route by checking the must-visit stops remaining to be visited, and prune the branch if it is worse than the optimal solution found so far. Each leaf node represents a full route, and we need to determine the optimal timing to visit each stop. For this purpose, we design a dedicated mixed integer linear program (MILP) based on a constructed time-location graph with duplicated edges.

We also use two pruning methods to improve TripCost. First, we develop dynamic programming (DP)-based approach to compute a tighter lower bound of the route at the leaf node before computing TripMILP, which enables us to avoid unnecessary MILP computations. Intuitively, each state of the DP is a tuple containing location and time, and the DP computes the minimum cost to reach the location at that time based on previous states. Second, for each (d, S) pair, we learn promising routes from the optimal schedule found for (d, S') where $S' \subset S$ and (d', S) where $d' \neq d$. By getting these promising routes, we avoid unnecessary tree search, thus the pruning efficiency increases significantly.

In part (c), we use a standard MILP for matching problem (same as in Alonso-Mora *et al.* [2017]) to find the cost-minimizing matching given the c_{dS} for all the feasible (d, S) pairs. The variables used in the MILP include x_{dS} and y_r. $x_{dS} = 1$ if (d, S) is matched and 0 otherwise; and $y_r = 1$ if rider r is unsatisfied and 0 otherwise.

Further, we formalize fairness by maximizing the lowest probability of matching, across all riders in the system. We also consider stability at the individual-level, i.e., individual rationality. A driver or rider is *individually rational* (IR) if he/she does not get worse utility by participating in the matching system.

When considering fairness and stability requirements as constraints, finding the optimal matching becomes more challenging and the system-wise efficiency of the optimal matching will decrease. We designed an algorithm based on column generation to find the optimal matching satisfying θ-fairness, i.e., the probability of a rider being matched is at least θ (as long as there is a driver that can be matched). We also modify the steps in the RTV-based algorithm to incorporate the IR constraint which can be described as a set of linear constraints.

3 Findings

We analyze the price of fairness and stability and show that the price of fairness and stability can be arbitrarily large, and moreover, fairness can result in a very unstable market without external incentives.

We evaluate the proposed algorithms through extensive experiments. Results show that solutions with enhanced fairness and stability can be found in shorter, or equivalent, computational times as solutions based on efficiency objectives exclusively. Furthermore, suitable external incentives can result in fair and stable outcomes.

4 Conclusion

Unlike prior work, this project focuses on user-centric matching in P2P ride-sharing. (a) We propose a model that incorporates users' preferred times; (b) we incorporate fairness and stability in our P2P ride-sharing model; (c) we provide theoretical and experimental results showing trade-offs between fairness, stability, and cost efficiency. While, in theory, the price of fairness and the price of stability can be arbitrarily large, experimental results show that fair and stable solutions can be obtained through moderate efficiency losses and in reasonable computational times.

Finally, we point out two potential future research directions. The first one is to design mechanisms to ensure incentive-compatibility, i.e., ensuring that there are no misreported values or misreported preferred times associated with uncoordinated incentives. Another direction is to study the trade-offs between fairness and stability in a broader area. Although the notion of fairness has been studied in many different contexts, there is little work investigating both fairness and stability in complex systems. Ultimately, incorporating both fairness and stability in analytic models can achieve more socially beneficial solutions and practical results.

References

Javier Alonso-Mora, Samitha Samaranayake, Alex Wallar, Emilio Frazzoli, and Daniela Rus. On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences*, 114(3):462–467, 2017.

Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.

Xiaohui Bei and Shengyu Zhang. Algorithms for trip-vehicle assignment in ride-sharing. In *Proc. AAAI*, 2018.

Gerardo Berbeglia, Jean-François Cordeau, Irina Gribkovskaia, and Gilbert Laporte. Static pickup and delivery problems: a classification scheme and survey. *Top*, 15(1):1–31, 2007.

Dimitris Bertsimas and John N Tsitsiklis. *Introduction to linear optimization*, volume 6. Athena Scientific Belmont, MA, 1997.

Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. The price of fairness. *Operations research*, 59(1):17–31, 2011.

Dimitris Bertsimas, Patrick Jaillet, and Sébastien Martin. Online vehicle routing: The edge of optimization in large-scale applications. *Operations Research*, 2019.

Kostas Bimpikis, Ozan Candogan, and Daniela Saban. Spatial pricing in ride-sharing networks. 2016.

Filippo Bistaffa, Alessandro Farinelli, Georgios Chalkiadakis, and Sarvapali D Ramchurn. Recommending fair payments for large-scale social ridesharing. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pages 139–146. ACM, 2015.

Anna Bogomolnaia and Hervé Moulin. A new solution to the random assignment problem. *Journal of Economic theory*, 100(2):295–328, 2001.

Ilan Reuven Cohen and David Wajc. Randomized online matching in regular graphs. Working paper, 2018.

Jean-François Cordeau, Gilbert Laporte, Martin WP Savelsbergh, and Daniele Vigo. Vehicle routing. *Handbooks in operations research and management science*, 14:367–428, 2007.

Jean-François Cordeau. A branch-and-cut algorithm for the dial-a-ride problem. *Operations Research*, 54(3):573–586, 2006.

Masabumi Furuhata, Maged Dessouky, Fernando Ordóñez, Marc-Etienne Brunet, Xiaoqing Wang, and Sven Koenig. Ridesharing: The state-of-the-art and future directions. *Transportation Research Part B: Methodological*, 57:28–46, 2013.

Kazuo Iwama and Shuichi Miyazaki. A survey of the stable marriage problem and its variants. In *International Conference on Informatics Education and Research for Knowledge-Circulating Society*, pages 131–136. IEEE, 2008.

Yonghe Liu and Edward Knightly. Opportunistic fair scheduling over multiple wireless channels. In *IEEE INFOCOM 2003. Twenty-second Annual Joint Conference of the IEEE Computer and Communications Societies*, pages 1106–1115. IEEE, 2003.

David F Manlove, Robert W Irving, Kazuo Iwama, Shuichi Miyazaki, and Yasufumi Morita. Hard variants of stable marriage. *Theoretical Computer Science*, 276(1-2):261–279, 2002.

Neda Masoud and R Jayakrishnan. A real-time algorithm to solve the peer-to-peer ride-matching problem in a flexible ridesharing system. *Transportation Research Part B: Methodological*, 106:218–236, 2017.

Erhun Ozkan and Amy Ward. Dynamic matching for real-time ridesharing. 2017.

Sophie N Parragh, Karl F Doerner, and Richard F Hartl. Variable neighborhood search for the dial-a-ride problem. *Computers & Operations Research*, 37(6):1129–1138, 2010.

Paolo Santi, Giovanni Resta, Michael Szell, Stanislav Sobolevsky, Steven H Strogatz, and Carlo Ratti. Quantifying the benefits of vehicle pooling with shareability networks. *Proceedings of the National Academy of Sciences*, 111(37):13290–13294, 2014.

Xing Wang, Niels Agatz, and Alan Erera. Stable matching for dynamic ride-sharing systems. *Transportation Science*, 2017.