

Towards Ridesharing with Passenger Transfers

(Extended Abstract)

Brian Coltin
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA
bcoltin@cs.cmu.edu

Manuela Veloso
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA
veloso@cs.cmu.edu

ABSTRACT

Ridesharing services have the potential to fill empty seats in cars, reduce emissions and enable more efficient transportation. We propose rideshare services which *transfer* passengers between multiple drivers. By planning for transfers, we increase the availability and range of the rideshare service, and reduce the total vehicular miles travelled by the network. We propose three heuristics to schedule rideshare routes with transfers. Each provides a tradeoff in effectiveness and computational cost. We demonstrate these tradeoffs and the advantage of transfers in simulation.

Categories and Subject Descriptors

I.2.8 [Problem Solving, Control Methods, and Search]: Scheduling

General Terms

Algorithms

Keywords

rideshare, transfers, auctions

1. INTRODUCTION

Dependence on personal automobiles is becoming increasingly costly due to accelerating climate change and rising gasoline prices. It is particularly wasteful when one realizes that most car seats are typically empty. Ridesharing schemes counter this wastefulness by matching passengers to drivers to drivers through the use of smart phones. Drivers offer transport to passengers in exchange for sharing fuel costs. There are currently over 600 ridesharing services [2].

A ridesharing problem is defined by a set of vehicles, a set of passengers, and a map. Each vehicle and passenger has a starting and ending location. Vehicles have capacities and a maximum number of total passengers they are willing to pick up. The goal is to assign each vehicle a path that delivers all passengers and vehicles to their destination while obeying the constraints and minimizing the total distance traveled.

In this work, we reduce fuel use even further by *transferring* passengers between vehicles. By planning for transfers,

passengers can travel further at less inconvenience to drivers. This planning problem is especially challenging because the number of possible assignments of riders to vehicles increases exponentially when transfers are allowed.

Limited prior work has been done on the ridesharing problem without transfers: auctions [1, 6], set cover approximation algorithms [5] and genetic algorithms with insertion heuristics [4] have previously been used to assign passengers to vehicles. To our knowledge, the only previous researchers to consider transfers in ridesharing used an evolutionary algorithm to plan a route only for a single passenger [3].

2. RIDESHARING ALGORITHMS

We use two algorithms to find ridesharing solutions without transfers: a greedy algorithm, in which we insert passengers into existing vehicle routes, and an auction in which passengers bid on vehicles to transport them. We introduce three novel algorithms to plan for ridesharing with transfers. We assume an algorithm to choose a transfer point between two vehicles' paths is given.

2.1 Greedy Algorithm

The first approach we propose greedily adds transfer points to an existing solution without transfer points. We form a queue of passengers and the vehicles transporting them. For each passenger and vehicle in the queue, we iterate through the other vehicles and find those that could take the passenger part of the way at a lower cost. If such a vehicle exists, we add the transfer to the vehicle which decreases the cost the most. We then add both halves of the split route back into the queue for further recursive splitting.

2.2 Auctioning Passengers

Our second approach is an auction which again starts from a solution without transfers. In each auction round, each passenger finds the vehicle which a transfer to would reduce the cost the most and places a bid for that vehicle. Each vehicle accepts the bid which decreases the total cost the most. New rounds continue until no bids are placed.

Like the greedy approach, the auction algorithm only considers a single transfer at a time. However, the auction is less greedy in the sense that the assignment does not depend on the ordering the passengers or vehicles are examined in.

2.3 Graph Search

Our final algorithm plans for transfers from the beginning. The graph search algorithm is greedy in the sense that it iterates through every passenger and plans the best

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.

Copyright © 2013, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

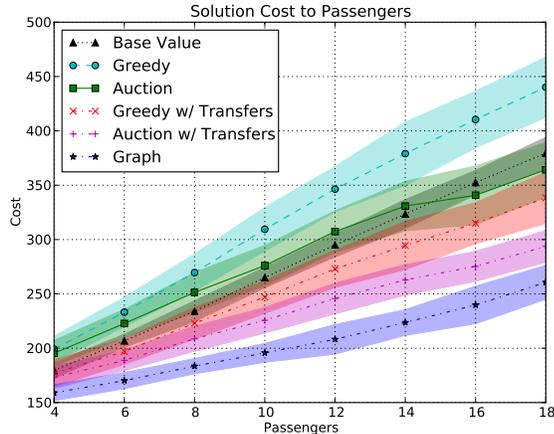


Figure 1: The solution cost found for each method in the Euclidean domain with $|V| = 20$, $C_v = 5$, $M_v = 7$, $B_v = 6$, and $c_T = 0$.

path for that passenger. To do so, we construct a directed multi-graph containing all the vehicles’ current routes, and the potential transfer points between edges on these routes. Graph edges represent segments of vehicles’ paths or transfers that the passenger could take on their routes. Edge weights represent the additional cost to vehicles of traversing that edge (i.e., edges already on vehicle routes have zero cost). The shortest path on the graph gives the route of least cost for the passenger.

Graph search is slower than the other algorithms, since the number of exchange points increases quadratically with the number of vehicle path edges. However, there is room for speed-ups with incremental graph construction.

3. SELECTED EXPERIMENTAL RESULTS

To verify the effectiveness of these three algorithms, we compared them in simulation. The experiments were performed on a Euclidean plane, with starting points and destinations chosen randomly such that passenger routes are of length at least 10 units, and vehicle routes are of length 5 to 7 units.

Figure 1 shows the costs of the solutions found by each of the algorithms in this scenario. The shaded regions denote the standard deviation across the fifty trials, and the “base cost” is what the cost in fuel would be if all of the passengers and drivers drove themselves in their own vehicles directly to their destinations. In this particular domain, the algorithms without transfers perform no better than the base cost, while with transfers we outperform the case where everyone drives themselves, reducing fuel usage. The greedy transfer algorithm, the auction transfer algorithm, and the graph-based algorithm each offer a successive improvement. With the graph-based algorithm and 18 passengers, transfers reduce the distance travelled by nearly 30%.

In Figure 2 we show the results of an experiment with the same parameters, except the number of vehicles changes rather than the number of passengers. Here, the approaches with transfers still significantly outperform the approaches without. The improvement, particularly with the graph algorithm, increases with the number of vehicles.

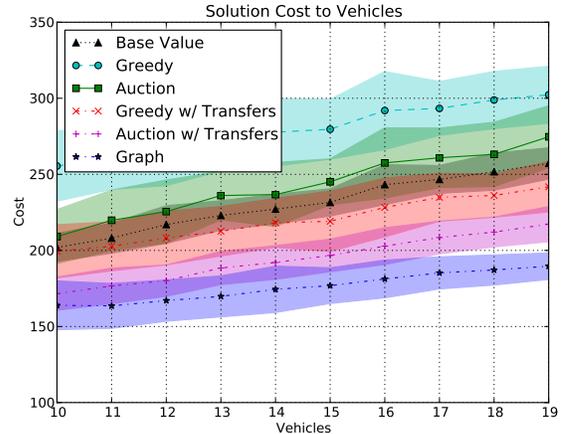


Figure 2: The solution cost found for each method in the Euclidean domain with $|P| = 10$, $C_v = 5$, $M_v = 7$, $B_v = 6$, and $c_T = 0$.

4. CONCLUSION

We have introduced the problem of ridesharing with transfers, and presented three algorithms to find solutions: a greedy approach, an auction approach, and an approach based on graph search. We have demonstrated that transferring passengers can reduce the distance travelled by nearly 30%. A wide range of future work remains, including evaluating these algorithms’ effectiveness in more realistic and *dynamic* settings, extending these algorithms to work in a distributed setting with only local information, and considering user time and convenience.

5. ACKNOWLEDGMENTS

This research was partially sponsored by the Office of Naval Research under grant number N00014-09-1-1031, and by the National Science Foundation under award number NSF IIS-1012733. The views and conclusions expressed are those of the authors only.

6. REFERENCES

- [1] S. Abdel-Naby, S. Fante, and P. Giorgini. Auctions negotiation for mobile rideshare service. In *Int. Conf. on Pervasive Computing and Applications*, 2007.
- [2] N. Chan and S. Shaheen. Ridesharing in north america: Past, present, and future. *Transport Reviews*, 2012.
- [3] W. Herbawi and M. Weber. Evolutionary multiobjective route planning in dynamic multi-hop ridesharing. In *Evolutionary Computation in Combinatorial Optimization*. 2011.
- [4] W. Herbawi and M. Weber. A genetic and insertion heuristic algorithm for solving the dynamic ridematching problem with time windows. In *Int. Conf. on Genetic and Evolutionary Computation*, 2012.
- [5] E. Kamar and E. Horvitz. Collaboration and shared plans in the open world: Studies of ridesharing. In *IJCAI*, 2009.
- [6] A. Kleiner, B. Nebel, and V. Ziparo. A mechanism for dynamic ride sharing based on parallel auctions. In *IJCAI*, 2011.