Promoting Peer-to-Peer Ridesharing Services as Transit System Feeders

Transportation Research Record, 2017

Neda Masoud (corresponding author)
Assistant Professor, Civil and Environmental Engineering
University of Michigan Ann Arbor, Ann Arbor, MI 48109
nmasoud@umich.edu

Daisik Nam
PhD Candidate, Civil and Environmental Engineering
University of California Irvine, Irvine, CA 92697
daisikn@uci.edu

Jiangbo Yu
PhD Candidate, Civil and Environmental Engineering
University of California Irvine, Irvine, CA 92697
jiangby@uci.edu

R. Jayakrishnan
Professor, Civil and Environmental Engineering
University of California Irvine, Irvine, CA 92697-3600
rjayakri@uci.edu

Abstract

Peer-to-peer ridesharing is a recently emerging travel alternative that can help accommodate the growth in urban travel demand and at the same time alleviate problems such as excessive vehicular emissions. Prior ridesharing projects suggest that the demand for ridesharing is usually shifted from transit, while its true benefits are realized when the demand shifts from single-occupancy vehicles. This paper studies the potential of shifting demand from private autos to transit by providing a general modeling framework that finds routs for private vehicle users that are a combination of peer-to-peer (P2P) ridesharing and transit.

The Los Angeles Metro red line is considered for a case study, since it has recently shown declining ridership trends. For successful implementation of a ridesharing system, strategically selecting locations for individuals to get on/off rideshare vehicles is crucial, along with an appropriate pricing structure for rides. A parametric study of the application of real-time P2P ridesharing to feed the LA Metro red line using simulated demand is conducted. A mobile application with an innovative ride-matching algorithm is developed as a decision support tool that suggests transit-rideshare and rideshare routes.

INTRODUCTION

One of the main issues faced by major cities in the US today is congestion. In addition to directly impacting travelers by increasing travel time and reducing travel time reliability, congestion leads to higher levels of Green House Gas (GHG) emissions which are damaging to people's health and the environment. One of the solutions as to how to reduce congestion is to eliminate vehicles from roads by putting individuals who are traveling along the same routes in the same vehicles.

Although carpooling has been around since the invention of cars, the nature of carpooling, which requires making arrangements in advance of the travel and committing to those arrangements for a period of time, makes it unattractive to many individuals. Dynamic ridesharing is a modernized form of carpooling that is on-demand, requires only a one-time commitment, and does not require making arrangements far in advance of the trip.

In recent years, ridesharing (including carpooling) in the US has experienced a slight increase in mode share, reaching a mode share of 11% in 2008 (*1*). Although this increase in ridesharing demand seems to be a step forward in the direction of a greener transportation system, this is not necessarily the case. The modal shift due to the introduction of ridesharing is as important. The benefits of ridesharing depend tremendously on this model shift. The benefits would be high in terms of reducing congestion and GHG emissions if the demand is shifted from single-occupancy vehicles to rideshare systems, but may not be significant if the ridesharing demand is being shifted from transit. In addition, introduction of ridesharing can lead to emergence of more complex multi-modal alternatives, such as the transit-rideshare mode.

Study of government-funded ridesharing systems indicates that ridesharing systems, as they work today, are competitive to transit systems (*2, 3*). The goal of this study is to assess the potential of ridesharing in being a complement to transit, feeding it, instead of shifting demand away from it. We analyze the potential of such multi-modal travel using a parametric study with simulated demand based on the current Southern California Assosiation of Governments (SCAG) model of a selected area (Los Angeles Metro red Line catchment area), and develop an app using an

advanced ride-matching algorithm. As a side benefit, this study also sheds light on the potential of ridesharing as a stand-alone mode to replace transit.

The reason why we consider the LA Metro red line (Figure 1(a)) to study potential rideshare-based demand augmentation is the noticeable reduction of ridership (shown in Figure 1(b)) in recent years. This indicates potential opportunities for additional demand-inducement strategies using rideshare options.
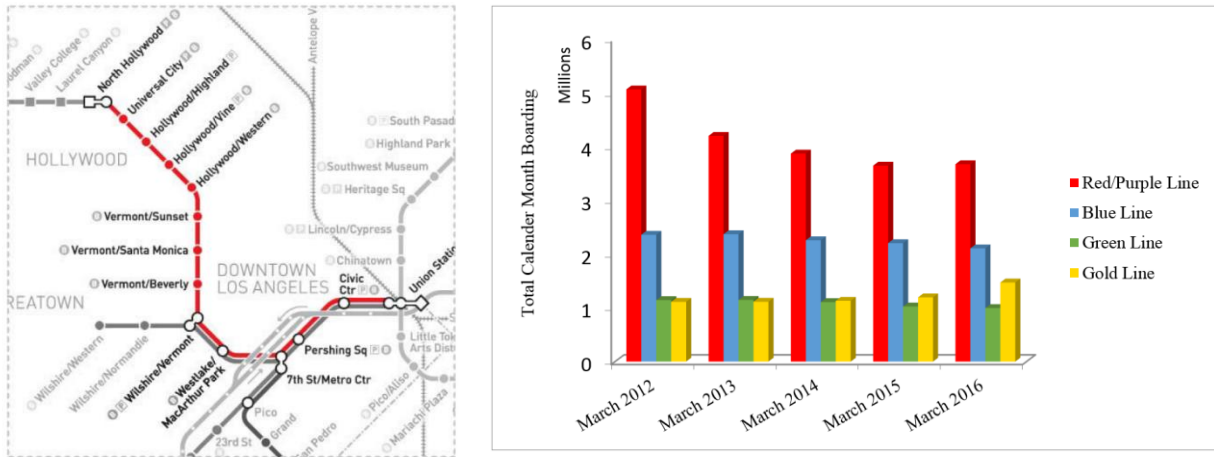
(a) Los Angeles Metro Red line   (b) Ridership trends in LA metro

FIGURE 1 Case study area: Los Angeles metro red line

The success of a multi-modal transit-rideshare system can be considerably influenced by the architecture of the designed system, namely locations where the ridesharing service is offered, price of ridesharing, and the matching method used by the system. In this study, we elaborate on this system architecture, and showcase the impact of such targeted architecture on transit ridership augmentation for the LA Metro red line.

LITERATURE REVIEW

Mobility options in urban areas have been traditionally duopolized by the private and public sectors. It has only been in recent years that public-private partnerships have started receiving attention. These new partnerships have emerged in light of the fact that there exist a wide spectrum of previously-neglected transportation options in between the fixed, inexpensive, and high-capacity public transportation option and the flexible, expensive, and low-capacity private sector option. Designing services that fall somewhere in the middle ranges of this spectrum calls for public-private partnerships, and is the key to offering flexible and yet relatively inexpensive transportation alternatives.

An example of this public-private partnership is Massachusetts Bay Transportation Authority's (MBTA's) use of taxis to connect individuals with disabilities to the transit lines in a pilot program (*4*). Another example is Car2Work, an innovative shared mobility concept that connects commuters to transit, introduced by Regue et al. (*5*). This model is motivated by Zev.Net, a limited shared-mobility deployment project launched collaboratively by University of California Irvine and Toyota Motor Sales in 2002.

Ride-hailing companies have also started building relationships with the public sector in recent years, and are expected to engage in more mutually beneficial arrangements with transit agencies in the future. Uber is one of the ride-hailing companies who has started making partnerships with public transportation agencies throughout the US in an attempt to fill the gap in their transit networks. Among other partnerships, Uber is currently working to address the first-mile, last-mile problem faced by DART riders in Dallas and MARTA riders in Atlanta (*6*). Lyft is another ride-hailing company who has recently announced a partnership with San Francisco Bay area's MTC (*7*).

Connecting individuals to transit using other modes of transportation calls for multi-modal matching algorithms. In a multi-modal (or multi-hop) matching algorithm, a passenger can travel from their origin to their destination by means of transferring between different vehicles/modes of transportation.

The first official formulation of a multi-modal passenger pick-up and delivery problem was offered by (*8*). Multi-modal routing has been incorporated in dial-a-ride problems as well, albeit in very limited settings. Studies conducted by (*9*), (*10*), and (*11*) showcase how devising multi-modal routes can enhance system performance in transportation systems that deal with both people and goods. In the context of P2P ridesharing, (*12*), (*13*), (*14*), and (*15*) have proposed optimization formulations/algorithms to solve the multi-hop ride-matching problem.

All the mentioned multi-hop matching algorithms can be used to connect individuals to transit, using other modes of transportation that could include ridesharing. The aforementioned algorithms, however, are not designed for real-time use in dynamic large-scale systems. For the purpose of this study, we use the Dynamic Programming (DP) algorithm proposed by Masoud and Jayakrishnan in (*16*) to solve the multi-hop matching problem, since this algorithm demonstrates promise in terms of scalability, and offers optimal solutions.
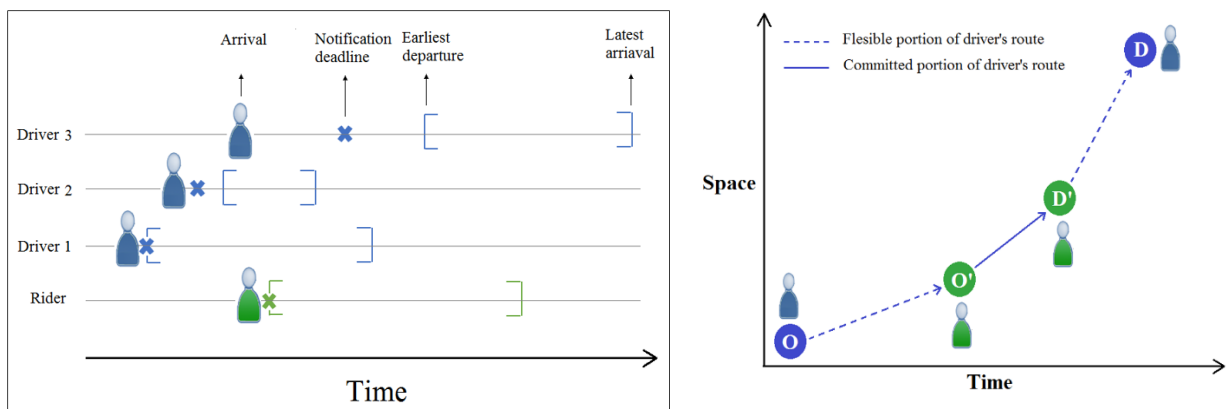
RIDESHARING SYSTEM

A ridesharing system comprises of a set of riders who are in need of transportation, and a set of drivers who are willing to use the empty seats in their vehicles to carry passengers in exchange for a monetary compensation. Let us define a set of pre-specified locations in the network, called "go-points", denoted by S_G, where individuals can start and end their trips. In addition, we define a set of "transfer points'", denoted by S_T, where riders can transfer between vehicles. Go- and transfer points, along with transit stations, form the set of "stations" in our system.

Upon registering in the system, a participant (rider or driver) is asked to provide information on their origin and destination go-points, a travel time window bounded from below by their earliest departure time from their origin go-point and from above by their latest arrival time at their destination go-point, and a notification deadline by which they need to be informed whether they have been matched or not. Figure 2 demonstrates a ridesharing instance. Drivers are asked to provide the capacity of their vehicles, and riders have the option to specify the maximum number of transfers they are willing to make.

In this study, we consider an on-demand ridesharing system in which a rider's arrival time at the system and their notification deadline coincide. Once a rider registers a request, the ridesharing system solves a matching problem that includes the rider, as well as all the drivers whose travel time windows intersect with the travel time window of the rider. In the ridesharing instance in

Figure 2, for example, all drivers are eligible to be included in the matching problem solved for the rider.

(a) An example of a ridesharing instance   (b) Partially fixed driver route

FIGURE 2 Ridesharing system

If a driver is matched with a rider, the part of the driver's route that is committed to the rider becomes fixed. Other parts of the driver's route, however, are flexible and can be optimized for the subsequent riders who join the system later. Figure 2(b) displays an example of such an instance. The driver in this figure (shown in blue) is traveling from origin O to destination D. This driver has been matched with a rider, but his/her route from O to O', and from the D' to D can still be optimized to put the driver in spatiotemporal proximity of other riders that request service at some later point in time. In other words, this previously matched driver now enters the ride-matching problem as three separate drivers, one that travels from O to O', one that travels from O' to D', and one that travels from D' to D. The travel time window for each of these three drivers is determined based on the committed portion of the driver's route, and his/her original travel time window. Note that the capacity of the vehicle going from O' to D' needs to be adjusted.

After determining drivers who are eligible to be included in a rider's matching problem, the system uses the multi-hop matching algorithm in (16) to solve the matching problem, and announces to the rider within a matter of seconds whether the rider has been matched, and the itinerary of the rider in the case of a successful match.

THE MULTI-HOP MATCHING ALGORITHM

As mentioned previously, in this study we use the Dynamic Programming (DP) algorithm proposed in (16). This algorithm finds a (multi-hop) path for a rider by optimally routing drivers. Transit lines can enter the model as inflexible drivers (since their routes and schedules are fixed and cannot be optimized).

The DP algorithm runs on a time-expanded network. By introducing stations, we are in fact discretizing the two-dimensional continuous-space network into a set of discrete locations. In addition, we discretize the study time horizon into a set of time intervals. We use 5-minute time intervals in this study. In a network discretized in time and space, we denote a node $n_i = (t_i, s_i)$ as a tuple of time interval and station, and a link as a tuple (t_i, s_i, t_j, s_j). Such a link can be

interpreted as a trip that starts at go-point s_i at time interval t_i, and ends at go-point s_j at time interval t_j.

For each rider, the algorithm identifies the set of links that can be traversed both by the rider and at least one driver/transit line given the time constraints of the rider and drivers/transit lines. Figure 3 shows an example of a time-expanded network for a rider. Nodes in this figure are shown as rectangles, and links as arcs connecting the nodes. In this example, the rider is starting his/her trip from go-point 5, and is traveling to go-point 12. The rider's earliest departure time from the origin go-point is the 3rd time interval, and his/her latest arrival time at the destination go-point is the 39th time interval. Drivers/transit lines that can potentially carry the rider on each link (i.e., leg of the trip) are denoted next to the link. The goal of the DP algorithm is to find the min cost path for the rider that connects origin O to destination D.
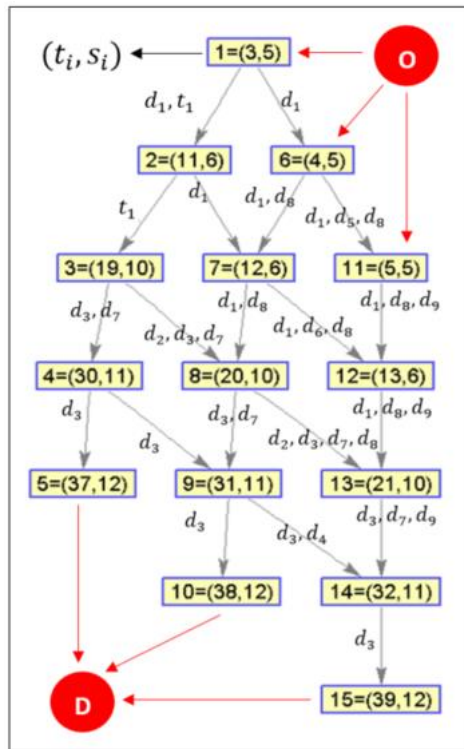

FIGURE 3 Time-expanded network

The cost function we use in the DP algorithm is the sum of four components: (*i*) a distance-based fare, (*ii*) dollar value of travel time, (*iii*) dollar value of additional penalty for waiting time, and (*iv*) dollar value of penalty for transfers. We consider the default value of $20/hr for value of time (VOT), $0.25/mile for distance-based fare, $1.5 fare for use of transit, $0.1 penalty for waiting for one time interval (in addition to the value of time), and $0.1 for penalty for each transfer. Admittedly, there could be other factors that contribute to the utility function of a rider that are not considered in this paper, including the cost of parking and the perceived discomfort and lack of reliability associated with not driving one's private vehicle. These factors could increase or decrease the matching rate from the numbers reported in this paper, depending on their relative importance to individuals.

CASE STUDY IN THE Los Angeles COUNTY

The success of a multi-modal transit-rideshare system can be considerably influenced by the architecture of the designed system, namely locations where the ridesharing service is offered, cost of ridesharing to consumers, and the matching method used by the system. In this section, we elaborate on this system architecture, and showcase the impact of such a targeted architecture on transit ridership augmentation for the LA Metro red line.

Stations

As mentioned previously, the station set is the union of go-points, transfer points, and transit stations. Due to the large size of the network, we only consider a subset of go-points as transfer points. We assume that each go-point belongs to a neighboring transfer point. Individuals can travel directly between the go-points attributed to the same transfer point. However, in order to travel between two go-points attributed to two different transfer points, an individual has to travel between the two transfer points. Note that despite traveling between transfer points, a transfer does not have to occur, i.e., an individual can enter and exit a transfer point in the same vehicle.

The SCAG region has a total of over 4000 TAZs that can be potential go-points. In order to identify go-points in our model, we studied the SCAG trip tables and identified OD pairs with significant demand levels (OD pairs with hourly trip rates higher than 10). This procedure resulted in a total of around 1000 go-points. We limited our analysis to single-occupancy auto demand only, because the focus of our study is to design a system that encourages modal shift from the drive-alone mode to rideshare and rideshare-transit alternatives.

The set of transfer points is a subset of the go-points set. We identified this subset using Algorithm 1 to ensure that transfer points (*i*) are located closer to go-points with higher levels of demand, and (*ii*) are distributed in the network as evenly as possible.

Algorithm 1 starts by initializing the set of transfer points, S_t, by the set of red line stations. Let us define the set of potential transfer points as S'_G, and initialize this set with the set of go-points, i.e., all go-points are potential transfer points. Next, in an iterative process, 40 additional transfer points were selected from the set of go-points and added to the set of transfer points. The number 40 was selected after multiple experimentations and arriving at the conclusion that this number provided a nice balance between the number of transfer points (too many transfer points increase the complexity of the problem), and the degree of network coverage (too few transfer points can cause insufficient network coverage) in our study area.

In each iteration, the algorithm selects a single go-point with the lowest total person-miles required for travel from the rest of the go-points in set S'_G. The selected go-point requires the least effort from individuals to arrive at it, in order to make a transfer, and therefore is set as a transfer point. In order to ensure that transfer points are not too close to each other and are well-spread throughout the network, the algorithm then eliminates from set S'_G the go-points within a 4 miles radius of this recently selected transfer point. The algorithm then moves forward to the next iteration, where the next transfer point is selected. The algorithm stops when set S'_G becomes empty, or the limit of 40 transfer points (besides the red line stations) is reached. Figure 4 displays the three sets of stations in the LA network.
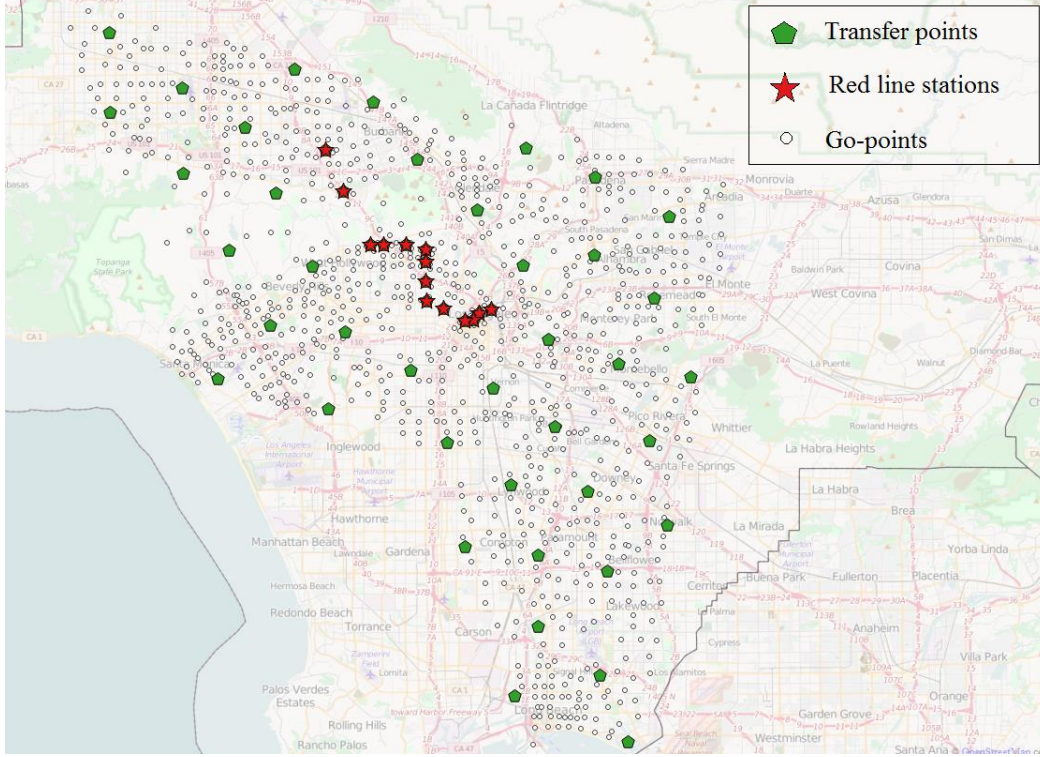
FIGURE 4 Three types of transfers in the LA network

ALGORITHM 1 Identifying transfer stations

$S_T = S_R$
$S'_G = S_G \setminus S_R$
For $s \in S_R$
    $del = \{k \in S_G : d_{s,k} \leq 2 \text{ miles}\}$
    $S_G = S_G \setminus del$
End For
Set $Done \leftarrow 0$
For $t = 40$
    While $Done = 0$
    For $s \in S'_G$
    $NS = \text{argmin}_{i \in S'_G} \{w^{*1}_{i,s}, d^{*2}_{i,s}\}$
    $S_T = S_T \cup NS$
    $del = \{k \in S'_G : d_{NS,k} \leq 4 \text{ miles}\}$
$S'_G = S'_G \setminus del$
    If $S'_G = \emptyset$
        $Done \leftarrow 1$
    End If
    End For
End While
End For
$* 1$ $w_{i,s}$: travel demand from station i to station s
$* 2$ $d_{i,s}$: travel distance from station i to station s

Link Sets

We introduce three families of link sets that connect different types of stations (i.e., go-points, transfer points, and red line stations). The first link set, displayed in Figure 5(a), connects transfer points to each other. The second link set, displayed in Figure 5(b), connects go-points to their corresponding transfer points. Figure 5(c) demonstrates the third link set, which connects the red line stations to their nearby go-points. This link set connects each of the go-points confined within a 2.5 mile radius of at least one of the red line stations to all the red line stations located within their 2.5 miles radius. Notice that we do not generate any ridesharing links that connect red line stations to each other, or to their nearby go-points, in order to have the itineraries use transit whenever possible, and achieve our goal of using ridesharing to feed transit, and not to replace it.

(a) Link set 1: Links connecting transfer points

(b) Link set 1: Links connecting go-points to their corresponding transfer points

(c) Links set 3: Links connecting red line stations to their neighboring go-points
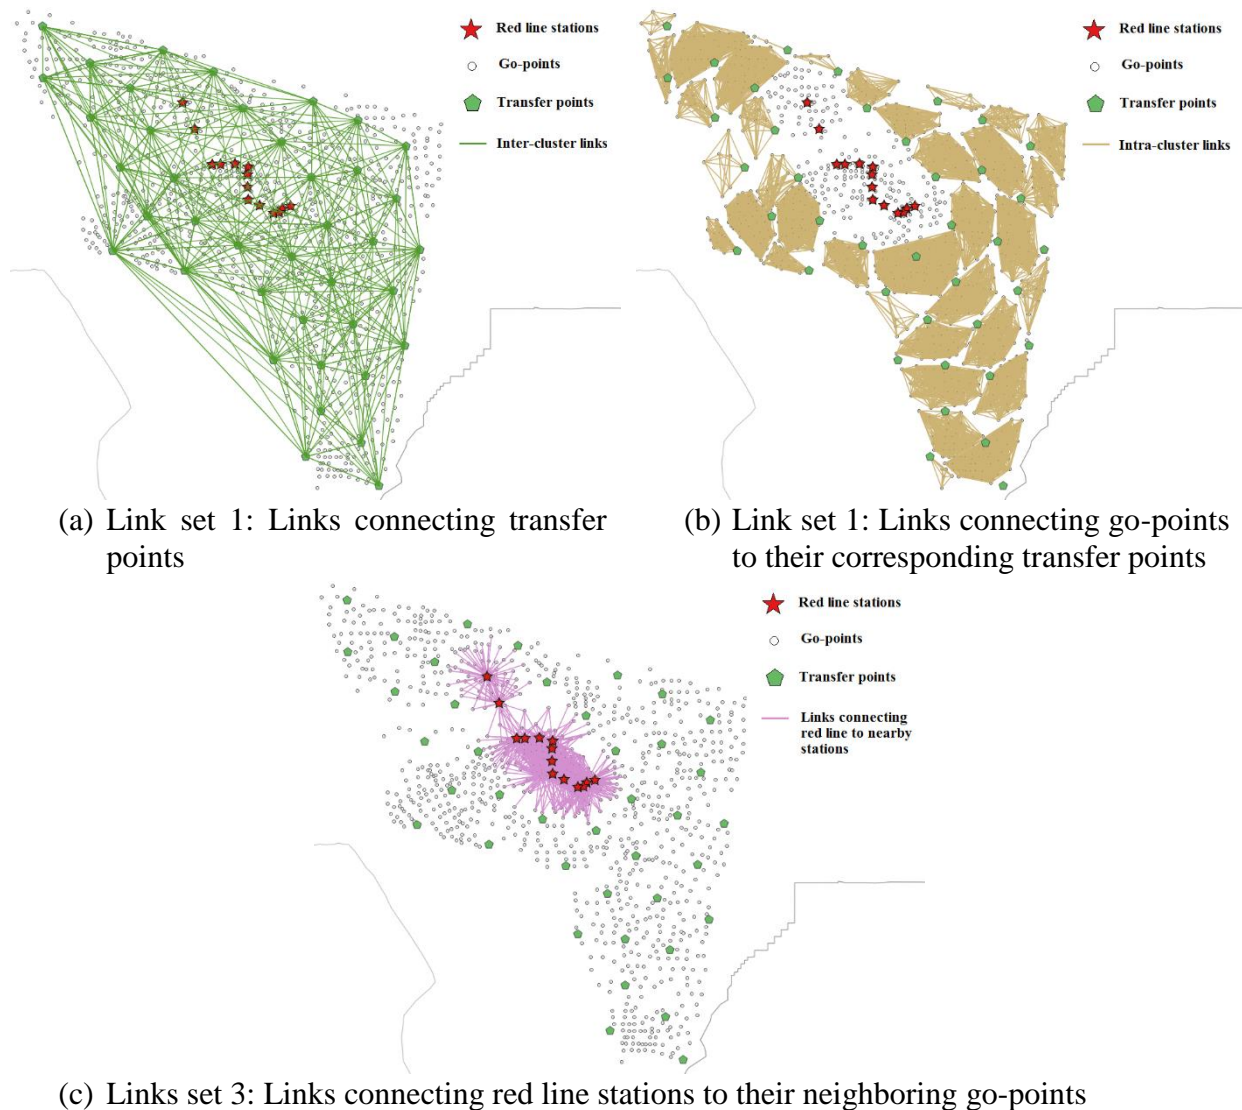
FIGURE 5 Link sets

Each go-point in the network is connected to at least one transfer point (link sets 2 and 3). In addition, all transfer points are connected to each other (link set 1). This indicates that there is a path between any two go-points in the network. Link set 2 indicates that the shortest path of a rider traveling between two go-points (that are not within a 2.5 miles radius of the red line) includes traveling to the transfer point corresponding to the origin go-point, from there to the transfer point corresponding to the destination go-point, and finally to the destination go-point itself. Furthermore, since all the go-points corresponding to the same transfer point are connected to each other, if a rider needs to make a short trip between two such go-points, no transfers are required. For practical reasons, it is assumed that transfers for trips that originate from or are destined to go-points within a 2.5 mile radius of the red line stations are limited to the metro red line stations only. Link set 3 connects such go-points to the red line stations directly.

Results

In this section, we study the modal shift from drive-alone to rideshare and rideshare-transit alternatives using simulations, and report on some of the features of the system. Simulations are conducted for the morning peak hour in LA. Origin-destination trip tables used in these simulations are obtained from the SCAG planning model, and spread throughout the three-hour morning peak period based on a uniform distribution.

For each simulation run, we randomly select our set of riders and drivers. In all simulation runs, we use 1,000 riders, but change the number of drivers from 1,000 to 160,000 in order to study the impact of the rider to driver ratio on the matching rate. We assume each vehicle has the capacity to carry 4 passengers, and do not a set a limit on the number of transfers. In addition, we assume that maximum ride times for participants are 20% higher than their shortest path travel times.

Matching Rate

In order to study the impact of the number of drivers on the matching rate, we ran a set of simulations, starting with the same number of riders and drivers (1000), and increasing the number of drivers to perform sensitivity analysis on increasing the supply level on the matching rate. The resulting matching rates are displayed in Figure 6. It is not an unrealistic assumption to presume that at least initially individuals might not be looking forward to leaving their vehicles behind and traveling with others as riders. However, participating in a ridesharing system as a driver is much more convenient, as drivers just need to register their trips (and may set a default time window so that they only need to login into the system when their regular daily schedules change). Hence, in these simulations, we study the impact of having more drivers participate in the system.

Figure 6(a) displays the percentage of served riders as a function of number of drivers. As intuition suggests, this percentage increases with the number of drivers. Percentage of served riders, however, grows with the number of drivers at a rate slower than linear. For example, with a 5,000 unit increase in the number of drivers (going from 5,000 to 10,000 drivers), we witness a 20% increase in the percentage of served riders. However, in order to experience another 20% increase in the percentage of served riders, we have to have a 10,000 unit increase in the number drivers (going from 10,000 to 20,000 drivers).

Figure 6(b) displays the number of served riders and matched drivers as a function of the number of drivers in the system. This figure sheds light on the performance of the system under different levels of supply (i.e., number of drivers). When the number of drivers is at its lowest, the number of served riders is about the same as the number of matched drivers, implying that most trips are

being served without transfers. (This conclusion is confirmed by looking at the number of transfers for each level of supply in Figure 7, as we will discuss in the following section.) At low levels of supply, the number of drivers is too small for multi-hop routes to be formed for riders. Up to a certain level (20,000 drivers), the difference between the number of matched riders and drivers (i.e., the horizontal distance between the two curves in Figure 6(b)) keeps increasing. Finally, when the supply level becomes large, and most of the demand is being served, there is no need for more costly multi-hop routes anymore, and the number of matched riders and drivers start to converge again.
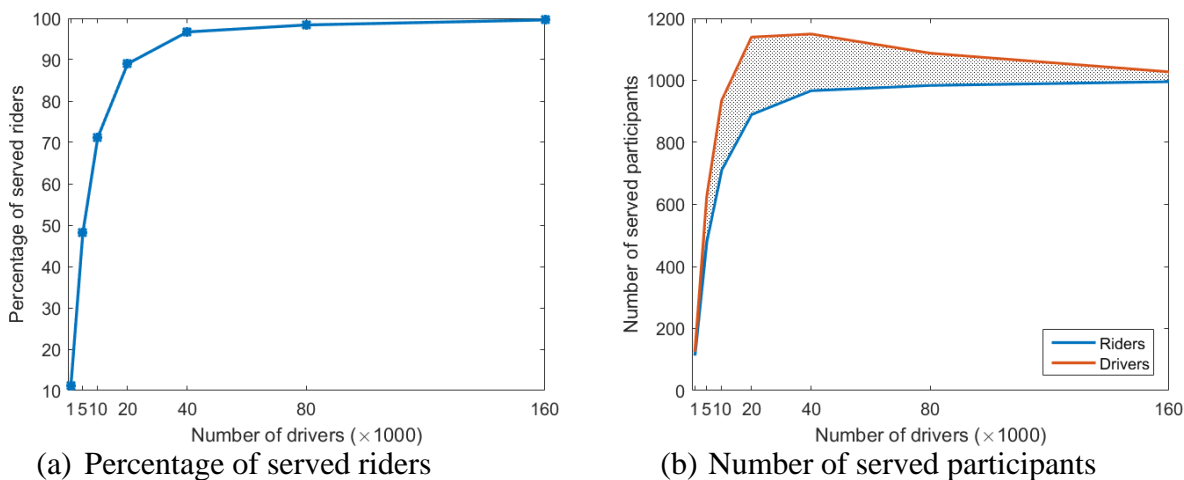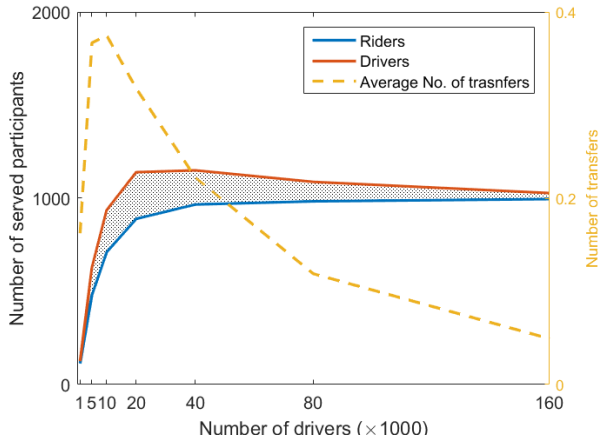

(a) Percentage of served riders   (b) Number of served participants

FIGURE 6 Matching rate as a function of the number of drivers
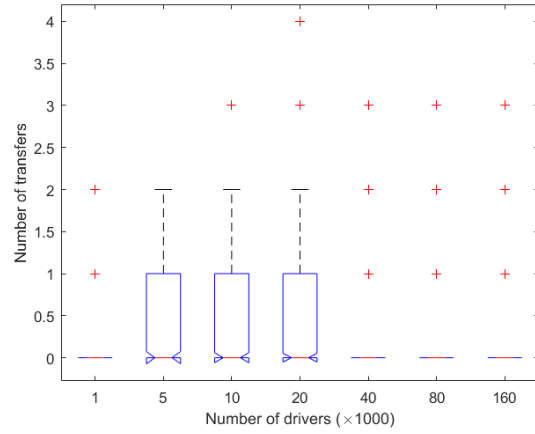
Number of transfers

Figure 7(a) demonstrates the number of transfers under different levels of supply. The number of transfers reported includes transfers between rideshare vehicles, or between rideshare and transit vehicles. This figure suggests that when the number of drivers is too low or too high, transfers are very limited, and most riders can be served with zero transfers (Figure7(b)). However, more transfers are required in the middle ranges.

As discussed in the previous section, at very low supply levels, there are not enough drivers in the system to form multi-hop routes, and at very high supply levels almost all ride requests can be served without any transfers, and therefore an overwhelming number of trips end up being single-hop. In the middle ranges, however, transfers are necessary to obtain higher matching rates. Figure 7 suggests that even in the middle ranges most riders experience zero transfers, with a few percentage experiencing 1 transfer. The maximum number of transfers ever witnessed was 4.
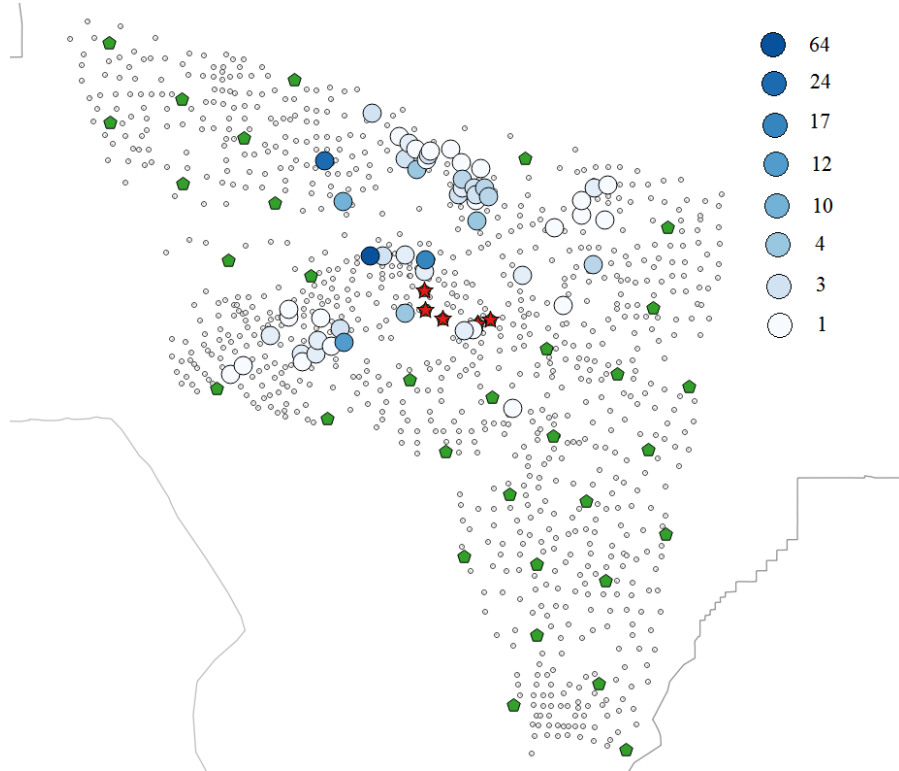
Figure 7(c) displays the most frequently used transfer points. This figure has been generated based on the simulation results for a ridesharing system with 20,000 drivers, since such a system was shown to have the highest number of transfers. This figure suggests that the most important transfer points coincide with some of the red line stations, which implies good decision making in determining the station locations by LA Metro. This figure can also be a guide in revising the transfer points in our models in the future studies.

(a) Relationship between the average number of transfers, and the number of served riders and drivers

(b) Boxplot of number of transfers

(c) Most frequently-used transfer points

FIGURE 7 Transfers

Vehicle Occupancy

Figure 8(a) shows the average vehicle occupancy as a function of the supply level. This figure suggests that the average occupancy of vehicles decreases as the number of drivers in the system increases, which is intuitive, since at lower levels of supply riders are more probable to share the limited resources. The maximum vehicle occupancy, however, follows the previously observed

trend of initially experiencing a rise, followed by a decline. Notice that the minimum vehicle occupancy is always higher than 2, since each matched driver carries at least one rider.
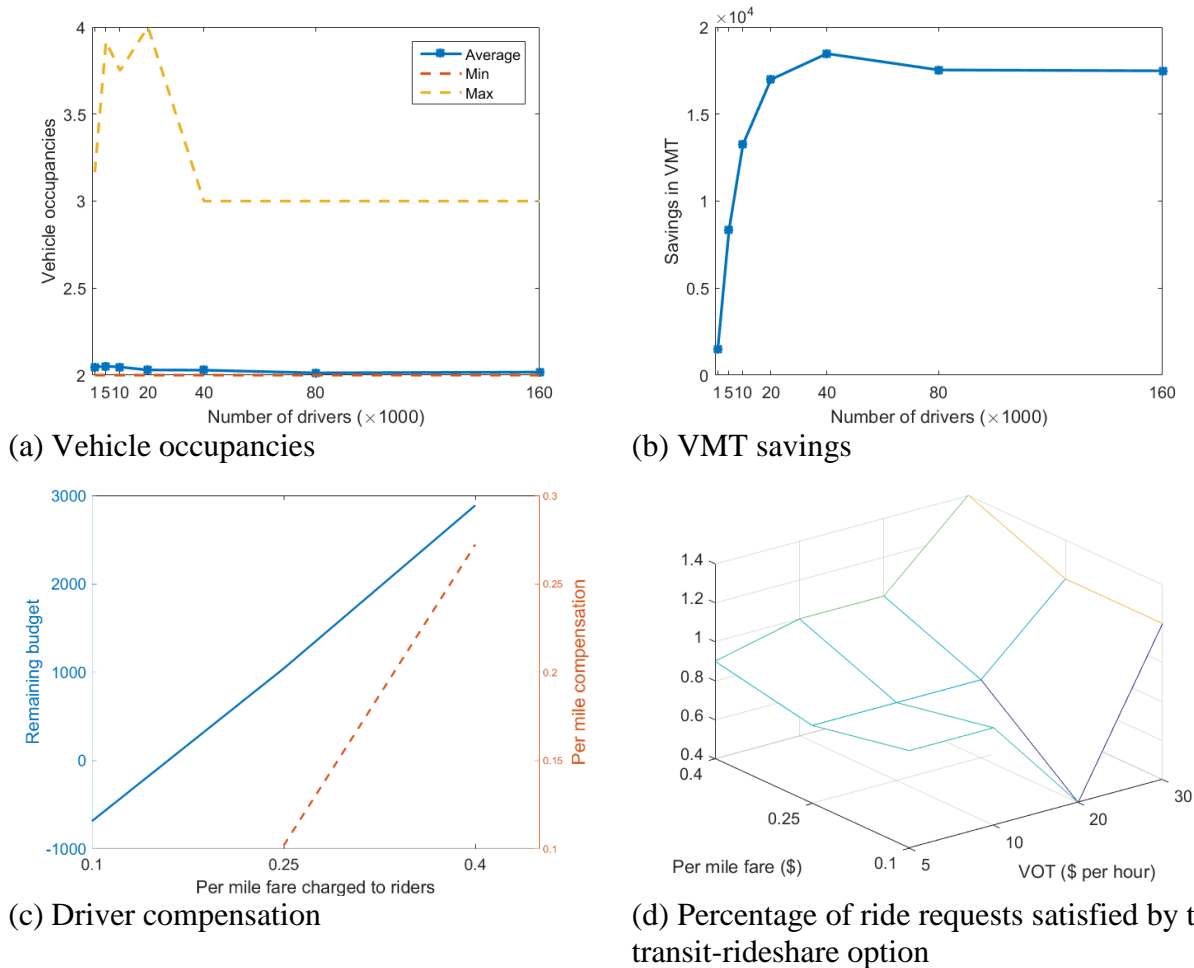

(a) Vehicle occupancies

(b) VMT savings

(c) Driver compensation

(d) Percentage of ride requests satisfied by the transit-rideshare option

FIGURE 8 Results of LA experiments

VMT Savings

Figure 8(b) displays the system-wide savings in vehicles miles traveled in the system under different levels of supply. Under lower levels of supply, where the matching rate is small, so are the VMT savings. VMT savings experience significant increase as the matching rate increases. However, at higher levels of supply, since drivers are abundant, vehicle occupancies drop (Figure 8(a)), and consequently VMT savings experience a slight decline. In general, the analyses conducted in this section suggest that with number of drivers falling in the range between 5,000 and 20,000, the system performance reaches its peak in terms of making the best use of resources.

Budget Balancedness

In this section, we look into the circumstances under which the ridesharing system can be budget balanced, and operate without need for outside subsidies. We assume that the system charges a distance-based fare to riders. From this total income, the system pays 56 cents to drivers for every additional mile they have to travel (compared to their shortest paths), and then distributes the rest

of the income to drivers on a per-mile basis for the duration of their trips where they have been transporting passengers.

Figure 8(c) shows the remaining budget after paying drivers for their additional idle travels. This figure displays the results of three simulation runs, with 10,000 drivers, and varying per-mile fares charged to riders. Results show that when riders are charged very little (10 cents per mile), the system will encounter a deficit. Per-mile fares of 25 and 40 cents, however, lead to some positive income, which can then be distributed to drivers based on their contribution to the system. Figure 8(c) indicates that when riders are charged 25 cents per mile, drivers can get paid 10 cents per mile, and when riders are charged 45 cents per mile, drivers can get paid 27 per mile. Under such circumstances, the system is budget balanced and does not require subsidies.

The Rideshare-Transit Alternative

Figure 8(d) shows the percentage of riders who use the transit-rideshare option, under different values for value of time (VOT), and distance-based fares. This figure shows that, as intuition suggests, use of transit increases with the distance-based fare (under all VOT values), since the transit fare is fixed. This figure also suggests that use of transit increases with VOT, which is expected, since the LA metro red line has a higher speed compared to the neighboring street network.

Although at first glance it might seem like the percentage of individuals using ridesharing as a means to connect to transit is not significant, it should be noted that the trip tables used for simulations are single-occupancy trip tables, and the individuals using the transit-rideshare option are actually increasing the current level of transit ridership. Furthermore, keep in mind that having as little as 1% of single-occupancy vehicles switch to transit would translate to a considerable increase in transit ridership. According to SCAG trip tables, about 1,000,000 single occupancy vehicles travel during the morning peak hours in LA County and within our study area in a working day. This adds up to a total of about 10,000 additional trips, just by the Metro red line, assuming that 1% of these single-occupancy trips will switch to the transit-rideshare alternative. This number of trips distributed evenly between the roughly 50 Metro lines running during the morning peak hours adds 200 passengers to each train.

Conclusion

In this paper, we proposed a system architecture that included strategically selecting a set of go-points for ridesharing, and a pricing scheme, to promote ridesharing as a solution for the first-mile, last-mile problem faced by the LA Metro red line. We showed that there exists a range of distance-based ridesharing fares for which people would prefer ridesharing to driving alone. Although the focus of the study and the drawn conclusions were centered on the role of ridesharing as a complement to transit, by selecting a study area that spreads way beyond the LA Metro red line, we were able to peak at the potential of ridesharing as a stand-alone mode as well.

The Los Angeles case study also provides insights on the pricing of ridesharing systems, and demonstrates how appropriate pricing levels could vary from region to region depending on the residents' financial welfare and their distribution of value of time. Furthermore, this study investigates under what levels of fares a ridesharing system can be budget-balanced, and operate without need for subsidies.

This study provides insights on the performance of the system under different levels of supply. Although it might seem counter-intuitive at first glance, for a fixed number of trip requests, there is an optimal range for the number of drivers that leads to the highest performance levels. Going beyond this range does not improve the matching rate, but deteriorates sustainability-related indices such as total vehicles miles traveled.

 In conclusion, this study demonstrates that if implemented correctly, a ridesharing system can in fact act as a feeder to public transportation, and hence bring about a greener transportation system and a less congested network, especially in urban settings.

Acknowledgements

References

[1] Santos, A., N. McGuckin, H. Y. Nakamoto, D. Gray, and S. Liss, Summary of travel trends: 2009 national household travel survey, 2011.

[2] Levofsky, A. and A. Greenberg, Organized dynamic ride sharing: The potential environmental benefits and the opportunity for advancing the concept. In Transportation Research Board 2001 Annual Meeting, 2001, pp. 7–11.

[3] Nelson Nygaard Consulting Associates and RideNow Inc, RideNow Evaluation Draft Final Report. In Prepared for the Alameda County Congestion Management Agency, 2006.

[4] Boston Globe, MBTA pilot taxi partnership could include Uber. Online: https://www.bostonglobe.com/metro/2015/11/14/mbta-launches-pilot-taxi-partnership-that-could-include-uber/6gTeEe8aJm5e6HEv9sdqNK/story.html, 2015, accessed: 2016-07-29.

[5] Regue, R., N. Masoud, and W. Recker, Car2work: A Shared Mobility Concept to Connect Commuters with Workplaces. Transportation Reserch Record, 2016.

[6] CITYLAB, Uber and Public Transit Are Trying to Get Along. Internet: http://www.citylab.com/cityfixer/2015/08/uber-and-public-transit-are-trying-to-get-along/400283/, 2015, accessed: 2016-07-29.

[7] Lyft Blog, Lyft Partners with California's MTC to Deliver New Carpooling Mode. Internet: https://blog.lyft.com/posts/lyft-mtc-511-carpooling, 2016, accessed: 2016-07-29.

[8] Cortés, C. E., M. Matamala, and C. Contardo, The pickup and delivery problem with transfers: Formulation and a branch-and-cut solution method. European Journal of Operational Research, Vol. 200, No. 3, 2010, pp. 711–724.

[9] Masson, R., F. Lehuédé, and O. Péton, The Dial-A-Ride Problem with Transfers. Computers & Operations Research, Vol. 41, No. 0, 2014a, pp. 12 – 23.

[10] Stein, D. M., Scheduling Dial-a-Ride Transportation Systems. Transportation Science, Vol. 12, No. 3, 1978, pp. 232–249.

[11] Liaw, C.-F., C. C. White, and J. Bander, A decision support system for the bimodal dial-a-ride problem. Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on, Vol. 26, No. 5, 1996, pp. 552–565.

[12] Herbawi, W. and M. Weber, Ant colony vs. genetic multiobjective route planning in dynamic multi-hop ridesharing. in Tools with Artificial Intelligence (ICTAI), 2011 IEEE Conference, 2011.

[13] Agatz, N., A. Erera, M. Savelsbergh, and X. Wang, Sustainable Passenger Transportation: Dynamic Ride-Sharing. ERIM Report Series Reference No. ERS-2010-010-LIS, 2009.N.

[14] Ghoseiri, K., DYNAMIC RIDESHARE OPTIMIZED MATCHING PROBLEM. Ph.D. Dissertation at University of Maryland, 2013.

[15] Masoud, N., and R. Jayakrishnan. A Decomposition Algorithm to Solve the Multi-Hop Peer-to-Peer Ride-Matching Problem. Transportation Research Part B: Methodological. V. 99, pp. 1-29, 2017.

[16] Masoud, N. and R. Jayakrishnan, A Real-Time Algorithm to Solve the Peer-to-Peer Ride-Matching Problem in a Flexible Ridesharing System, 2015.