USING BILATERAL TRADING TO INCREASE RIDERSHIP AND USER PERMANENCE IN RIDESHARING SYSTEMS

Neda Masoud
Ph.D. Candidate
Department of Civil and Environmental Engineering and
Institute of Transportation Studies
University of California, Irvine
Irvine, CA, USA, 92697
Email: nmasoud@uci.edu
Corresponding author

Roger Lloret-Batlle
Ph.D. Candidate
Department of Civil and Environmental Engineering and
Institute of Transportation Studies
University of California, Irvine
Irvine, CA, USA, 92697
Email: rlloretb@uci.edu

Abstract

One of the main obstacles that has challenged peer-to-peer (P2P) ridesharing systems in operating as stand-alone systems is reaching a critical mass of participants. Toward this goal, we propose the *P2P ride exchange* mechanism to increase customer retention. This mechanism gives riders the opportunity to purchase other riders' itineraries, while providing suitable alternative rides to the sellers and therefore maximize the number of rider requests ridesharing systems can satisfy. The proposed mechanism maximizes expected user surplus, is robust towards selfish user manipulation, and has very low information requirements. If implemented correctly, P2P ride exchange effectively increases the number of served riders and enhances customer loyalty by engaging customers in the ride-matching process. The paper only focuses on the potential of the proposed schemes and does not delve into important aspects such as the regulatory changes as well as socio-political arguments against such systems that are beyond the scope of the paper.

Key words: Bilateral trading, P2P Ridesharing, Mechanism Design, P2P Exchange

INTRODUCTION

Peer-to-peer (P2P) ridesharing is a shared mobility alternative in which peer drivers and riders (passengers) share the space in the drivers' personal vehicles. The term *peer-to-peer* points out the fact that drivers are not hired by companies to transport passengers, but are rather using their personal vehicles to carry out personal activities, and this makes them peers to riders. The term *dynamic* highlights the fact that customers can join the system at any moment in time, and do not have to book their trips in advance.

P2P ridesharing manages to eliminate vehicles from roads by getting people who travel in the same direction in the same vehicle. It benefits drivers, riders, the community, and the environment. Drivers receive monetary compensation for the service they provide while following their own daily schedules, and riders are charged less than other transportation alternatives, such as taxies.

As opposed to traditional service businesses where servers belong to the business and their number is proportional to the demand for service, in P2P ridesharing systems servers are also customers. Therefore, it is important for the system operator to attract the right proportion of riders and drivers.

Another feature of P2P ridesharing systems is that drivers typically have specific locations where they start and end their trips, and tight travel time windows to carry out their trips. This limits the spatiotemporal coverage of the network by drivers. Therefore, in order to serve a higher number of riders, the system needs to increase the spatiotemporal coverage of the network by increasing the number of drivers. To motivate, attract and retain a high number of drivers, a high number of riders is required. Therefore, the number of customers in a P2P ridesharing system should pass a certain critical mass, with a specific required proportion of drivers to riders, in order for the system to be able to operate without outside help.

A ride-matching algorithm is the engine of a P2P ridesharing system, and it decides how drivers and riders should be paired together. Except for very simple and non-efficient ridesharing systems (which we will discuss later), the ride-matching problems are generally hard to solve. A good ride-matching method is one that can provide the highest number of matches, in an attempt to engage the highest number of customers and provide the critical mass required.

Customer experience is another factor that plays a great role in the success of P2P ridesharing systems, especially during initial phases. A customer (rider or driver) may give the system a chance by registering in the system a few times, but if he/she is not able to participate, there is a possibility that such a customer would never return to the system. Therefore, it is essential for a P2P ridesharing system to involve and retain as many customers as possible.

Customers in many transportation systems are served on a first-come first-served (FCFS), or otherwise pre-ordered basis. An example of this rule is how green signal time is distributed among conflicting directions at intersections. For P2P ridesharing systems in which customer retention is especially important, considering riders on a FCFS basis is an inefficient use of the very limited available resources (drivers). Dropping the FCFS principle, however, leads to high solution times for the resulting matching problem, and is not an appropriate implementation strategy for dynamic real-time systems. In this paper, we introduce the *P2P ride exchange* as a mechanism to improve the number of matches in a FCFS-based matching solution while maintaining the competitive solution time.

In a system where P2P ride exchange is implemented, riders will still be considered for service on a FCFS basis. Upon joining the system, a rider will be offered the best available itinerary, according to certain criteria which we will discuss later. However, a rider will be given the chance to buy another rider's itinerary under specific circumstances. Purchasing an itinerary from a previously matched rider is in fact reversing the FCFS rule. This exchange of rides is accompanied with an exchange of money through the system. Since the objective of the system from implementing P2P ride exchange is to increase the total number of matched riders, only riders for whom an alternative itinerary is available will receive a proposal to sell their current itineraries, in exchange for monetary compensation.

There are, however, considerable regulatory obstacles to overcome for such P2P exchange or trade schemes to be used in transportation systems. The legal battles faced by ridesharing companies are well-known, transportation supply being considered a public good, any breaking of the traditional FCFS operational paradigms also could face objections based on socio-political arguments of inequity across users. While important, such topics are considered beyond the scope of this paper that focuses only on showing the performance potential of the proposed schemes.

RELATED WORK

P2P ridesharing systems are a member of the family of shared-use mobility alternatives. There is an abundance of work in the literature on the benefits ridesharing systems have to offer in terms of reduced direct and indirect cost to the environment and the society *(1, 2, 3, 4)*.

Despite these benefits, ridesharing systems have been hard to be run as stand-alone systems. Furuhata et al. *(5)* conduct a thorough survey of different types of ridesharing systems, and discuss some of the challenges that has prevented these systems from reaching their potential, despite the improvements in communication technologies, prevalence of GPS-enabled cell-phones, and ease of developing cellphone applications that greatly facilitate participating in ridesharing systems. Ultimately, for a ridesharing system to operate successfully, it has to attract a minimum number of customers. An essential challenge practitioners face is finding the most effective way to build this critical mass *(6, 7, 8)*. James Shield, previously of Carma Technology Corporation that develop carpooling applications, is quoted in *(9)* to believe that although there is no definitive answer to this question, attracting a higher number of drivers, increasing marketing efforts, improving the technology, and attempting to use a societal, behavioral approach to engage people and make habits are all valid approaches.

Research in the field of marketing has found customer satisfaction, among other factors, to be a great predictor of customer retention rate *(10, 11, 12)*. A satisfied customer not only has a higher probability of returning to the system, but generates positive word of mouth (WOM) that helps in attracting new customers *(13)*. Research has shown that WOM is a more important factor when it comes to deciding on services, rather than goods *(14)*. In addition, the cost of customer acquisition is about five times the cost of customer retention *(15)*, suggesting that the few first experiences of a customer with the system plays a central role on its long-term success. In light of this research, it is very important for ridesharing systems, especially in their initial stages, to serve as many ride requests at possible. Trying to encourage a high number of drivers to participate in the system is another goal of a ridesharing system, albeit not as important as the first one. The reason is that first, drivers who participate in ridesharing usually receive a base fare regardless of the extent of their contribution. For a given level of demand (ride requests), there is an optimal amount of supply (drivers) above which the contribution of additional supply is only marginal, and therefore attracting drivers with only marginal contribution is not financially wise. Second, drivers in ridesharing systems are traveling to attend their personal activities. Even if they are not included in the system on a regular basis, entering their fixed daily schedules in the system only once could earn them extra revenue.

The ride-matching algorithm used by a ridesharing system plays an important role on the number of riders the system can serve. The simplest form of ride-matching algorithm matches each driver with a single rider, if the two share the same origins and destinations *(16)*. This problem can be formulated as a maximum cardinality matching problem in a bipartite graph, and solved quickly using efficient algorithms *(17)*. The more sophisticated ride-matching problems are capable of routing drivers' vehicles in order to put them in vicinity of riders *(16, 18, 19, 20, 21, 22, 23, 24)*, allocating more than one rider to a single driver *(18, 19, 20, 22, 23, ,25, 26)*, proposing multi-hop itineraries for riders, where they can transfer between drivers *(21, 24, 25, 27, 28)*, and finally considering multiple riders and drivers in the same problem *(24, 25)*.

A many-to-many ride-matching problem, where a rider can transfer between multiple drivers, and a driver can carry multiple passengers at any moment in time, is the most comprehensive ride-matching problem, and can yield the highest number of matches. Not surprisingly, a many-to-many problem is also the hardest matching problem to solve. In a one-to-many problem, a rider can transfer between drivers, but the matching problem is solved for one rider at a time. Although this problem is still computationally hard to solve, it has one less degree of flexibility compared to a many-to-many problem, and therefore can be solved in a shorter period of time.

In addition to serving as many riders as possible, the responsiveness of the system to dynamic ride requests could play a role in customer satisfaction as well. In general, ridesharing systems have two basic operating strategies. The first strategy is to secure itineraries for riders on a FCFS basis. This strategy coupled with an algorithm that can solve the (one-to-many) matching problem in a short period of time creates a ridesharing system that is capable of handling dynamic ride requests in real-time.

The second strategy is to re-optimize the system on a rolling-horizon basis. In this type of systems, the (many-to-many) ride-matching problem that needs to be solved in each re-optimization period is larger in size, and less likely to be solved in real-time, especially in more sophisticated many-to-many ridesharing systems. On the other hand, these systems can serve a potentially higher number of riders.

There is an evident trade-off between the two factors that influence customer satisfaction: serving higher number of riders, and system responsiveness. In this paper, we design the P2P ride exchange mechanism as a tool to improve the number of served riders from the optimal solution of a one-to-many system, while maintaining its higher responsiveness, and demonstrate the role this mechanism plays in customer retention in ridesharing systems.

ONE-TO-MANY RIDESHARING ALGORITHM

We solve the one-to-many ridesharing problem using the dynamic programming (DP) algorithm proposed in *(29)*. This algorithm is suitable for the purpose of P2P ride exchange, because first, using this algorithm, real-life size problems can be solved in a very short period of time (a fraction of a second in most settings), and second, all feasible solutions to the problem are retrievable using the set of trees that are generated while solving the problem. In this section, we provide a short review of the algorithm.

Let us define a graph $G = (N, L)$. Each vertex $n \in N$ in this graph is a tuple $(s, t) \in S \times T$, where S is the set of stations in the network, and includes all trip origins and destinations and a set of fixed transfer locations, and T is the set of time intervals during the study time horizon. An edge $\ell = (n_1, n_2) = (s_1, t_1, s_2, t_2) \in L$ in this graph corresponds to trip between stations s_1 and s_2 that begins at interval t_1 and ends at interval t_2. Each participant $p \in P$ (rider or driver) upon registering in the system provides information on their origin and destination stations (OS_p, DS_p respectively), and their travel time window, $[ED_p, LA_p]$, where ED_p is the earliest departure time from the origin station, and LA_p is the latest arrival time at the destination station. For each participant p, a graph G_p, can be constructed based on these parameters. The purpose of the algorithm is to search on this graph for a minimum cost path that starts from station OS_p, and ends at station DS_p. We define the cost of a path as a weighted linear combination of the in-vehicle travel time, waiting travel time, and number of transfers. Note that a rider can use multiple vehicles/modes to accomplish his/her trip.

PEER-TO-PEER RIDE EXCHANGE

Dynamic ridesharing systems should have the capability of matching riders and drivers in real-time. Since participants have rather tight travel time windows, as soon as a rider joins the system, the attempt to find a match for him starts by running the DP algorithm described in the previous section. If all the itineraries generated by the algorithm are infeasible due to their conflicts with itineraries of the previously assigned riders (i.e. if the itineraries use the same drivers, but through different paths), then the system evaluates the possibility of a trade. In this section, we show through an example the benefits of a P2P exchange program, discuss the conditions under which trade can happen, and devise a mechanism that ensures a fair trade.

Let P_r denote the set of itineraries for rider r. Each itinerary has a value that is determined by a pre-specified objective function (the DP objective function), based on which the itineraries within P_r are ranked. Let p_i^r denote the i^{th} itinerary of rider r, and $d(p_i^r)$ denote the set of drivers who contribute to itinerary p_i^r. Note that there is no need to know all members of set P_r in advance, but we will generate them as (and if) needed. Furthermore, let p_k denote the itinerary of the assigned driver k.

Once rider r joins the system, the system uses the DP algorithm to generate a set of trees from which members of set P_r can be retrieved. The system starts by evaluating members of set P_r in order of their ranking. If an itinerary with no conflicts with the driver itineraries is found, this itinerary will be assigned to rider r. If the system exhausts all members of set P_r, and is not successful in finding a non-conflicting itinerary for rider r, then it considers the possibility of a trade.

Assume that rider 1 enters the system, and has two itineraries: $P_1 = \{p_1^1, p_2^1\}$, where $d(p_1^1) = \{d_1\}$ and $d(p_2^1) = \{d_2\}$. The left hand side picture in Figure 1 shows the rider and his itinerary set. Assuming that the minimum cost itinerary for this rider is the first one, this itinerary will be announced to both rider 1 and driver 1.

Now, assume that rider 2 joins the system. Because rider 1's itinerary has been fixed, there are no feasible itineraries for rider 2. However, rider 2 has a chance to buy rider 1's itinerary, if rider 1 has not started his trip yet. The right hand side picture in Figure 1 shows this scenario after the trade. In this trade, rider 2 buys rider 1's assigned itinerary, and by doing so releases driver 1, who in turn forms a feasible itinerary for rider 2. Rider 1 switches to a less convenient itinerary in exchange for a monetary compensation. This trade's contribution to customer retention is double-folded. Not only are both riders served, but now both drivers are participating in the system as well.

Note that it was possible to obtain the same optimal solution by solving a many-to-many ride-matching problem that is capable of considering both riders at the same time. There are, however, two issues with such an approach: (1) An optimal matching algorithm that could consider both riders at the same time is computationally hard to solve (for larger problem sizes), and therefore cannot yield solutions in real-time. (2) Even if the system is equipped with a many-to-many ride-matching algorithm that can yield solutions in a moderate period of time, such a system has to work on a rolling horizon basis, and be re-optimized periodically. If rider 2 enters the system after the re-optimization with rider 1 is done, and the next re-optimization is scheduled for after rider 1 departed from his origin station, rider 2 would still be left without a ride.

The system studies the possibility of a trade if the following three conditions hold. First, the buyer does not have any feasible itineraries. Second, the seller has an alternative feasible itinerary to his current one, and third, both parties will be better off with the trade than without it.

The monetary transfer from the buyer to the seller covers the additional cost the seller has to undertake due to switching itineraries. This cost includes the additional monetary cost due to a potentially increased travel distance, and a compensation to the seller for a potentially increased travel time. A proportion of this money will be used by the system operator to cover the cost of the seller's more expensive new itinerary, and the rest will be transferred to the seller himself.

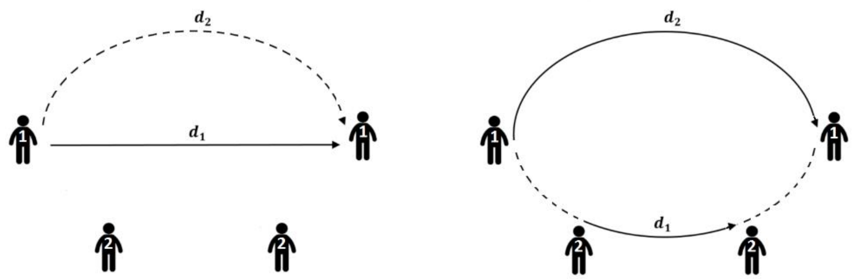

FIGURE 1 An example of a trade. The left-hand side picture shows the system without the trade, and the right hand side shows the system after the trade. The solid lines represent the itineraries assigned to riders.

The Scope of the Trade

Assume a set of itineraries P^r for rider r. Drivers contributing to itinerary i are stored in set $d(p_i^r)$. Let us divide members of set $d(p_i^r)$ into two mutually exclusive sets, $d_a(p_i^r)$ and $d_f(p_i^r)$. Drivers in set d_a have been previously assigned to other riders, but their corresponding riders' trips have not started yet. Drivers in set d_f are free, and have not been assigned to any riders. p_i^r is a feasible itinerary for rider r if one of the following conditions hold: (1) $d_a(p_i^r) = \emptyset$, i.e. none of the drivers that contribute to the itinerary are assigned to other riders, and (2) $\forall k \in d_a(p_i^r), p_i^r(k) \in p_k$, i.e. drivers in set $d_a(p_i^r)$ can still follow their previous itineraries. If none of these two conditions hold, then the system tries to find a good candidate for a trade amongst the assigned riders.

To find the candidates for a trade, the system has to first identify the itinerary that rider r is interested in. It starts from the best itinerary, i.e. p_1^r, and moves to the next itinerary if the trade on the current itinerary is not possible.

In order for the system to offer itinerary p_i^r to rider r, it has to free all the drivers in set $d_a(p_i^r)$ from their previous assignments. Therefore, the system has to find all the riders who are using these drivers, and try to find alternative itineraries for them as well. These riders form the sellers in the first level of trade (Figure 2(a)). In order for the system to propose an exchange to a rider r' in the first level of trade, it should find an alternative itinerary for the rider first. This task can be accomplished by identifying the set of assigned drivers for the rider (d'_a), finding the rest of the riders whose itineraries are affected by these drivers, and finally finding alternative itineraries for them as well. This procedure continues until the system reaches a level of trade where all riders have itineraries with free drivers (or non-conflicting assigned drivers).

The system will then start proposing trades to riders, starting from those in the last level of trade. In order for a trade to be approved at any level, all the riders at that level should approve the trades proposed to them. For the n^{th} level of trade to take place, the trade at level $n + 1$ should have been approved. Once all riders in a given level of trade approve the proposed trades, the system can move to a higher level of trade (moving upwards in Figure 2(a)). Therefore, it is clear that the higher is the levels of trade, the less likely it is for rider r to obtain itinerary p_i^r.

Another complication with this approach is that even if even one rider does not approve the trade at a certain level, the trade cannot happen. In this case all the riders in the same and lower levels who have accepted the trades proposed to them have to go back to their previous itineraries. Therefore, in order to simplify this procedure and make it easy to implement in practice, this paper only considers trades in settings where the level of trade is 1, and the number of riders in the first level of trade is limited to 1 as well, i.e. the set of assigned drivers affect only the itinerary of a single previously assigned rider (Figure 2(b)). These simplifications limit the trade to be between two individuals only: the buyer, r, and the seller, r'.

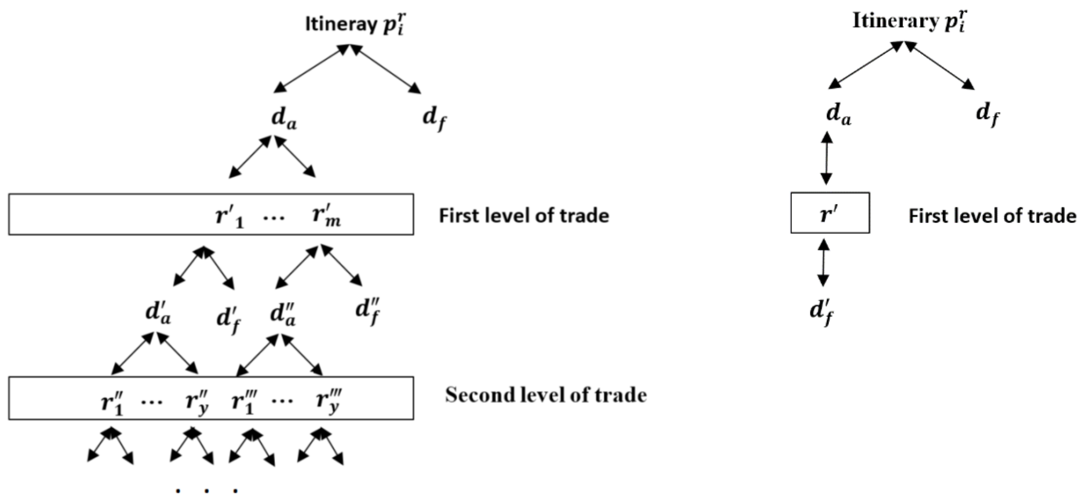

(a) The general level of trade **(b) Level of trade considered in this paper**

FIGURE 2 Levels of trade

P2P Ride Exchange Mechanism

Besides ensuring that the trade makes both parties better off, the designer (operator) should also ensure that the trading parties cannot manipulate the outcome of the trade. Since they hold private information not known to the operator, this could lead to an inefficient outcome. This issue is addressed by modelling the trade from a mechanism design perspective. Informally, a mechanism is a method that defines rules for a game with incomplete information (Bayesian game) to influence agents' behavior and reach a particular goal, which in this case is efficiency maximization. Excellent introductions to mechanism design can be found in *(30, 31)*. The basic definitions are provided next, but a complete understanding of mechanism design may require reading the above introductory references.

Let $I = \{1, \dots, n\}$ be the set of agents. Each agent has a type (value of time), $\theta_i \in \Theta_i$ which is private. $\Theta = \times_{i \in I} \Theta_i$ is the type profile set. Agent i has the (quasilinear) utility function: $u_i(\theta_i, \theta_{-i}; \theta_i) = v_i(k(\theta_i, \theta_{-i}); \theta_i) - p_i(\theta_i, \theta_{-i})$. Where $v_i(\cdot)$ is his valuation and $p_i(\cdot)$ is the price charged to him. The types before the semicolon are the types announced to the designer, while the type at its right is the agent i's actual type. A (direct revelation) mechanism is composed of two interrelated functions. Firstly, an allocation function $k: \Theta \to K$ that maps the type space to an outcome set K. That is, for every announced type profile, an allocation $k(\boldsymbol{\theta}) \in K$ is given. In the case of a bilateral trade, K is composed of the two allocations (trading states): either there is trade or there is not. The allocation rule that maximizes the sum of agents' valuations is the efficient allocation rule, $k^*(\boldsymbol{\theta}) \in K$. Secondly, there is a payment function $p: \Theta \to \mathbb{R}^N$. This function assigns a transfer amount to every agent i in accordance with its announced type θ_i.

Mechanism design defines concepts that address how the strategic interests of agents are satisfied. The main one is truthfulness, or incentive compatibility, which states that truthful bidding forms an equilibrium. In other words, any participating agent is always better off by truthfully eliciting its type rather than lying, subject to others telling the truth. A mechanism (k, p) is (Dominant-Strategy) Incentive Compatible (DSIC) if:

$$v_i(k(\theta_i, \theta_{-i}); \theta_i) - p_i(\theta_i, \theta_{-i}) \geq v_i\big(k(\hat{\theta}_i, \theta_{-i}); \theta_i\big) - p_i(\theta_i, \theta_{-i}), \forall i \in I, \forall \theta_i \in \Theta_i, \forall \theta_{-i} \in \Theta_{-i}, \forall \hat{\theta}_i \in \Theta_i$$

Besides truthfulness, a designer is interested in the users' willingness to participate in the mechanism, called individual rationality. A mechanism (k, p) is Ex-Post Individual Rational (EPIR) if:

$$v_i(k(\theta_i,\theta_{-i});\theta_i) - p_i(\theta_i,\theta_{-i}) \geq \bar{u}_i(\theta_i) \quad \forall i \in I, \forall \theta_i \in \Theta_i, \forall \theta_{-i} \in \Theta_{-i}$$

Here $\bar{u}_i(\cdot)$ is agent i's utility from not participating in the mechanism. If EPIR is satisfied, an agent would be willing to truthfully participate in the mechanism rather than stay out. Finally, the designer may be interested in the mechanism being self-sufficient from a budget point of view. In this way, the mechanism should not require an external subsidy to achieve the desired outcome. This condition is known as balanced budget. A mechanism (k,p) is weakly budget balanced (WBB) if:

$$\sum_{i\in I} p_i(\theta_i,\theta_{-i}) \leq 0$$

If the sum of payments is equal to zero, Strict Budget Balancedness (SBB) is satisfied.

A "pessimistic" approach is followed *(32)*, in that the designer calculates the expected surplus that is guaranteed among all the admissible strategy profiles of the agents, with the actual strategies played being unknown to the designer. A strategy is said to be admissible if it is not weakly dominated. It is assumed that the operator has a prior on the private information from agents, but the agents themselves do not have a prior of the other agents' type, as in the case of weaker truthfulness concepts *(33)*. This framework is very convenient for our purposes, since the mechanism has to be designed far in advance, with no previous experience or learning on trading outcomes from either the users' or the designer's part, while, at the same time, it guarantees an increase in of the number of served riders to achieve user permanence in the system.

The trade is modelled as a bilateral trade with private information *(33, 34)*. Let $I = \{1,2\}$ be the set of agents, $i = 1$ being the seller and $i = 2$ being the buyer. Each agent has a type $v_i \in [\underline{v}_i, \overline{v}_i]$. The mechanism lies in the space $(q,p) \in [0,1] \times \mathbb{R}$, where q is the probability of trade and p is the payment from the buyer to the seller. Bilateral trading, by definition, satisfies the strict budget balance property, thus the seller has utility $u_1 = v_1 q + p$ and the buyer $u_2 = v_2 q - p$. For clarity in the exposition, we use the following change in notation $c_1 \stackrel{\text{def}}{=} -v_1$. Trade occurs when $v_2 > p > c_1$: both agents are proposed the price p and if both agree, the trade takes place. The designer's surplus function to be maximized is $w\big((q,p);\theta\big) = (v_2 - c_1)q$.

Proposition 1 *(32)*: A posted-price mechanism with price τ for the bilateral trade setting is a revelation mechanism (q,p) such that:

$$(q,p) = \begin{cases} (1,\tau) & if\ v_2 > \tau > c_1 \\ (0,0) & otherwise \end{cases}$$

This mechanism is DSIC, EPIR, SBB and guarantees the following expected surplus $W(\tau)$:

$$W(\tau) = \iint_{(c_1,v_2)=(\underline{c_1},\tau)}^{(c_1,v_2)=(\tau,\overline{v_2})} (v_2 - c_1)\phi(c_1,v_2)\ \ dv_2 dc_1$$

which is convex with regard to the price τ. Therefore, an optimal price τ^* exists and it is unique. ϕ has to be the continuous density of an absolutely continuous (w.r.t to a Lebesgue measure on $\mathbb{R}^2$) function Φ. This assumption is fulfilled by most continuous probability distributions, such as the lognormal.

Instead of valuing an object by a scalar as in the original bilateral trading environment, riders value their allocation (assigned ride) by its generalized cost, which is an affine transformation of their type. This cost is the product of the travel time and the sum of the value of time θ_i, plus the fare per unit of time δ_r. These valuations are normalized with regard to the initial situation (no trade) to fit the bilateral trading original setting: c_1 and v_2 are in fact the valuation difference between states "trade" and "no trade". Thus, when there is no trade, both c_1 and v_2 are zero. When there is a trade, $c_1 = \theta_1(t'_{r1} - t_{r1}) + c'_{r1}t'_{r1} - c_{r1}t_{r1}$ and $v_2 = \theta_2(t_{out} - t_{r2}) + c_{out}t_{out} - c_{r2}t_{r2}$. Here, t_{r1}, t'_{r1}, t_{r2}, and t_{out} refer to the travel time of the seller's current and new itineraries, travel time of the buyer's itinerary, and the travel time of the buyer's outside alternative, respectively, and c_{out} is the cost of the outside option to the buyer.

Surplus integral bounds $\underline{c_1}$ and $\overline{v_2}$ depend on the relative magnitude of the time and ride cost which are input parameters. For instance, if $t_{out} < t_{r2}$, the travel time difference is valued negatively and $\overline{v}_2 = v_2(0)$. In this particular case (with also $t'_{r1} > t_{r1}$), the optimal posted price is calculated as:

$$p^* = \underset{p}{\operatorname{argmax}} \left(\iint_{(c_1,v_2)=(c_1(0),p)}^{(c_1,v_2)=(p,v_2(0))} (v_2 - c_1)\phi(c_1, v_2) \;\; dv_2 dc_1 \right)$$

Each expected surplus integral is solved by numerical simulation and the optimal price is found by golden section search. Since the VOT distributions of the agents are assumed to be independent, the VOT joint distribution is the product of the two marginal distributions. The actual empirical VOT distribution used in our study comes from a survey on households conducted in Stockholm, Sweden in 2005 *(35)*. In that research, the Stated Preferences (SP) choice scenarios are composed of car alternatives that differ on attributes such as travel times and travel costs. Since only the main statistics are available in the publication, the distribution is recalibrated as a lognormal distribution given these statistics. Its parameters are location $\mu = 2.16$ and scale $\sigma = 0.40$.

Pricing

There are many factors that should be taken into consideration in determining the fare for ridesharing services. Setting the right price is essential to the success of a ridesharing system, and deserves the designing of a separate mechanism which ensures that no incentive exists for drivers and riders to falsely report their preferences in order to affect the amount of transaction.

The fare a rider is charged in our system is made of two components. The first component is a time-based fee. Assume that rider r's itinerary is composed of a set of traveling links L^*, and that on each link $l \in L^*$, n_l number of individuals (including the driver) share the same vehicle with the rider. The cost of travel on each link is equally shared by the individuals who travel on the link. Therefore, the time-based fee of rider r will be $\sum_{l \in L^*} \frac{\delta . d_l}{n_l}$. Note that a system operator could easily replace the time-based free with a distance-based free, or consider a combination of both.

The second part of the fair is a fixed base fair. Since drivers may have to divert from their shortest/preferred paths in order to accommodate riders, they need to be compensated for this extra travel. We calculate the base fare based on the average extra travel time drivers have to spend in the network, assuming an average speed of 40 mph, and a payment of 60 cents per mile. These fares could vary for different times the day, and days of the week, based on the composition of the ridesharing system, i.e. number of participants, and driver to rider ratio. Although a pricing scheme that can distribute fares among drivers based on their contribution to the system may be fairer, in the interest of simplicity we use the more preliminary pricing scheme introduced in this section.

NUMERICAL STUDY

In order to study the impact of the P2P exchange mechanism of a ridesharing system, multiple random instances of the ridesharing problem are generated and solved. Problems are generated in a 49 square mile network. The network is divided into two non-intersecting sections (clusters) that represent the central business district, and the residential areas. It is assumed that trips start from a random location in the first cluster, and end at the second (to represent the PM peak hour). A station is randomly generated within each square mile block of the grid. Stations serve as transfer hubs.

Since efficiency and level of service in a ridesharing system highly depend on the spatiotemporal distribution of trips, we generate different problem sets with different degrees of spatiotemporal proximity by varying total number of participants, and temporal distribution of trips. In order to vary the temporal proximity of trips, we vary the length of the time period within which trips start.

We study three sets of problems. The first set, representing a system with high spatiotemporal proximity, comprises of 1000 participants whose trips start within a 15 minute time interval. The second set, representing a system with moderate spatiotemporal proximity, comprises of 1000 participants whose trips start within a 30 min time period, and the third set, representing a system with low spatiotemporal proximity, comprises of 500 participants whose trips start within a 15 minute time interval

Base Fares

As mentioned in the previous section, this study uses the same base fares for all drivers in a given time period, for example, during weekday morning peak hours. These fares may vary from location to location, and depend on the system composition (number of participants, and ratio of drivers to riders). In this section, base fares are generated for the three problem sets introduced earlier.

Figure 3 displays the number of served riders and involved drivers, and the amount of the base fares for the three problem sets. Each graph in this figure summarizes the output of multiple ridesharing instances with different number of riders and drivers, and is the result of averaging over three randomly generated problems. Figures 3(a) and 3(b) demonstrate the performance of a ridesharing system with low spatiotemporal proximity. The best performance is obtained when the number of riders and drivers are almost equal (250), and even at this peak, only less than 48% of the riders can be served. This low spatiotemporal proximity leads to higher base fares for riders (about \$1.5 for the case of 250 riders and drivers). Figure 3(b) suggests that initially, as the number of rider requests increases, the base fair decreases. Towards the end of the graph, however, since the number of drivers in the system is very small, they have to divert from their shortest paths even further, which causes the base fare to start going up. This trend in base fares can be witnessed in all graphs, regardless of the degree of spatiotemporal proximity.

Figures 3(c) and 3(d) show the performance of a system with moderate spatiotemporal proximity. At its peak, this system can satisfy about 55% of the demand, and the base fare turns out to be about \$1.1 per rider, which not surprisingly is lower than the previous case.

Figures 3(e) and 3(f) demonstrate the performance of a system with high spatiotemporal proximity. Figure 3(e) shows that the number and percentage of served riders is higher compared to the two previous cases. For example, when the system has 500 riders and 500 drivers, about 350 riders can be served (70%). The higher number of served riders results in lower base fares. For example, in the case of 700 riders and 300 drivers (where still 70% of the demand can be served), the base fare can go down to less than \$1 per rider.

Note that base fares in Figure 3 are computed by solving the ridesharing problems without implementing the P2P exchange mechanism. In the next section, we use these fares to determine the feasibility of potential exchanges.

Value of Exchange

In this section, we use the DP algorithm in *(29)* combined with the P2P ride exchange mechanism to find the maximum number of served riders in the three sets of ridesharing problems introduced earlier.

For each participant, a VOT is generated using the lognormal distribution described earlier. Each individual is assumed to have a separate transportation alternative outside of the system, with a travel time equal to the shortest path travel time between the individual's origin and destination. The distance-based cost of the outside alternative is assumed to be equal to that of the ridesharing system. With these assumptions in place, the three previously introduced problem sets are solved. Results are presented in Table 1.

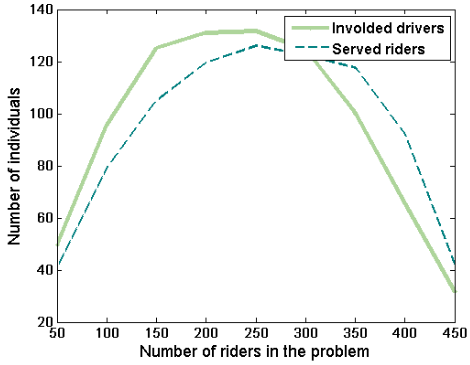

(a) Low spatiotemporal proximity: 500 participants. Trips generated within a 15 min time period.

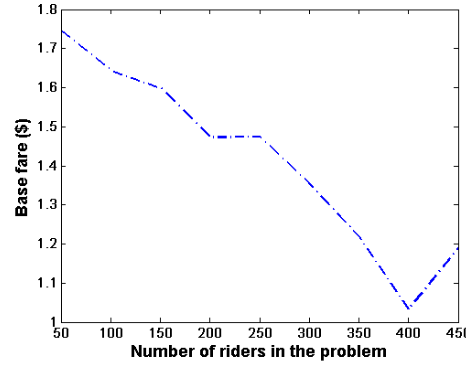

(b) Low spatiotemporal proximity: 500 participants. Trips generated within a 15 min time period.

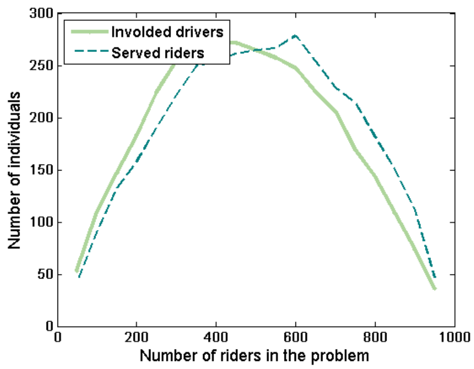

(c) Moderate spatiotemporal proximity: 1000 participants. Trips generated within a 30 min time period.

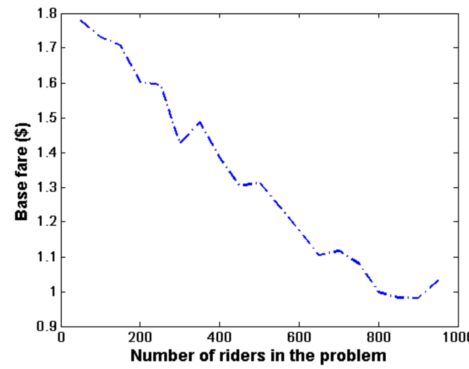

(d) Moderate spatiotemporal proximity: 1000 participants. Trips generated within a 30 min time period.

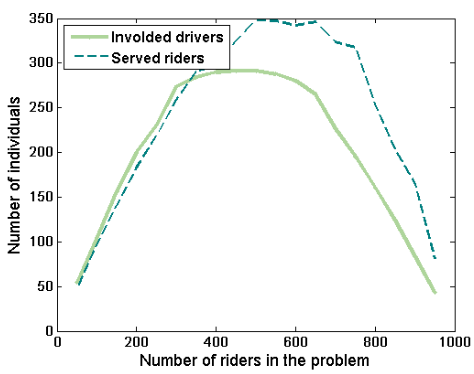

(e) Ridesharing instance with 1000 participants. Trips generated within a 15 min time period.

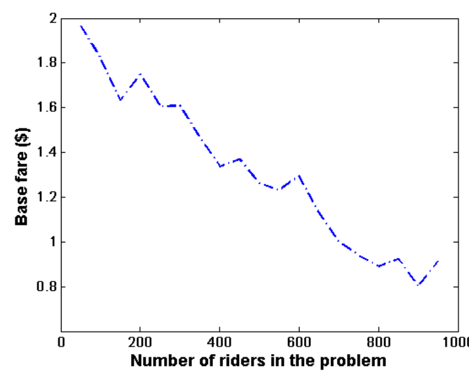

(f) Ridesharing instance with 1000 participants. Trips generated within a 15 min time period.

FIGURE 3 Base fares for ridesharing systems with different levels of spatiotemporal proximity between trips

For each problem instance in Table 1, the release period is the length of the time period during which all trips start. Each problem is solved for three different instances, and the average values are reported.

Table 1 suggests that as the spatiotemporal proximity of participants in a ridesharing system increases, so does the number of served riders, and the number of exchanges. In the problem instances with low spatiotemporal proximity, number of exchanges is considerably smaller compared to the problems with moderate and high spatiotemporal proximity. This suggests that the initial mass of participants in the systems positively reinforces the number of successful exchanges. Although the number of exchanges might not seem too high, there are two points that should be taken into consideration. The first point is that the higher number of served riders due to the exchange does not have a one-to-one impact on the performance of the system. It is true that the number of served riders increases only by the number of successful exchanges, but the impact on customer retention and the reputation of the system should also be taken into consideration.

To study customer retention, we use the same set of riders in each of the three instances generated for problems in Table 1 (i.e. only the driver set has changed for each problem instance). It is assumed that a rider does not return to the system if he/she has three failed experiences. By simulating a three day experience for each rider, the number of retained customers due to the exchange has been obtained and presented in Table 1. To compute customer retention, it is assumed that a rider will consider using the system again if he/she has at least one successful experience. In addition to creating a positive experience for these riders and increasing the probability of them returning to the system, using P2P exchange could eliminate the possible negative WOM that could have been generated by these riders had they not been served, and could even replace them with a positive WOM.

No. of participants	No. of Riders	No. of Drivers	Release period (min)	No. of satisfied riders (including exchange)	No. of exchanges	No. of involved drivers	No. of additional drivers	No. of retained riders due to exchange	Potential number of additional first level exchanges
Low spatiotemporal proximity									
500	100	400	15	79	1	93	3	3	7
500	250	250	15	126	5	129	3	15	13
500	400	100	15	92	1	65	1	3	12
Moderate spatiotemporal proximity									
1000	200	800	30	158	5	179	5	11	12
1000	500	500	30	265	10	262	6	28	22
1000	800	200	30	182	7	141	3	20	15
High spatiotemporal proximity									
1000	200	800	15	183	6	195	5	14	18
1000	500	500	15	349	12	283	9	30	26
1000	800	200	15	254	5	160	1	15	20

TABLE 1 P2P Exchange and customer retention

The second point is the possibility of increasing the scope of the trades. This study concentrated on the simplest possible scenario for trades, where there is a single buyer and a single seller. Table 1 displays the additional number of possible exchanges that could take place if all trades in the first level of trade were considered. Moving even further down the levels of trade could have a considerably larger impact on the number of exchanges. These exchanges translated into customer retention rate suggest the possibility of promising contributions of P2P exchange into the success of ridesharing systems. In order to reach to these levels, however, more sophisticated mechanisms need to be devised.

CONCLUSION

We introduced P2P ride exchange as a mechanism to obtain high system efficiency in dynamic P2P ridesharing systems. Although the application of P2P exchange in this paper was limited to the simplest scenarios, the results suggest that extending this mechanism to more complicated scenarios could contribute to the efficiency of ridesharing systems even further, and lead to higher customer retention rates. In addition, it was observed that there is a positive correlation between the spatiotemporal proximity of trips in a ridesharing setting, and the number of successful trades by P2P exchange.

The mechanism employed is robust towards selfish manipulation from users and helps in increasing ridership with minimal information requirements from both operator and users. Despite these low information requirements, an increase in the probability of getting a ride is achieved, leading to further user permanence into the system, even in situations where the system operation has not yet started and no data is available. Once more data is available, a new mechanism that takes into consideration observed user interaction may be designed.

We conclude with a comment on the essential paradigm that lies behind the entire concept, namely that the design facilitates the easing of the traditional first-come-first-served (FCFS) rule in the consumption of transportation supply by users. While the tradition rule is considered the most fair, it is arguably in place mostly because the lack of information exchange possibilities always prevented any consideration of "even fairer" paradigms based on user preferences. This paper is an early attempt at analyzing such a paradigm for ride-sharing systems. We repeat that the regulatory stumbling blocks as well as socio-political inequity arguments against the proposed system may be substantial and that we have not attempted to address such issues that are beyond the scope of this paper, which focuses only on illustrating the potential success of the proposed scheme if such hindrances can be overcome.

REFERENCES

1. Chan, N. D., and S. A. Shaheen. Ridesharing in North America: Past, present, and future. *Transport Reviews*, Vol. 32, No. 1, 2012, pp. 93–112.
2. Morency, C., The ambivalence of ridesharing. *Transportation*, Vol. 34, No. 2, 2007, pp. 239–253.
3. Kelley, K. Casual carpooling enhanced. *Journal of Public Transportation*, Vol. 10, No. 4, 2007, pp. 119–130.
4. Heinrich, S. *Implementing real-time ridesharing in the San Francisco Bay Area*. Ph. D. thesis, San Jose State University. 2010.
5. Furuhata, M., M., Dessouky, F., Ordóñez, M. E. Brunet, X. Wang, and S. Koenig. Ridesharing: The state-of-the-art and future directions. *Transportation Research Part B: Methodological*, Vol. *57*, 2013, pp. 28-46.
6. Cervero, R., and B. Griesenbeck. Factors influencing commuting choices in suburban labor markets: a case analysis of Pleasanton, California.*Transportation Research Part A: General*, Vol. 22, No. 3, 1988, pp. 151-161.
7. Brereton, M., and S. Ghelawat. Designing for participation in local social ridesharing networks: grass roots prototyping of IT systems. In *Proceedings of the 11th Biennial Participatory Design Conference*, 2010, pp. 199-202).
8. Raney, S. San Francisco to Silicon Valley, California, Instant Ridesharing with Transfer Hub. In *Transportation Research Record: Journal of the Transportation Research Board*, No. 2143, Transportation Research Board of the National Academies, Washington, D.C., 2010, pp. 134-141.
9. Gaynor, T., *Solving The Critical Mass Problem. What's The Simplest Thing That Could Work?*, 15/01/2015,https://carmacarpool.com/blog/2013/01/15/solving-the-critical-mass-

problem-what%E2%80%99s-the-simplest-thing-that-could-possibly-work/, Accessed 7/29/2015.
10. Gustafsson, A., M. D. Johnson, and I. Roos. The effects of customer satisfaction, relationship commitment dimensions, and triggers on customer retention. *Journal of marketing*, Vol. 69, No. 4, 2005, pp. 210-218.
11. Rust, R. T. and A. J. Zahorik, A. J. Customer satisfaction, customer retention, and market share. *Journal of retailing*, Vol. 69, No. 2, 1993, pp. 193-215.
12. Ranaweera, C., J. Prabhu, J. On the relative importance of customer satisfaction and trust as determinants of customer retention and positive word of mouth. *Journal of Targeting, Measurement and Analysis for marketing*, Vol. 12, No. 1, 2003, pp. 82-90.
13. Söderlund, M. Customer satisfaction and its consequences on customer behavior revisited: The impact of different levels of satisfaction on word-of-mouth, feedback to the supplier and loyalty. *International Journal of Service Industry Management*, Vol. 9, No. 2, 1998, pp. 169-188.
14. Buttle, F. A. Word of mouth: understanding and managing referral marketing. *Journal of strategic marketing*, Vol. 6, No. 3, 1998, pp. 241-254.
15. Pfeifer, P. E. The optimal ratio of acquisition and retention costs. *Journal of Targeting, Measurement and Analysis for Marketing*, Vol. 13, No. 2, 2005, pp. 179-188.
16. Agatz, N. A., A. L. Erera, M. W. Savelsbergh, and X. Wang. Dynamic ride-sharing: A simulation study in metro Atlanta. *Transportation Research Part B: Methodological*, Vol. 45, No. 9, 2011, pp. 1450 - 1464.
17. Alt, Helmut, et al. Computing a maximum cardinality matching in a bipartite graph in time $(n\text{^}1.5\sqrt{}(m/\log n))$. *Information Processing Letters*, Vol. 37, No. 4, 1991, pp. 237-240.
18. Calvo, R. W., F. de Luigi, P. Haastrup, and V. Maniezzo. A distributed geographic information system for the daily carpooling problem. *Computers & Operations Research*, Vol. 31, No. 13, 2004, pp. 2263-2278.
19. Baldacci, R., V. Maniezzo, and A. Mingozzi. An exact method for the car pooling problem based on lagrangean column generation. *Operations Research*, Vol. 52, No. 3, 2004, pp. 422-439.
20. Teodorovic, D. and M. DellOrco. Bee colony optimization-a cooperative learning approach to complex transportation problems. In *Advanced OR and AI Methods in Transportation: Proceedings of 16th Mini-EURO Conference and 10th Meeting of EWGT*, 2005, pp. 51-60.
21. 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.
22. Herbawi, W. and M. Weber. A genetic and insertion heuristic algorithm for solving the dynamic ridematching problem with time windows. In *Proceedings of the fourteenth international conference on Genetic and evolutionary computation conference*, 2012.
23. Di Febbraro, A., E. Gattorna, and N. Sacco. Optimization of Dynamic Ridesharing Systems. *Transportation Research Record: Journal of the Transportation Research Board*, No. 2359, 2013, pp. 44-50.
24. Masoud, N., and R. Jayakrishnan. A decomposition Algorithm to solve the multi-hop peer-to-peer ride-matching problem. In *Transportation Research Board 94th Annual Meeting*, No. 15-2755, 2015, 2015.
25. Ghoseiri, K. *Dynamic rideshare optimized matching problem*. PhD Dissertation at University of Maryland, 2013.
26. Schaub, T., G. Friedrich, and B. O'Sullivan. ECAI 2014: 21st European Conference on Artificial Intelligence, Volume 263, 2014.
27. Herbawi, W. and M. Weber. Ant colony vs. genetic multiobjective route planning in dynamic multi hop ridesharing. In *Tools with Artificial Intelligence (ICTAI)*, 2011a.
28. Herbawi, W. and M. Weber. Comparison of multiobjective evolutionary algorithm for solving the multiobjective route planning in dynamic multi-hop ridesharing. In *Proceedings of the 11th European conference on Evolutionary computation in combinatorial optimization*, Torino, Italy, 2011b, pp. 84 - 95.

29. Masoud, N., and R. Jayakrishnan. A Real-Time Algorithm to Solve the Peer-to-Peer Ride-Matching Problem in a Flexible Ridesharing System. UCI-ITS-WP-15-1. Pdf: http://www.its.uci.edu/its/publications/papers/ITS/UCI-ITS-WP-15-1.
30. Mas-Colell, A., M. D. Whinston, and J. R. Green, J. R. Microeconomic theory. *The Canadian Journal of Economics*. Vol. 21, No. 2, 1988, pp. 436-438.
31. Nisan, N., Ronen, A. Algorithmic Mechanism Design, *Games and Economic Behavior*, 2001, Vol. 35, pp. 166-196.
32. Yamashita, T. Implementation in Weakly Undominated Strategies: Optimality of Second-Price Auction and Posted-Price Mechanism, *The Review of Economic Studies*, 2015, Vol. 82, No. 3, pp. 1223-1246.
33. Myerson, R. B., and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, Vol. 29, No 2, 1983, pp. 265-281.
34. Hagerty, K. M., and W. P. Rogerson. Robust trading mechanisms. *Journal of Economic Theory*, Vol. *42*, No. 1, 1987, pp. 94–107.
35. Abou-Zeid, M., Ben-Akiva, M., Bierlaire, M., Choudhury, C., and Hess, S., Attitudes and Value of Time Heterogeneity, *Applied Transport Economics: A Management and Policy Perspective*, Van de Voorde, E. and Vanelslander, T. (eds.), Uitgeverij De Boeck nv, Antwerp, Belgium, 2010, pp. 523-545.