

A Mechanism for Dynamic Ride Sharing based on Parallel Auctions

Alexander Kleiner, Bernhard Nebel and Vittorio Amos Ziparo

University of Freiburg, Freiburg, Germany, {kleiner,nebel}@informatik.uni-freiburg.de
Algorithmica Srl, Rome, Italy, ziparo@algorithmica.it *

Abstract

Car pollution is one of the major causes of greenhouse emissions, and traffic congestion is rapidly becoming a social plague. Dynamic Ride Sharing (DRS) systems have the potential to mitigate this problem by computing plans for car drivers, e.g. commuters, allowing them to share their rides. Existing efforts in DRS are suffering from the problem that participants are abandoning the system after repeatedly failing to get a shared ride. In this paper we present an incentive compatible DRS solution based on auctions. While existing DRS systems are mainly focusing on fixed assignments that minimize the totally travelled distance, the presented approach is adaptive to individual preferences of the participants. Furthermore, our system allows to tradeoff the minimization of Vehicle Kilometers Travelled (VKT) with the overall probability of successful ride-shares, which is an important feature when bootstrapping the system. To the best of our knowledge, we are the first to present a DRS solution based on auctions using a sealed-bid second price scheme.

1 Introduction

Road transportation is one of the major challenges for the new millennium, as the current transportation models are both not environmentally and socially sustainable. On the one hand, car pollution is a major source of greenhouse emissions (see for example [European Commission, 2010]) that, in turn, is a cause of global climate change. On the other hand, traffic congestions are a world-wide problem and are believed to yield to a reduction of people’s quality of life [European Conference of Ministers of Transport, 1999]. Apparently, there is plenty of room for improvements, for example, when considering the fact that on average cars carry 1.6 passengers, i.e., only an estimated 25% of emissions are caused by people traveling, and the rest by moving “empty seats” [Jeanes, 2010]. Many researchers, spanning over multiple disciplines

such as transportation [Morency, 2007], economics [Brownstone and Golob, 1992], and behavioral, social and environmental psychology [Nordlund and Garvill, 2003], have identified ride sharing as a solution to the inefficiency of current transportation models.

To date, DRS systems (e.g., www.avego.com and www.participate.com) are not capable of retaining a critical mass of users. The experience from previous projects (see dynamicridesharing.org for a discussion on the topic) highlights that major drawbacks of existing systems are related to *flexibility in usage*, e.g., allowing people to decide with whom they share a ride, and to *high availability*, as, on average, if users do not get a ride within the first three attempts, they will not return. In order to provide high availability of service, a DRS system must be able to provide adequate incentives to users and, in particular, to drivers. Kamar and Horvitz [Kamar and Horvitz, 2009] recently proposed the *ABC* DRS system that dynamically generates shared plans and computes fair payments using a VCG mechanism [Nisan *et al.*, 2007] as incentives for the users. However, besides fair payments (which are not incentive compatible) their system does not provide the required flexibility, as users have no influence on their assignments.

In this paper we present a DRS solution based on auctions with a second-price scheme. In contrary to existing solutions, the proposed mechanism facilitates both individual user preferences for ride sharing partners and personal valuations for specific rides. In colloquial terms, passengers are bidding for increasing their ranking, and thus visibility to drivers, whereas drivers can select passengers according to their preferences. From the perspective of a single driver, potential passengers can require different detours for being picked-up.

One important feature of our system is that it allows to tradeoff the minimization of Vehicle Kilometers Travelled (VKT), i.e. greenhouse emissions, with the overall probability of successful ride-shares. On the one hand side, the auction mechanism automatically increases the chances of passengers bidding with higher valuation since their bid might cover the detour costs of several drivers. On the other hand side, longer detours are increasing the overall VKTs. We show empirically that the variance in valuation and thus the potential detour distances impacts both the probability for ride-shares and the amount of saved VKTs. In fact, we show that the lower the variance the closer the system performs to-

*The authors are grateful to Robert Mattmüller for the extensive discussions and useful suggestions.

wards the optimal assignment in terms of VKT savings. This opens the door for regulating the tradeoff, which is an important feature particularly for bootstrapping the system. For example, permitted detour distances can be kept unbounded at the beginning and then stepwise be decreased during the life-time of the system in order to minimize VKTs.

The presented approach is related to the generalized second-prize auctions (GSP) [Edelman *et al.*, 2007] (also known as Position Auctions [Varian, 2007]), used by Yahoo! and Google to assign advertisement slots within web search results [Battelle, 2005]. In general, it is very difficult to obtain both efficiency and incentive compatibility in a GSP, as for example the aforementioned ad auctions are not incentive compatible. We show that our approach is incentive compatible for passenger bids.

The paper is structured as follows. First, in Section 2, we provide a brief overview of our DRS system. Then, in Section 3, we provide a formal characterization of our system and we show that it is incentive compatible. In Section 3.3, we describe our approach on dealing with the dynamic nature of the problem. In Section 4, we describe an optimal, yet strict, assignment method that we use as a baseline for the evaluation of our system. In Section 5, we provide empirical results, and in Section 6, we conclude with some discussion.

2 System Overview

The target system of our DRS approach is a smartphone application for dynamic ride-sharing that operates in real-time, does not require pre-declared paths or pickup points, and can be used from virtually any location at any time. The system promotes ride-shares among people with connections within social networks (e.g., *Facebook*) in order to favor ride-shares among users that trust each other. From a passenger’s perspective, the DRS system is like an “eBay of rides”. When the application is launched on a smartphone, the user is asked to enter a goal location, a deadline, and to place a bid. Then, the system determines drivers that are online and compatible for this ride based on the length of the detour to be performed, and whether deadlines will comply. The passenger can then select from the list of compatible drivers, which are presented with additional information such as user ratings and social network status. Finally, these drivers will receive the bid, which is considered as a binding commitment of the passenger, i.e., to accept the first driver that accepts. After the ride, the passenger has the possibility to review the driver, for example, by a rating (zero to five stars), and optionally a comment.

From a driver’s perspective, the DRS system may be thought as a GPS car navigation system connected to a social network. The driver, just before starting to drive, declares a goal location, and a deadline. Then, the system computes the shortest path and gives directions to the driver for reaching the goal. During the ride the driver may be informed, e.g. vocally, about ride sharing offers from passengers nearby to his route. For example, the system might announce that there are three bids and the best one offers a profit of €10 for a detour of 15 minutes (computed by the system). The system may also announce that there is a bidding passenger that

is a remote friend according to a social network graph, e.g. *Facebook*, and that other drivers rated him on average with four stars out of five. The driver may accept the first bid or request information on the next bid. Note that the cognitive load of the driver can be reduced by a set of individually adjustable preferences that filter-out uninteresting offers beforehand. When the driver accepts an offer, the application changes the route according to the optimal joint route of both. At the end of the trip, the system automatically performs the payment and the driver has also the possibility to review the passenger.

The implementation of the system is still ongoing work. For our feasibility study we implemented a server platform than can simulate offers and demands from drivers and passengers based on the road network of any city in the world. Furthermore, we implemented a client application running on the *Android* mobile phone architecture for accessing this data from the road.

3 Multi-Agent DRS System

The DRS system may be thought of as a multi-agent system composed of selfish agents that dynamically enter and exit the system. Agents can enter the system either as passengers or as drivers, and exit when they share a ride, reach their goal locations, or meet their deadline. For the purpose of this paper, we assume that only one passenger can be assigned to (i.e., share a ride with) one driver. The proposed system can generally be extended for assigning multiple passengers to a single driver. However, then the computational complexity for computing all possible roundtrips (in the worst case $(n-1)!/2$) increases, which necessitates the deployment of heuristic algorithms, such as for example ant colony optimization [Dorigo and Gambardella, 1997].

The goal of our system is twofold. On the one hand, to minimize the greenhouse emissions by minimizing VKTs and on the other hand, to increase the number of ride-matches by providing an incentive compatible mechanism.

3.1 Agents

We denote by \mathcal{A}_t the set of agents that are in the system at time t . Each agent $a \in \mathcal{A}_t$ has associated a type $\langle s_{(a,t)}, g_{(a,t)}, dl_{(a,t)} \rangle$, where $s_{(a,t)}$, $g_{(a,t)}$ and $dl_{(a,t)}$ are, respectively, the current position, the goal position and the deadline for reaching the goal of a at time t .

The set \mathcal{A}_t of agents in the system at time t is partitioned in two sets:

- the set of drivers \mathcal{D}_t at time t , i.e. the set of people that have a car and are willing to drive it;
- the set of passengers \mathcal{P}_t at time t , i.e. the set of people that either do not have a car, or they have it, but are not willing to drive it.

Figure 1 shows the single paths of the driver and passenger if driving alone (left side) and the path required if they decide to share a ride (right side). It can easily be shown that in some cases the joint route is significantly shorter than the sum of both. For example, in case that both have the same goal at a far distance and their starting locations are located in the same neighborhood.




Figure 1: Single routes of passenger (green) and driver (blue), overall path required for ride share (orange).

The estimated duration of a ride share is, in general, different for the passenger and the driver as the driver has to travel an extra path. We define these durations as follows:

- $duration_p(d, p)$, i.e., the time required to drive by car the shortest path going through $s_{(d,t)}$, $s_{(p,t)}$ and $g_{(p,t)}$;
- $duration_d(d, p)$, i.e., the time required to drive by car the shortest path going through $s_{(d,t)}$, $s_{(p,t)}$, $g_{(p,t)}$ and $g_{(d,t)}$.

In order for a passenger and a driver to share a ride, their deadlines must be compatible with respect to the time required to drive to their goal locations. For each point in time t and for each passenger $p \in \mathcal{P}_t$, we define the set $\mathcal{C}_{(p,t)}$ composed of all drivers $d_i \in \mathcal{D}_t$ that are compatible with p , i.e. $(duration_p(d, p) < dl_{(p,t)} - t) \wedge (duration_d(d, p) < dl_{(d,t)} - t)$.

Let us define c_{km}^d as the cost per km of a driver. Moreover, we define $len(p, p')$ as the length of the shortest path that leads from position p to position p' by car in km . For convenience of notation, we define the length of the path for the passenger $lp(p) = len(s_{(p,t)}, g_{(p,t)})$, the length of the original path of the driver $ld(d) = len(s_{(d,t)}, g_{(d,t)})$, and the overall path of the ride share $lt(d, p) = len(s_{(d,t)}, s_{(p,t)}) + len(s_{(p,t)}, g_{(p,t)}) + len(g_{(p,t)}, g_{(d,t)})$.

A ride share $\langle d, p \rangle_t$ is the assignment to a passenger $p \in \mathcal{P}_t$ of a driver $d \in \mathcal{C}_{(p,t)}$ at time t . The valuation of the ride share $\langle d, p \rangle_t$ is

$$v_d(\langle d, p \rangle_t) = (ld(d) - lt(d, p)) * c_{km}^d,$$

for driver d , while the passenger p values a ride from $s_{(p,t)}$ to $g_{(p,t)}$ with

$$v_p(\langle d, p \rangle_t) \geq lp(p) * c_{km}^d.$$

Notice that the valuation of the ride share for a driver is always not positive, as he may usually have to drive a longer distance to fetch the passenger and deliver him to his goal location.

3.2 Auction Mechanism

The auction mechanism is used to provide a ranking for each driver of the bids received from passengers. Agents when entering the system announce their types. Additionally, passengers announce a bid b_p for their rides, which they are willing to pay. In the end, the mechanism ensures that passenger p will pay an amount $pay_p \leq b_p$. From specific user preferences and a social distance measure derivable from *Facebook* or similar means, we assume a function $val_d(p)$ returning the

value driver d assigns to passenger p . Based on that, the utility for a ride share $\langle d, p \rangle_t$ for the passenger is

$$u_p(\langle d, p \rangle_t, b_p) = v_p(\langle d, p \rangle_t) - pay_p$$

and for the driver is

$$u_d(\langle d, p \rangle_t, b_p) = v_d(\langle d, p \rangle_t) + pay_p + val_d(p).$$

For each passenger p , the system computes the target drivers to whom the bid of p is forwarded:

$$\mathcal{T}_{(p,t)} = \{d_i \in \mathcal{C}_{(p,t)} \text{ s.t. } u_d(\langle d, p \rangle_t, b_p) > 0\}.$$

The bids are ranked for each driver based on the utility functions $u_d(\langle d, p \rangle_t, b_p)$. If the value $val_d(p)$ has captured all non-monetary preferences of the driver, the system could just select the passenger on the first rank as the winner of the auction. Notice that translating preferences into a single-dimension utility value is quite common in economics. Of course, one could also use multi-criteria preferences with the effect that outcomes are not comparable.

Alternatively, we consider the case that the final decision on who gets the ride is made by the driver by selecting a passenger from the list. In an auction, the bids of the passenger are ordered by the utility of the driver, provided the passengers would actually pay their bid: $b_{p_1}, b_{p_2}, \dots, b_{p_n}$. The passenger that wins an auction, has to pay according to a second price scheme. The winner p_i with bid b_{p_i} has to pay to the driver pay_{p_i} , which is the sum of

- $-v_d(\langle d, p_i \rangle_t)$ i.e., the detour cost;
- $v_d(\langle d, p_{i+1} \rangle_t) + b_{p_{i+1}} + val_d(p_{i+1})$ i.e., the utility of the ride share with passenger p_{i+1} , provided that p_{i+1} is the next passenger in the ranking of the auction and that p_{i+1} would have paid the bid $b_{p_{i+1}}$;
- $-val_d(p_i)$ i.e. the negative valuation of the presence of p_i for the driver.

Notice that negative payments are not possible. In order to be an eligible passenger (see the definition of the set $\mathcal{T}(p, t)$), the utility has to be positive. In particular, even if we would have only the one eligible passenger, he would have to pay the price the he bid - as is the usual way in 2-nd price auctions when there is only one bidder.

3.3 Computation of Dynamic Assignments

In the dynamic ride sharing setting, announcements of passengers (e.g. bids) and drivers are arriving continuously at arbitrary times. Hence, matches between passengers and drivers have to be computed many times during the day which implies that future requests are unknown during each planning cycle. Figure 2 depicts this problem from the perspective of a single driver that receives announcements (circles in the figure) with deadlines (diamonds in the figure) from passengers $\{p_1, \dots, p_5\}$. Regardless of the assignment mechanism, the problem is to decide an appropriate point in time for computing and committing the assignment between drivers and passengers. By *computing* we specifically mean to decide the best assignment given the current knowledge, e.g., to compute the current ranking of bids for each driver. By *committing* we mean to let the assignment come into effect, either

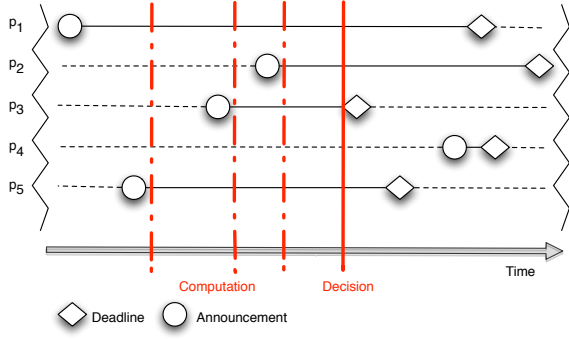


Figure 2: Solving the dynamic ride sharing problem by a rolling horizon decision making approach.

automatically or by presenting a ranked list of potential passengers to the drivers, which then finally commit a decision by selecting one or none of the passengers on the list.

A common method to solve problems of this kind is the deterministic rolling horizon approach [Sethi and Sorger, 1991], by which assignments are computed using all known information within a planning horizon, but decisions are only committed when necessitated by a deadline. As depicted by the dotted vertical lines in Figure 2, computations are triggered when the planning horizon grows due to deadlines from new passenger announcements. As shown by the bold vertical line in Figure 2, a decision is made shortly before the deadline of passenger 3 expires. Notice that deadlines always expire before passengers and drivers change from being compatible to being incompatible due to $(duration_p(d, p) > dl_p - t) \vee (duration_d(d, p) > dl_d - t)$.

3.4 Incentive Compatibility

One desirable property of such a system is that it is incentive compatible, i.e. that it is rational for each user to state his type and preferences truthfully. Let us consider the driver's type first.

The starting position $s_{(a,t)}$ is something the driver hardly can cheat about because it is known by the system. The deadline and the goal location may be something the driver could lie about. However, stating an earlier deadline only limits the number of possible passengers and stating a later deadline might lead to a late arrival. Lying about the goal could potentially change the payment to the driver. However, since the payment is computed automatically, this computation could be done when the driver has finished his trip. If his goal location has changed, the auction payment could be re-evaluated. So, lying about the goal would not help.

Using the perspective of the passenger, the type is again something one hardly would be untruthful about. It does not make sense to lie about starting and goal location. Similarly, the deadline should be stated truthfully because otherwise one might limit the number of drivers or risks of being too late. This is true, at least, if we consider one such auction. However, if we look at the entire system taking time into account, then stating an earlier deadline could help because the end of the auction is determined by the deadline (see Section 3.3).

If stating an earlier deadline causes a passenger who comes later to be excluded, one could win an auction by lying. Of course, this seems to be unavoidable. However, one could at least punish the untruthful statement in cases when a passenger loses an auction by disallowing a similar ride with a later deadline in the next hour by the system.

Concerning the bid, the modified second-prize scheme guarantees incentive compatibility in case of the automatic assignment. Neither a higher nor a lower bid would change the payment in case of winning the auction - provided one bids more than the second bidder. For all losing bidders, a lower bid would not help and a higher bid could hurt. So the statement of the true valuation is a dominant strategy.

Now let us consider the case that the driver selects one passenger from the list. Assuming that with lower ranks there is a strictly decreasing probability that the driver selects a passenger, the expected utility could be defined as the probability of being chosen times the associated utility. Then it looks as if we end up with the generalized second-prize auction (GSP) [Edelman *et al.*, 2007]. Note, however, that in our kind of auction there is only one winner who has to pay the full price, while in a GSP everybody wins something and also has to pay something. So, the process of choosing somebody not on the first rank must either be viewed as a non-rational action or as a process where non-stated preferences play a role. Here a passenger, who has all information, could clearly gain something from placing bids that are below his true valuation.

4 Optimal Assignment

In this paper, the optimal assignment will be used as a baseline for evaluating the proposed auction mechanism. When computing the optimal assignment, the goal is to maximize the total amount of VKT savings in order to reduce the overall travel costs and greenhouse emission. This is a common approach in existing ride sharing systems, e.g., for the ride sharing of commuters [Kamar and Horvitz, 2009].

The optimal assignment between passengers and drivers can be computed by finding the maximum weighted bipartite matching. We consider a matching as feasible only when the travel distance of the joint path of passenger and driver is smaller than the sum of the travel distances of both individual paths, i.e., when the matching yields a reduction of the total travel distance. Therefore, we have to extend the cost matrix by the cases where passengers and drivers travel on their own (assuming they can use a Taxi or their own car).

Equation 1 shows the extended cost matrix where $c(d_i, p_j) = lt(d_i, p_j)/2$ denotes the half joint cost for the matching of driver d_i and passenger p_j , $c(d_i) = ld(d_i)$ the single cost of driver d_i , and $c(p_j) = lp(p_j)$ the single cost of passenger p_j .

$$\begin{matrix}
 & \mathbf{d}_1 & \dots & \mathbf{d}_m & \mathbf{p}_1 & \dots & \mathbf{p}_n \\
 \mathbf{p}_1 & c(d_1, p_1) & \dots & c(d_m, p_1) & c(p_1) & \dots & \infty \\
 \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\
 \mathbf{p}_n & c(d_1, p_n) & \dots & c(d_m, p_n) & \infty & \dots & c(p_n) \\
 \mathbf{d}_1 & c(d_1) & \infty & \infty & c(d_1, p_1) & \dots & c(d_1, p_n) \\
 \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\
 \mathbf{d}_m & \infty & \infty & c(d_m) & c(d_m, p_1) & \dots & c(d_m, p_n)
 \end{matrix} \quad (1)$$

From the cost matrix in Equation 1 the optimal assignment can efficiently be computed in $O(N^3)$ by a modified version of the *Hungarian algorithm* [Kuhn, 1955], where N is the maximal number of rows and columns of the cost matrix.

5 Experimental Results

We developed an experimental platform to validate the hypothesis that an auction based system leads to a higher probability of ride-shares between passengers and drivers when compared to methods optimizing on the VKTs only. We implemented the optimal assignment by using the Hungarian method with the cost matrix shown in Section 4 and the auction approach according to the description in Section 3.2.

5.1 Test Platform

The experimental platform simulates the behavior of drivers and passengers, each having his own (randomly generated) preference over closing a deal. Planned routes of passenger and drivers are randomly sampled from a uniform distribution over city maps by generating for each passenger $p \in \mathcal{P}_t$ and driver $d \in \mathcal{D}_t$ the tuples $(s_{(p,t)}, g_{(p,t)})$ and $(s_{(d,t)}, g_{(d,t)})$, respectively. Uniform sampling can be considered as a conservative approach since in a real setting the locations would be much more correlated, e.g., clustered due to commercial areas, industrial areas, and shopping centers.

For sake of simplicity we assume that all drivers and passengers have compatible deadlines, i.e., that any driver can deliver any passenger in time. Note that for both the optimal approach and the auction approach the incompatibility between a driver d_i and passenger p_j can easily be expressed by setting the joint costs $lt(d_i, p_j) = \infty$.

We utilized OpenStreetMaps (openstreetmap.org) data together with the *MoNav* planner that offers exact routing without heuristic assumptions at very little computational demand due to a routing core based on contraction hierarchies [Vetter, 2010]. The planner is used at each time step to compute the joint cost $lt(d_i, p_j)$ for the matching of driver d_i and passenger p_j , and the single costs $ld(p_j)$ and $ld(d_i)$ of driver d_i and passenger p_j , respectively, for each $p_j \in \mathcal{P}$ and $d_i \in \mathcal{D}$. As a testbed for our experiments we used the German cities of Freiburg and Berlin with a size for sampling locations of $22km^2$ and $170km^2$, respectively. Figure 3 depicts an example output of our system.

As performance metrics we counted during each run the frequencies of successful ride agreements and the savings of VKTs by comparing the joint kilometers against the kilometers driven when driver and passenger are traveling on their own. Experiments were conducted with varying ratios among drivers and passengers. Results are averaged over 100 runs for each combination of number of drivers (1-50) and number of passengers (1-50).

5.2 Optimal Assignment

In this experiment we evaluated the probability of matches between drivers and passengers when applying the optimal assignment on the map of Freiburg. In the optimal assignment matches are only possible when the cost of the joint route of driver and passenger is below or equal to the sum of both routes travelled individually.



(a)



(b)

Figure 3: Output of our experimental platform displayed on Google maps: (a) the single trips of two passengers and drivers, (b) their shared rides computed by the auction mechanism.

A commercial DRS based on the optimal assignment could be implemented so that each participant pays a monthly base fee in advance. Then, these fees could be used for supporting individual ride matches, i.e., lowering the joint cost artificially by covering extra detour costs of the drivers. Figure 4 summarizes the matching probability from conducted experiments with varying ratios between the number of passengers and the number of drivers. Mean and variance are presented at different degrees of system support, i.e., the maxi-




Figure 4: Average probability of ride matches during optimal assignment with increasing system support tested on the city map of Freiburg, Germany.

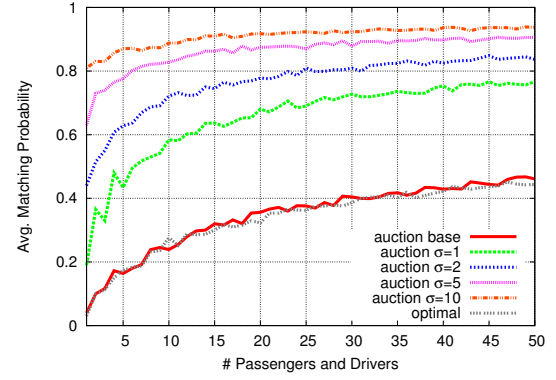
mal amount of detour kilometers that would be paid by the system for enabling the match.

These results show clearly that without financial support from the outside no satisfying matching probability can be achieved which in turn can increase the number of people abandoning the system. Adding system support increases the number of matchings, however, it is unclear in a practical setting how to weight it for specific driver-passenger pairings, and how to define the base fee. Results presented in the following section are related to the auction system that in contrary automatically adjusts to individual user valuations.

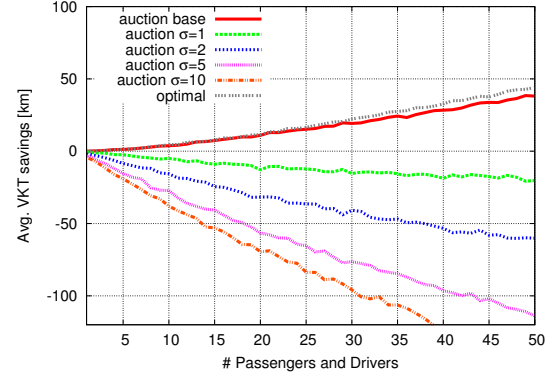
5.3 Optimal vs Auction Based Assignment

In this section we compare the auction-based mechanism with the optimal assignment. In the auction setting we simulated the individual valuations of passengers by random sampling from a half-Gaussian distribution over the interval $lp(p) * c_{km}^{taxi} \geq v_d(d, p)_t \geq lp(p) * c_{km}$, where c_{km}^{taxi} denotes the taxi price per km (€3.1) and c_{km} the base cost per km for a private driver (€0.3). The half-Gaussian distribution was chosen to model that with high probability passengers are willing to pay an amount close to the minimal costs for a single ride and with low probability an amount that is close to the price of the same ride with a taxi. For each randomly generated set \mathcal{D}_t and \mathcal{P}_t a sequence of auctions was initiated with one for each driver $d \in \mathcal{D}_t$. Passengers $p \in \mathcal{P}_t$ participated in every compatible auction and the results of the auctions were computed sequentially. Figure 5 and Figure 6 depict the VKT savings and the probabilities for matching rides on the Freiburg and Berlin map, respectively. We compared the optimal assignment and the auction method using different variances σ for the half-Gaussian distribution.

The results clearly show that with increasing variance, i.e. higher willingness of the passengers to pay more than the base costs, the probability of matchings increases. This is obvious since higher investments have more potential for covering longer detours. However, with increasing variance also the amount of VKTs savings decreases since longer detours are taken by the drivers. Interestingly, when limiting the maximal valuation of passengers to zero, i.e. bids are exclusively



(a)



(b)

Figure 5: Experiments on the Freiburg city map with increasing number of drivers and passengers: (a) The percentage of matchings. (b) the savings of total VKT (Vehicle Kilometers Travelled).

the base costs (denoted by *auction base* in the figures), nearly the same amount of VKT savings are reached as with the optimal assignment. Therefore, our approach allows to trade-off near to optimal VKT savings for higher system availability by limiting the amount of allowed detour costs, as needed. This can easily be implemented by removing drivers from the set of compatible drivers in Section 3.2 if their detours are above the limit.

6 Discussion

In this paper we have presented a novel DRS system that: 1) allows more flexibility than state of the art systems, as it gives more freedom to users in the choice of their ride-sharing partner; 2) is robust in the sense that it promotes truth-telling about most aspects and fairness in payments; and 3) allows to trade-off VKT savings with the probability in finding matches between drivers and passengers.

Moreover, despite giving freedom to users to deviate from the optimal solution, experimental results show that the performance of our system is very close to optimal. In particular, it emerges that, depending on the maximum detour we allow in the system, we can tradeoff the probability of having a matching (i.e. a ride-share) versus minimization of VKT.




Figure 6: Experiments on the Berlin city map with increasing number of drivers and passengers: (a) The percentage of matchings. (b) the savings of total VKT (Vehicle Kilometers Travelled).

This result can be exploited to tune the system in order to achieve a critical mass, and thus high availability of service. For example, when the system is first launched one could allow for higher detours in order to maximize the matches and, thus, promote the creation of a critical mass of users. Then, detours would be limited as a long term solution in order to minimize VKT.

Current limitations of our system are the restriction to dual ride shares, i.e., one passenger and one driver, as well as the lack of a large-scale experiment proving the performance of our system in practice. In the near future, we plan to develop realistic simulations based on existing traces of user behaviors and field experiments with real users, in order to better understand the impact of limiting detours with respect to the achievement of critical mass. Moreover, we plan to address the issues related to the creation of ride-shares composed of more than two users and the impact of heuristic methods for computing the round-trips.

References

[Battelle, 2005] John Battelle. *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture*. Portfolio Hardcover, September 2005.

[Brownstone and Golob, 1992] David Brownstone and Thomas F. Golob. The effectiveness of ridesharing incentives: Discrete-choice models of commuting in southern california. *Regional Science and Urban Economics*, 22(1):5–24, 1992.

[Dorigo and Gambardella, 1997] M. Dorigo and L. M. Gambardella. Ant colony system: A cooperative learning approach to the traveling salesman problem. *IEEE Transactions on Evolutionary Computation*, 1:53–66, 1997.

[Edelman et al., 2007] Benjamin Edelman, Michael Ostrovsky, , and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.

[European Commission, 2010] Energy European Commission. *Statistical pocketbook 2010, EU energy and transport in figures*. Publications Office of the European Union, 2010.

[European Conference of Ministers of Transport, 1999] European Conference of Ministers of Transport. *Traffic Congestion in Europe*. ECMT Round Tables. OECD Publishing, 1999.

[Jeanes, 2010] Phil Jeanes. Can total car emissions be reduced enough by 2020? *Traffic Engineering & Control*, PA Associates, January 2010.

[Kamar and Horvitz, 2009] Ece Kamar and Eric Horvitz. Collaboration and shared plans in the open world: studies of ridesharing. In *Proceedings of the 21st international joint conference on Artificial intelligence*, pages 187–194, San Francisco, CA, USA, 2009. Morgan Kaufmann Publishers Inc.

[Kuhn, 1955] H. W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.

[Morency, 2007] Catherine Morency. The ambivalence of ridesharing. *Transportation*, 34(2):239–253, 2007.

[Nisan et al., 2007] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. The vickrey-clarke-groves approach. In *Algorithmic Game Theory*, pages 559–563. Cambridge University Press, 2007.

[Nordlund and Garvill, 2003] A. M. Nordlund and J. Garvill. Effects of values, problem awareness, and personal norm on willingness to reduce personal car use. *Journal of Environmental Psychology*, 23(4):339–347, December 2003.

[Sethi and Sorger, 1991] Suresh Sethi and Gerhard Sorger. A theory of rolling horizon decision making. *Annals of Operations Research*, 29:387–415, 1991.

[Varian, 2007] H. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, December 2007.

[Vetter, 2010] Christian Vetter. Fast and Exact Mobile Navigation with OpenStreetMap Data. Diploma thesis, Faculty of Computer Science Institute for Theoretical Computer Science, Algorithmics II, University of Karlsruhe, Germany, 2010.