A Marketplace for Spatio-temporal Resources and Truthfulness of its Users

Ouri Wolfson
Department of Computer Science
University of Illinois at Chicago
Chicago, U.S.A
Tel: 1-312-996-6770
Email: wolfson@uic.edu

Jane Lin
Department of Civil Engineering
University of Illinois at Chicago
Chicago, U.S.A
Tel: 312-996-3068
Email: janelin@uic.edu

ABSTRACT
In this paper we propose a platform called Spatio-Temporal rEsources Marketplace (STEM) for users of spatio-temporal transportation resources such as electric vehicle charging stations or parking spaces. STEM exploits the gap between optimum and equilibrium in game theory in order to improve the position of users in competing for such resources, as well as improve social welfare. It does so via a proposed payment scheme called Guaranteed-Agent-Gain (GAG). In this scheme users pay STEM, but if STEM refunds its profit to the users, then the scheme amounts to peer-to-peer transactions. We prove that in many cases GAG does not need to be subsidized in order to deliver its gains to travelers and society.

We show that GAG raises the issue of user truthfulness, in the sense that users can gain from being untruthful about their private information, such as their current location. Thus we also analyze the application of Vickrey-Clarke-Groves (VCG) mechanisms studied in economics to the problem of location truthfulness in STEM. We show that these mechanisms can indeed induce location truthfulness, but at the cost of reduced revenue.

General Terms
Algorithms, Performance, Economics

Keywords
Optimum, equilibrium, parking, electric-vehicles, matching, resources

1. INTRODUCTION
On average, people traveling during morning and evening rush hours in urban areas experienced 34 hours of delay annually in 2010; in urban areas with population over 3 million that delay goes up to 52 hours [1]. Between 1990 and 2007, Vehicle-Miles-Travel (VMT) grew 41% [2], and the projected total VMT in 2050 is 4,834 billion miles, an increase of 60% over 2007. Congestion is related to scarcity of resources, such as road space and parking. In one business district of Los Angeles, researchers found that vehicles searching for parking traveled a distance equivalent to 38 trips around world, produced 730 tons of carbon dioxide, and burned 47,000 gallons of gasoline in one year [3].

IWCTS'14, November 4, 2014, Dallas, TX, USA.

In this paper we propose to alleviate these problems by bridging the gap between equilibrium and optimum states. More specifically, it is well known that equilibrium is a stable state for transportation systems. Intuitively, this means that the system is settling into a state in which no user can unilaterally improve her performance. Unfortunately, the equilibrium state is often much worse overall than the optimum state, i.e. the state that maximizes overall social welfare. Moreover, the gap between equilibrium and optimum is potentially huge. Specifics depend on the particular system. But, for example, we have shown that for parking the gap between optimum and equilibrium is unbounded in the worst case [4], and is about 20% on average [5]. Imagine the potential of reducing travel-time by 20%!

We propose here to move from equilibrium towards optimum by peer-to-peer (P2P) financial transactions that will guarantee that every user will not be worse off than in equilibrium; and often, as will be demonstrated in the paper, each user will be better off. These transactions will be executed by a software app on the mobile devices called *Spatio-Temporal rEsources Marketplace (STEM)*. STEM may be supported by some central, i.e. cloud-based, infrastructure, but it can also be implemented in a completely distributed fashion.

We introduce STEM in the context of spatio-temporal transportation resources. These are resources in geo-space, or time-intervals, that may be available or unavailable to a user depending on other users. Examples of such resources include parking slots, taxi-cab customers, share-bikes, share-bike racks, ride- and car-shares, electric vehicles charging stations, and landing slots at an airport.

We introduce a payment scheme called Guaranteed-Agent-Gain (GAG), and demonstrate a subset of configurations where STEM transactions are guaranteed to leave every user, and society overall, in a better off situation than in equilibrium (Theorem 0). These are configurations where the cost of using a resource is the driving time to the resource, and the value of time is the same for all users (imagine a resident searching for parking in a neighborhood with a homogeneous socio-economic class).

A drawback of the proposed scheme is that it is vulnerable to strategic manipulation, e.g., users lying about their location in order to gain an advantage. We discuss several approaches to address this problem, and elaborate on one of them, namely Vickrey-Clarke-Groves (VCG) mechanisms for truthful auctions. We show that, in a model adapted from VCG where users pay for resources based on their declared distance from them, and slightly different than the GAG model, payments can be defined in a way that incentivizes truth-telling.

The rest of this paper is organized as follows. In section 2 we introduce the model for spatio-temporal resources, and in section 3 we introduce, discuss, and demonstrate the GAG payment scheme. We also prove that for a subset of configurations GAG transactions are guaranteed to leave every user and society in a better situation than in equilibrium. In section 4 we discuss truthfulness, and in section 5 we show intuitively that the concepts carry over to routing. In section 6 we discuss relevant work, and in section 7 we provide the conclusion and discuss future work.

2. THE MODEL

A configuration consists of a set of mobile agents (e.g. vehicles) V={1,2,...,n} in geo-space and a set of static available resources (e.g. parking slots) R={1,2,...,m}. Agents do not have an extent, thus at any time they are located at points in geo-space. The resources are either all spatial or all temporal. If spatial, then they are static and occupy point locations in geo-space. Temporal resources are time intervals. For example, landing time-slots at an airport-runway are temporal resources. A resource can be used by a single agent at a time, and the agent has to reach the resource in order to use it. An assignment is a mapping from agents to resources.

At the outset all agents arrive simultaneously at various starting locations. At that time, for each agent *i* (*i*∈V) there is a driving-time, D_{ij}, to reach resource *j* (*j*∈R). Additionally, there is a natural cost C_{ij} associated with agent *i* using resource *j*. This natural cost may, for example in the case of parking, be simply the driving-time to the slot, or, a weighted average of (the driving time to the slot) + (the walking time from the slot to the driver's final destination).

The locations of the resources are known to a software platform called the Spatio-Temporal rEsources Marketplace (STEM). STEM may reside centrally in the cloud, or a copy of it may reside in a tamper-resistant fashion in the mobile device of each agent.

At the outset, each agent *i* sends STEM the information necessary to compute *i*'s driving-time to each resource and *i*'s natural cost for each resource. For example, in the case of parking, if the natural cost is a weighted average of (the driving time to the slot) + (the walking time from the slot to the driver's final destination), then each agent sends STEM its starting location and its driver's final destination.

When receiving this information from all the agents, STEM computes all D_{ij}'s and C_{ij}'s, and then it computes an assignment and a payment scheme. The assignment **M** will minimize the total driving time. In general, such an assignment is optimal from a social-good point of view, but the problem is that some agents may incur a higher natural cost than necessary from their individual stand point of view.

For example, consider the configuration in Figure 1 and assume that the costs of using the resources are simply the driving-times in an arbitrary road-network (thus the triangle inequality does not need to be satisfied). Then the minimum driving-time assignment **M** is **M**={(v_1, s_2), (v_2, s_1)} with a total cost of 7. But v_1 will not collaborate since it can reach s_1 before v_2 and incur a cost of just 1, lower by 1 than in **M.** In other words, in a free-for-all competition in which agents do not use the services of STEM the system will converge to an equilibrium assignment **E** [6], where, intuitively, **E** is an assignment in which no agent can unilaterally improve its cost. In this example, **E**={(v_1, s_1), (v_2, s_2)} with a total cost of 9.

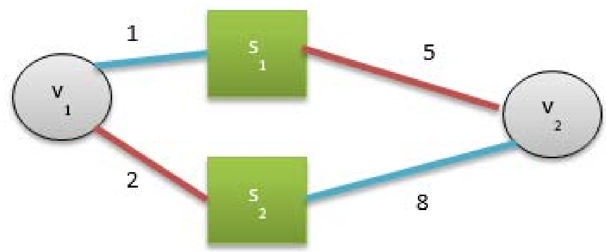

Figure 1: A configuration of vehicles (agents v_1 and v_2) and parking slots (resources s_1 and s_2), where the edge-labels denote driving-times in the road network.

Now we formally define the minimum and equilibrium assignments. For this purpose we assume in an assignment, if multiple agents are mapped to the same resource *r*, then the one having the shortest driving time to *r* gets it; and the remaining ones are left *unallocated*. An unallocated agent *v* incurs a very high natural cost, *α*. For example *α* may be higher than the sum of all the pairs of natural costs. Intuitively, it means that there is a high penalty for an agent not assigned a resource that they can capture. Obviously, if the number of resources is lower than the number of agents, any assignment will have some agents that are unallocated. The *cost* of an agent *v* (*v*∈V) in an assignment **A**, denoted *C*(*v*, **A**), is the natural cost of *v* in using the resource assigned to *v* in **A**. Suppose that *v* is assigned resource *r* (*r*∈R) in **A**, then *C*(*v*, **A**) = C_{vr}. If *v* is unallocated in **A**, then *C*(*v*, **A**) = *α*. The *cost of an assignment* **A** is $\sum_{\forall v \in V} C(v, \mathbf{A})$.

An assignment **M** is a *minimum-cost* assignment if any other assignment **B** has a cost that is not higher than that of **M**. In [4] we have shown that given a finite set of agents and a finite set of resources, a minimum-cost assignment can be computed in strongly polynomial-time by representing it as a minimum-cost network flow on a bipartite graph [7].

Formally, an assignment **E** is an *equilibrium* assignment if for every agent *j*, and for every other assignment **B** that differs from **E** only in the assignment of *j*, *C*(*j*,**E**) ≤ *C*(*j*,**B**). In [5] we have shown that an equilibrium assignment can also be found in polynomial time using the Gale-Shapley deferred acceptance algorithm.

An interesting point to observe is that if the number of agents is higher than the number of resources, i.e., n>m, then different sets of agents may be allocated a resource in **E** and **M**. In other words, different sets of agents may remain unallocated in the two assignments. To see this, consider Figure 2, which is just Figure 1 with one more vehicle, v_3, added. In this configuration, the minimum cost assignment is **M**={(v_1,s_2), (v_2,s_1)} with v_3 unallocated; and the equilibrium assignment is **E**={(v_1,s_1), (v_3,s_2)} with v_2 unallocated.

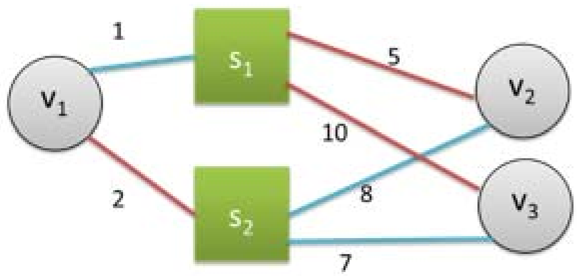

Figure 2: A configuration of vehicles and parking slots, where different subsets of the vehicles park in equilibrium (v_1 and v_3) and optimum (v_1 and v_2) assignments.

3. PAYMENTS FOR EQUILIBRIUM-TO-OPTIMUM CONVERSION

Since some agents may be better off in **E** than in **M**, they may not cooperate with STEM. This means that from each agent's viewpoint its necessary natural cost is the cost in an equilibrium assignment, and it will not cooperate if its cost in an assignment produced by STEM is higher than necessary.

To compensate for this situation and induce cooperation, STEM uses the following Guaranteed-Agent-Gain (GAG) scheme for converting equilibrium to optimum. GAG is a payment scheme that guarantees to each agent v that its overall cost in **M** will not be higher than its natural cost in **E**. This is how GAG works.

i. If for some agent v, $D_v = C(v,\mathbf{E}) - C(v,\mathbf{M})$ is negative, meaning the natural cost of v in **E** is smaller than in **M**, then STEM pays v an amount equal to $|D_v|$ in dollars. This is to compensate for the increase in v's natural cost by moving v from **E** to **M**.

ii. If D_v is positive, it means that v benefits by moving from **E** to **M**. Then v pays back STEM D_v in dollars and its overall cost in **M** is still no worse than that in **E**.

Thus, the GAG payment scheme guarantees that each agent v pays an adjusted cost, i.e. $C(v,\mathbf{M}) + D_v$, which is not higher than v's natural cost in equilibrium, $C(v,\mathbf{E})$.

Observe that the GAG payment scheme requires a \$-payment for a natural cost. Finding the \$-value of the natural cost is an implementation detail, but as an example, assume that an agent v's cost of using a parking slot is simply v's driving time to the slot. If so, the \$-payment for the natural cost is simply the v's value of time, e.g. \$1/minute. Of course, this valuation can be adjusted to account for gas expended in driving to the slot, and also for the time to walk from the slot to the final destination.

Assume now that STEM proceeds with the assignment **M** and the GAG payment scheme, i.e. it announces these to the agents, only when the *\$-income,* i.e. the sum of D_v's received from the agents in (ii), is not lower than the *\$-outcome*, i.e. the sum of D_v's paid out to the agents in (i). Otherwise STEM does not mediate the competition, and tells the agents to compete for the resources as they currently do, i.e. without the mechanism proposed in this paper.

Definition 1. A minimum-cost assignment M combined with the GAG payment scheme is called *viable* if for M and GAG the \$-income is not lower than the \$-outcome.

Intuitively, if a minimum cost assignment combined with the GAG payment scheme is viable, then STEM does not need to subsidize the mechanism. The next theorem shows that many natural situations (or configurations) are viable.

Theorem 0: If the natural costs are the driving times, and if the value of time is the same for all agents, then for every configuration of agents and resources, the minimum-cost assignment **M** combined with the GAG payment scheme is viable.

Proof: The proof follows from Theorem 4 in [5], and is based on the fact that the total system cost of the minimum-cost assignment **M** is not higher than the total system cost of the equilibrium assignment **E**. That is, the following inequality always holds:

$$\Sigma_v\, C(v,\mathbf{E}) \geq \Sigma_v\, C(v,\mathbf{M}) \tag{1}$$

Then

$$\Sigma_v\, [C(v,\mathbf{E}) - C(v,\mathbf{M})] \geq 0 \tag{2}$$

i.e., $$\Sigma_v\, D_v \geq 0 \tag{3}$$

We rewrite (3) into the following:

$$\Sigma_+ D_v + \Sigma_- D_v \geq 0 \tag{4}$$

where $\Sigma_+ D_v$ represents the summation of all positive D_v's and $\Sigma_- D_v$ all negative D_v's. Based on Definition 1, $\Sigma_+ D_v$ is simply the total \$-income of STEM from the agents and $\Sigma_- D_v$ is simply the total \$-outcome from STEM to the agents. (4) says the total \$-income is always no less than the total \$-outcome. Therefore, the GAG payment scheme combined with the minimum-cost assignment **M** is viable. []

Observe that for the purpose of computing the assignments and the payment scheme it does not matter whether STEM is implemented centrally in the cloud, or distributed on the mobile devices of the agents. If distributed, all mobile devices will receive the same information and compute the same assignment.

Further assume that some *refund scheme* is followed, i.e. a scheme that refunds STEM's profit to the agents. Distributing the profit evenly among the agents is one such a refund scheme. In this case, the distributed version of STEM simply assigns slots and implements peer-to-peer (P2P) financial transactions. For example, for the configuration of Figure 1, assume that for both vehicles (i.e. agents) the natural costs are the driving times, and that the value of time for each vehicle is \$1/minute. Then, in the GAG payment scheme v_2 pays STEM \$3 (i.e. D_{v2} in Table 1), and STEM pays v_1 \$1 (D_{v1}). The even-distribution refund-scheme means that the \$2 profit of STEM is distributed evenly between v_1 and v_2. Namely, v_2's payment to STEM is reduced by \$1, from \$3 to \$2; and v_1's payment received from STEM is increased by \$1 to \$2. The net effect of this is that \$2 is transferred in a P2P transaction from v_2 to v_1. As a result, the overall cost in GAG to both vehicles is equal to $C(v,\mathbf{M}) + D_v$ + Refund, which translates to \$0 for v1 and \$7 for v2. Both represent gains compared to \$1 and \$8 respectively for v1 and v2 in the equilibrium assignment. Furthermore, the equilibrium assignment, which would have occurred without STEM, was converted to an optimum assignment, a conversion which benefits society at large as well (less total driving time means less pollution, congestion, gas consumption, etc.).

Observe that although some central coordination occurs through STEM, due to the refund scheme STEM does not make a profit, i.e., all the money paid by agents is paid out to agents, and in this sense the transactions are P2P. However, observe that in the general case, these are not necessarily binary transactions. This means that the payment of one agent may be paid out to more than one other agent, or the combined payment of three agents may be paid out to four other agents.

Table 1. Overall cost to agents in GAG in Figure 1 with even refund.

Agent	$C(v,\mathbf{E})$ ($)	$C(v,\mathbf{M})$ ($)	D_v ($)	Even refund ($)	Total add'l cost ($)	Overall cost in GAG ($)
v_1	+1	+2	-1	-1	-2	0
v_2	+8	+5	+3	-1	+2	+7

4. TRUTHFULNESS

In the previous section we assumed that the agents provide STEM with the correct information to compute the D_{ij}'s and C_{ij}'s. However, observe that agents can gain from being untruthful. For example, consider Figure 1 again, suppose that the natural costs are the distances, and suppose that v_2 reports its location as being at distance 0 from s_1. If this were true, then the minimum assignment $\mathbf{M}$ = {(v_1,s_2), (v_2,s_1)} is also the equilibrium assignment. Therefore in the GAG payment scheme the $-income and the $-outcome are both 0. Furthermore, **M** is the minimum assignment in both, the truthful and the untruthful configurations. However, in the untruthful case v_1 is cheated out of a payment by v_2. Recall in Table 1 that v_1 would have been paid $2 in the truthful configuration, and is not paid anything in the untruthful one. This is an undesirable situation, since the fact that cheating is possible may deter agents from participating in the scheme.

One way to combat cheating is to give STEM direct access to GPS. This authorization may be combined with the authorization to impose fines if cheating is detected; and cheating can be detected by failure to arrive at the resource at the expected time. So, if v_2 reports that it is at driving-time 0 from s_1, but arrives 5 minutes later than expected, such a fine can be imposed; particularly since the GPS trace indicates that v_2 was not at the reported location. Furthermore, this method would not help if the natural cost includes accessing time to the final destination after parking.

An alternative way is based on Vickrey-Clarke-Groves (VCG) mechanisms that price the resources in a way that incentivizes truthful location-revelation. The adaptation of such a mechanism to the resource assignment problem addressed in this paper is described for the rest of this section.

Assume now that the location of each agent is only known to itself, and it may disclose to the STEM a false location. Assume further that for each agent *i* there is a value associated with using a resource *j* and that value is 1, and its natural cost for doing so is its driving-time to the resource, i.e., $C_{ij} = D_{ij}$. Then the D_{ij}-adjusted-value to *i* of using resource *j*, called DAV_{ij}, is $1\text{-}D_{ij}$. Namely, if $D_{ij} \geq 1$ then *j* is too far, and *i* would rather not use it, and DAV_{ij} is set to zero. In other words, *j* has no value to *i*. In this case we say that *j* is *infeasible* for *i*, otherwise it is *feasible*. In the case of temporal slots, a slot is infeasible if it is either too close to the present time and the airplane cannot get to it in time to land, or the slot is too far from the present time and will necessitate circling too long.

Observe that this model is slightly different than the one in the previous section, particularly in the assumption that the value of the resource is 1. We didn't make this assumption previously, but it is made here in order to adapt the VCG mechanisms to our resource assignment problem.

So assume that STEM's objective is to maximize the total DAV. This means maximizing the social welfare, and by the way the DAV is defined, this is maximized when the sum of the D_{ij}'s is minimized. In other words, maximizing the sum of the DAV(i,j)'s is equivalent to minimizing the sum of D_{ij}'s for the feasible pairs (agent, resource).

Maximum-total-DAV is solved by a maximum matching in the following bipartite graph G. G has resources and agents as nodes, and an edge between each pair of an agent *i* and a feasible resource *j*; this edge has weight = DAV_{ij}. See maximum weighted bipartite matching in [8].

Let DAV_i be the DAV of agent *i* in the maximum-total-DAV assignment **M**. Let B_{ji} be the DAV of agent *j* in an assignment **M'** of maximum total DAV that includes all the resources but does not include agent *i*.

Definition 2. Pricing Scheme PA: Price paid by agent *i* to STEM in an assignment **M** of maximum total DAV, called PA_i, is $\sum_{j\neq i} B_{ji} - \sum_{j\neq i} \mathrm{DAV}_j$.

Theorem 1: PA is: 1) truthful (i.e., the best strategy for each agent is to declare its true location, which means its true value for each resource), 2) individually rational (i.e., $\mathrm{PA}_i \leq \mathrm{DAV}_i$), and 3) each $\mathrm{PA}_i \geq 0$

Proof sketch: The theorem follows from the VCG theorem, with the Clarke pivot rule. More specifically, it follows from theorem 9.17 and Lemma 9.20 in [9]. 9.17 addresses a model in which there is a set of alternatives A (corresponding to the possible assignments in our model), and a valuation function V_i(a) of each player (agent) for each alternative a in A. V_i(a) corresponds to DAV_{ij}, where *j* is the resource assigned to agent *i* in assignment a. 9.17 indicates that a set of pricing schemes (mechanisms) is truthful, and 9.20 further refines these by indicating when these schemes are also individually rational and do not pay money to the players. Intuitively, this happens when agent i pays an amount that is "equal to the damage that he causes the other players – the difference between the social welfare of others with and without i's participation. In other words, the payments make each player internalize the externalities that he causes [9]". In our model, these translate into pricing scheme PA.[]

To see that payment scheme PA induces truth-telling, consider again the configuration of Figure 1. We assume that the value of using a resource is $10, and the value of time for each driver is $1/minute. In this case DAV_{11}=9, DAV_{12}=8, DAV_{21}=5, DAV_{22}=2, B_{12}=9, B_{21}=5, DAV_1=8, DAV_2=5. Thus in the assignment M = {(v_1,s_2), (v_2,s_1)} the price paid by v_1 is $\mathrm{PA}_1 = B_{21} - \mathrm{DAV}_2 = 0$. And the price paid by v_2 is $\mathrm{PA}_2 = B_{12} - \mathrm{DAV}_1 = 1$. Observe that if v_2 lies and says she is very close to (e.g. driving time 0.1 minutes from) s_1, and v_1 tells the truth, the maximum-total-DAV assignment would not change, and v_2's price would still be 1. Intuitively, the reason for this is that the price paid by v_2 depends on the damage that her assignment in M causes the other drivers. This is similar to Vickrey's second price auction, where the price paid by the winning player does not depend on the value she declared, but on the value declared by the 2nd highest bidder. In other words, the winner's price depends on the damage she causes the other players, which is the value to the 2nd highest bid.

Similarly, if v_1 lies and says he is very close to s_1, his price would still be 0 because the maximum-total-DAV assignment would still be the same and $PA_1 = B_{21} - DAV_2$.

Now observe that incentivizing truthfulness has its price. Specifically, STEM's revenue suffers due to the incentive that it provides for truthfulness. To see this, let the Naïve payment scheme PN be one in which each agent *i* declares to STEM its true location, and thus its DAV for each resource j (i.e. DAV_{ij}), and pays the price equal to the DAV of the assigned resource. As previously, STEM makes the assignments which maximize total welfare. Thus, since each agent declares its true location, then the payment of agent *i*, denoted PN_i, is DAV_i. Then:

$ThePriceOfTruthfulness = PA_i - PN_i = \sum_{j\neq i} B_{ij} - \sum_{j\neq i} DAV_j - DAV_i = \sum_{j\neq i} B_{ij} - \sum_j DAV_j < 0.$

The last inequality holds since the second sum is the value of the maximum-cost assignment, whereas the first sum is the value of some assignment that does not even include all the agents (i.e. it does not include *i*).

This means that each agent pays less under PA than under PN, and the difference is the price that the STEM pays to induce each agent to be truthful. Again, this is similar to the situation in which the Vickrey second-price auction is compared with the naïve auction, i.e. the one where each agent declares and pays his value for the item; in the auction case, the winner also pays less in Vickrey auction than in naïve auction. Specifically, in Vickrey's 2nd price auction, the house revenue is not the highest bid, but the 2nd highest.

The implication of this observation is that even if the PA pricing scheme and the GAG pricing scheme can be reconciled, i.e. the PA pricing scheme can be adapted to the GAG model, then Theorem 0 would probably not hold anymore.

5. EXTENSION TO ROUTING

The gap between optimum and equilibrium exists in routing as well, i.e. routing settles into an equilibrium state, which is often suboptimal. To see that consider the configuration in Figure 3 that represents the adaptation of the Braess paradox for our purpose.

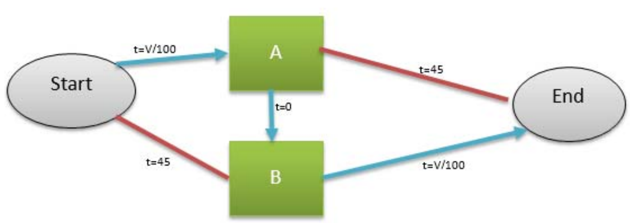

Figure 3: Routing 4000 vehicles from Start to End through the network with indicated travel times. t=V/100 means that the travel time on the link depends on the number of vehicles; it is the number of vehicles V divided by 100. So if 4000 vehicles use the link, the time to traverse the link is 40 minutes. t=45 means that the time is 45 minutes, independently of the number of traveling vehicles.

Specifically, if 4000 vehicles go from start to end their equilibrium route is: Start > A > B > End, with a total travel time of 85 minutes for each vehicle. This is an equilibrium route in the sense that if a vehicle deviates from this route, its total travel-time will increase. In contrast, the minimum travel-time routes are as follows. 2000 vehicles to go: Start > A > End, and the other 2000 to go: Start > B > End. In this case the travel-time for each vehicle is 65 minutes.

Here also, automatic negotiation and transactions among the STEM agents on the mobile devices of the 4000 vehicles can convert the equilibrium to an optimum. Observe that recent trends in transportation, geo-spatial computation, and wireless communication make such transactions seem a natural progression. For example, more and more people use smartphones and in-vehicle systems for navigation, and thus enter their destination at the start of a trip, even in familiar environments in order to get traffic information. Moreover, v2v communication seems imminent in the developed world [10].

6. RELATED WORK

In this paper we assume that all agents can receive information about available resources. Such information can indeed be obtained by already existing research work and technologies on monitoring and sensing open parking slots. Examples of research works dealing with detection of open parking slots include the use of ultrasonic sensor technology to determine the spatial dimensions of open parking slots [11], and the use of wireless sensors that are used to track open parking slots in a parking facility [12]. Beyond simple detection of slots, [3] show how to couple detection with sharing of the parking slot information in a mobile sensor network by presenting a methodology for vehicles driving past curbside parking slots to detect open ones, as opposed to having to spend on equipping each parking slot with wireless sensors for monitoring. These mobile sensors generate a map of parking slot availability.

In [4, 5], we introduced the so-called *Parking Slot Assignment Games (PSAG)* to analyze various parking related problems in competitive settings. The parking problem was studied in a centralized context as well as in the context of a distributed model with individual selfish agents, and a relationship between the Nash equilibrium and stable marriage assignments [13] was established in [5]. When drivers are selfish and cannot be controlled by a central authority, it is well accepted that the overall system converges to the Nash equilibrium since it describes a situation where they cannot improve on their incurred costs. In [5] we discussed pricing of resources (specifically parking slots) to convert the equilibrium to an optimum. In this paper we carry this work a step further by introducing financial transactions and truthfulness in bridging the gap from equilibrium to optimum.

Pricing of resources to obtain some system-wide objectives as studied in this paper has been considered in the past in other contexts for transportation applications. In the transportation literature this is commonly known as ``congestion pricing'' [14]. The most common type of congestion pricing is that of toll-like prices assessed on major urban areas or major roads to decrease the demand of entering to these areas and roads, and pricing strategies of similar type has been famously implemented in the central business district of Singapore [15] and in other major cities across the world. In principle, road tolling also attempts to bridge the gap between optimum and equilibrium [16, 17].

However, tolling is a form of taxation. In contrast, in this paper we propose to bridge this gap in a revenue neutral way, by P2P transactions among users.

This paper investigates the pricing problems in the context of algorithmic game theory which has a rich history, see textbooks such as [9] for further details.

7. CONCLUSION AND FUTURE WORK

In this paper we proposed STEM, a platform that facilitates transactions among users of spatio-temporal resources such as taxi cab customers, parking slots, and EV charging stations. Under the GAG payment scheme introduced in this paper, we showed that often such transactions can enable the transition from an equilibrium state to an optimum state, in a way that benefits the users as well as society (e.g. the environment). Furthermore, Theorem 0 showed that under the GAG payment scheme, society does not need to subsidize this transition from equilibrium to optimum.

The GAG payment scheme raises the issue of truthfulness. Specifically, in the vanilla GAG scheme users need to disclose to STEM private information such as their location, and they may be able to gain from being untruthful. Therefore we considered STEM truthfulness using the Vickrey-Clarke-Groves mechanisms studied in economics, and we showed how location truthfulness can be induced using such mechanisms. Nevertheless, truthful disclosure using VCG mechanisms comes at a price of reduced revenue, and therefore we conjecture that applying VCG will require a subsidy by society to sustain GAG. This conjecture is the subject of future work.

Another subject of future work is extension of the GAG payment scheme to sequential, rather than simultaneous arrival of the agents. From a practical point of view, is there a problem of scalpers, i.e. drivers circling around parking spaces to be paid by bona fide users of parking spaces?

Acknowledgments: This research was supported in part by the US Department of Transportation National University Rail Center (NURAIL); Illinois Department of Transportation (METSI); and National Science Foundation grants IIS-1213013, CCF-1216096, DGE-0549489, IIP-1315169.

8. REFERENCES

[1] Urban Mobility Report. 2011. Texas Transport Institute

[2] 2010 Pocket Guide to Transportation, Bureau of Transportation Statistics, US DOT.

[3] S. Mathur, T. Jin, N. Kasturirangan,J. Chandrashekharan, W. Xue, M. Gruteser, and W. Trappe. 2010. Parknet: Drive-by sensing of road-side parking statistics. In MobiSys, San Francisco, CA, June 2010

[4] Ayala, D., Wolfson, O., Xu, B., Dasgupta, B., and Lin, J. 2011. Parking Slot Assignment Games. In *Proc. of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (ACM GIS),* Chicago, IL, Nov. 2011, pp.299-308.

[5] Ayala, D., Wolfson, O., Xu, B., Dasgupta, B., and Lin, J. 2012. Pricing of Parking for Congestion Reduction. In *Proc. of the 20th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (ACM GIS)* Redondo Beach, CA, Nov. 2012.

[6] Ayala, D., Wolfson, O., Xu, B., Dasgupta, B., and Lin, J. 2012. Parking in Competitive Settings: A Gravitational Approach. In *Proceedings of the 13th International Conference on Mobile Data Management (MDM)* Bengaluru, India, July 2012.

[7] Cook, W, J., Cunningham, W, H., Pulleyblank, W, R., and A. Schrijver.1998. Combinatorial Optimization. *John Wiley & Sons*.

[8] http://en.wikipedia.org/wiki/Maximum_matching

[9] Nisan, N., Roughgarden, T., Tardos, E, and Vazirani, V. 2007. Algorithmic Game Theory. *Cambridge University Press, New York, NY, USA*.

[10] Wireless Wheels, the Economist, 2014. http://www.economist.com/news/leaders/21615599-connected-cars-will-make-driving-safer-cleaner-and-more-efficient-their-introduction-should-be?frsc=dg%7Cc. Sept 6, 2014

[11] Park, W., Kim, B., Seo, D., Kim, D., Lee, K. 2008. Parking Space Detection Using Ultrasonic Sensor in Parking Assistance System. *IEEE Intelligent Vehicles Symposium Eindhoven, The Netherlands*, June 2008

[12] Panja, B., Schneider, B., and Meharia, P. 2011. Wirelessly Sensing Open Parking Spaces: Accounting and Management of Parking Facility, *AMCIS 2011 Proceedings*

[13] Gale, D., and Shapley, L, S. 1962. College admissions and the stability of marriage, *The American Mathematical Monthly,* vol.69, pp.9-15

[14] Wie, B, W., and Tobin, R, L. 1998. Dynamic congestion pricing models for general traffic networks. *Transportation Research Part B – Methodological,* vol.32, pp.313-327

[15] Phang, S., and Toh, R, S. 2004. Road Congestion Pricing in Singapore: 1975 to 2003. *Transportation Journal,* vol.43, pp.16-25

[16] Arnott, R., Palma, A, D., and Robin, L. 1990. Economics of a Bottle Neck. *Journal of Urban Economics, 27, 11-130*

[17] Gonzales, E, J., and Daganzo, C, F. 2012. Morning commute with competing modes and distributed demand: User equilibrium, system optimum, and pricing. *Transportation Research Part B 46 (2012),* pp.1519–1534