TEL AVIV UNIVERSITY אוניברסיטת תל-אביב
The Raymond and Beverly Sackler Faculty of Exact Sciences
The Blavatnik School of Computer Science

Flow Equilibria via Online Surge Pricing

Thesis submitted in partial fulfillment

of graduate requirements for the degree

“Master of Sciences” in Tel Aviv University

School of Computer Science

By

Lior Shultz

Prepared under the supervision of Prof. Amos Fiat and Prof. Yishay Mansour

March 2018

Abstract

We explore issues of dynamic supply and demand in ride sharing services such as Lyft and Uber, where demand fluctuates over time and geographic location. We seek to maximize social welfare which depends on taxicab locations, passenger locations, passenger valuations for service, and the distances between taxicabs and passengers. Our only means of control is to set surge prices, then taxicabs and passengers maximize their utilities subject to these prices.

We study two related models: a continuous passenger-taxicab setting, similar to the Wardrop model, and a discrete (atomic) passenger-taxicab setting. In the continuous setting, every location is occupied by a set of infinitesimal strategic taxicabs and a set of infinitesimal non-strategic passengers. In the discrete setting every location is occupied by a set of strategic agents, taxicabs and passengers, passengers have differing values for service.

Afterwards we expand the continuous model to a time-dependent setting and study the corresponding online environment.

The utility for a strategic taxicab that drives from u to v and picks up a passenger at v is the surge price at v minus the distance from u to v. The utility for a strategic passenger at v that gets service is the value of the service to the passenger minus the surge price at v.

Surge prices are in passenger-taxicab equilibrium if there exists a min cost flow that moves taxicabs about such that (a) every taxicab follows a best response, (b) all strategic passengers at v with value above the surge price r_v for v, are served and (c) no strategic passengers with value below r_v are served (non-strategic infinitesimal passengers are always served).

This thesis computes surge prices such that resulting passenger-taxicab equilibrium

maximizes social welfare, and the computation of such surge prices is in poly time. Moreover, it is a dominant strategy for passengers to reveal their true values.

We seek to maximize social welfare in the online environment, and derive tight competitive ratio bounds to this end. Our online algorithms make use of the surge prices computed over time and geographic location, inducing successive passenger-taxicab equilibria.

Table of Contents

*

Acknowledgments

I would like to thank my supervisors Prof. Amos Fiat and Prof. Yishay Mansour for guiding me in the rewarding process of research.

Due to my limiting circumstances they have forfeited much of their personal time to meet and discuss the various work presented in this thesis, for that I am truly thankful. They have allowed me to experience many different methods of work and approaches and I cannot imagine this process without either of them. Thank you very much!

I Introduction

In the sharing economy[1] individual self-interested suppliers compete for customers. According to PWC, the *sharing economy* is projected to exceed 300 billion USD within 8 years. Lyft and Uber are prime examples of such systems. According to [16, 11] it is the users who gain the majority of the surplus from such systems, and significantly so. Contrawise, many studies suggest negative societal issues in the sharing economy (e.g., see [20, 12, 21, 3]).

Unlike salaried employees of livery firms, drivers for Uber (and other "gig" suppliers) are free to decide when they are working and what calls/employment to accept. *E.g.*, drivers can refuse to accept a call if it is too far away. To increase supply (and reduce demand) Uber introduced "surge pricing" which is a multiplier on the base price when demand outstrips supply. The surge price can be different at different locations.

In the past pricing schemes resulted in what was theorized to be negative work elasticity [6]. In their work it is suggested that drivers impose upon themselves "income targets". This means that drivers will work until they reach their target income for the day causing them to extend their hours in times of low payouts. Recent studies suggest that this is false, surging prices in times of peak demand seems to conjure positive work elasticity [8], allowing supply and demand to balance more efficiently.

[1] Also known as the "gig" economy.

1.1 Network Model, Surge Pricing, Utility, and Passenger-Taxicab Equilibria

Our goal is to maximize social welfare, defined as the sum of valuations of the users serviced by taxicabs, minus the cost associated with providing such service. We do so by setting surge prices (one per location), and let the system reach equilibrium. Our surge pricing schemes have several additional features such as envy freeness.

We consider two related settings:

- A continuous setting where supply and demand consist of infinitesimal quanta, *supply* and *demand* are modeled as fractional quantities at locations. This is analogous to the non-atomic traffic model used in Wardrop equilibria [24].
  - Here we assume that the taxicabs are strategic and respond to changing surge prices whereas passengers are non-strategic so that demand is insensitive to price (alternately, one may view these passengers as having high value for service).
  - The cost for a taxicab at location x to serve a customer at location y is the distance from x to y.
  - Our goal here is to set a surge price r_x at every location x so as to incentivize taxicabs to act in a way that maximizes social welfare, *i.e.*, all possible demand is serviced while the sum of distances traversed is minimized.
- A discrete setting where both taxicabs and passengers are strategic, and every taxicab and passenger is associated with some location.
  - In this setting both demand and supply may change as a function of the surge price. Every passenger has a value for service and every taxicab has a cost for service at a given location, *e.g.*, the distance to the location.

- Our goal here is to maximize social welfare (the sum of the values for the served customers minus the sum of costs of the taxicabs to do so).
- At every location x, we set a surge price r_x, that incentivizes taxicabs to serve passengers in a manner that maximizes social welfare.
- Moreover, maximizing social welfare is not only in equilibrium but also envy free.
- Every passenger at x whose value is strictly greater than r_x is served, and no passenger at x with value strictly less than r_x is served.

We define the utility for a taxicab at x to serve a passenger at y as the surge price at y, r_y, minus the distance from x to y. A passenger at x with value v has utility $v - r_x$ to be served by a taxicab, and utility zero if she takes no taxicab. Clearly, a passenger at x with $v - r_x < 0$ will refuse to take a taxicab.

We introduce the notion of a *passenger-taxicab equilibria*, for both continuous and discrete settings. A flow is a mapping from the current supply to some new supply. A flow has an associated cost which is the sum over edges of the flow along the edge times the length of the edge. A flow f is said to be a min cost flow that maps the current supply to the new supply if it achieves the minimal cost for moving the current supply to the new supply (this cost is also called the min earthmover cost).

A passenger-taxicab equilibria consists of a vector of surge prices $r = \langle r_x \rangle$, where r_x is the surge price at location x, current supply $s = \langle s_x \rangle$, new supply $s' = \langle s'_x \rangle$ and demand $d = \langle d_x \rangle$, such that, for any min cost flow from s to s', every taxicab and every passenger maximize their utility. *I.e.*, no taxicab can improve its utility by doing anything other than following the flow, every passenger at x who has value greater than r_x is in d_x and is served. Every passenger at x who has value less than r_x is not served.

The surge prices r_x are poly time computable. In the continuous setting this is polynomial in the number of locations, in the discrete setting this is polynomial in the number of passengers and taxicabs.

1.2 Maximizing Social Welfare in an Online Setting via Surge Pricing

We consider an online setting based on the continuous setting, where the time progresses in discrete time steps. In each time step the following occurs: First, a new demand allocation appears. Second, the online algorithm determines a new supply. Given an allocation of supply and demand, the demand served at a location is the minimum between the supply and demand at the location. The social welfare is the difference between the total demand served and the total movement cost, summed over all locations and time steps. The main new crux of our model is that the online algorithm (principal) can not impose a new supply allocation, but is limited to setting surge prices. If flow f is a flow equilibrium arising from these surge prices — strategic suppliers follow flow f. Our results on surge prices for flow equilibria imply that the online algorithm has flexibility in selecting the desired supply.

Trivially, for any metric, a simple algorithm that randomizes the start setting and doesn't move achieves a $\Theta(1/k)$ competitive ratio, where k is the number of locations. However, If the costs of moving from any location to any other location is 1, we give an optimal competitive ratio of $\Theta(\sqrt{1/k})$. If the demand sequence has the property that at any time and location the demand does not exceed $1/\rho$ ($\rho \geq 1$), then we show a tight competitive ratio bound of $\Theta(\sqrt{\rho/k})$. For more general metric spaces we show mainly negative results. Specifically, if all the distances are $1+\epsilon$ we show that the competitive ratio is no better than $(1+\epsilon)^2/(\epsilon k)$, which implies an optimal competitive ratio of $\Theta(1/k)$ for $\epsilon = \Theta(1)$.

Another extension we consider is when the average difference between successive demand vectors is bounded by δ (in total variation distance). In this case we show that simply matching supply to the current demand gives a competitive ratio of $1-\delta$ and show that the competitive ratio can not be better than $1-\delta/4$ (in the case that all the distances are 1).

1.3 Related Work

It has been observed in taxicab services that a mismatch between supply and demand, along with first-in-first-out scheduling of service calls, without restricting the "call radius", results in reduced efficiency and even market failure [1, 26]. This happens because taxicabs are dispatched to pick up customers at great distance because no closer taxicab is currently available, more time is wasted traveling to pick up clients, and the system performance degrades. Recent papers [9, 7] study how changing surge prices over time allow one to avoid such issues. These papers do not consider the issue of having geographically varying surge prices.

Assuming a stochastic passenger arrival rate, [2] uses a queue theoretic approach to model driver incentives in the system. The paper considers a simplistic dynamic pricing scheme, where there are two different pricing schemes for each node depending on the amount of drivers at said node. This model is compared to a simple flat rate. Drivers are assumed to calculate their incentives over several rides. The paper concludes that the dynamic pricing scheme can only achieve the welfare of the flat rate. However, the dynamic pricing scheme allows for the manager to have more room for error in calculating what the optimal rates are.

A central problem in handling a centralized taxi system involves routing empty cars between regions . Within the centralized mechanism, [5] shows that, assuming stochastic arrival of passengers, an optimal static strategy (*i.e.*, one that does not change it's routing policy based on current shortages) can be calculated by solving a linear programming problem.

Recently, and independently, a similar problem was studied in [18]. In their model, selfish taxicabs seek to maximize revenue over time. There is no explicit cost for travel, one loses opportunities by taking long drives. They derive prices in equilibria that maximize the sum of passenger valuations, but ignore travel costs. In contrast, we ignore the time dimension and focus on the passenger valuations and travel costs.

Competitive analysis of online algorithms [23, 22, 15] considers a worst case sequence of online events with respect to the ratio between the performance of an online algorithm and the optimal performance. In a centralized setting, task systems, [4], can be used to model a wide variety of online problems. Events are arbitrary vectors of costs associated with different states of the system, and an online algorithm may decide to switch states (at some additional cost). A strategic version of this problem, for a single agent, was considered in [10] where a deterministic incentive compatible mechanism was given. The competitive ratio for incentive compatible task system mechanisms is $O(1/k)$ where k is the number of states. We cannot use the incentive compatible task system mechanisms from [10] for two reasons: (1) in our setting there are a large number of strategic agents (many Uber drivers) split amongst a variety of different [task system] states (locations) rather than one such agent in a single state, and (2) the suppliers have both profits (payments) and loss (relocation).

Competitive analysis of the famous k-server problem [19] has largely driven the field of online algorithms. A variant of the k-server problem is known as the k-taxicab problem [13, 25]. Although the problem we consider herein and the k-taxicab problem both seek efficient online algorithms, and despite the name, the nature of the k-taxicab problem is quite different from the problem considered in this thesis. In the k-taxicab problem a single request occurs at discrete time steps and a centralized control routes taxicabs to pick up passengers, seeking to minimize the distances traversed by taxis while empty of passengers. Taxicabs are not selfish suppliers, and all requests must be satisfied. This is quite different from our setting where both demand and supply are spread about geographically, there are many strategic suppliers, and not all demand must be served.

II Model and Notation

2.1 *The Continuous Passenger-Taxicab Setting*

We model the network as a finite metric space $G = (V, E)$, where $\ell_{u,v} \geq 0$ is the distance between vertices $u, v \in V$. *I.e.*, $\ell_{u,v}$ is the cost to a taxicab to switch between vertices u and v. Infinitesimally small taxicabs reside in the vertices V.

Demand and supply are vectors in $[0,1]^{|V|}$ that sum to one. Given demand d and current supply s, we incentivize strategic taxicabs so that current supply s becomes new supply s' which services the demand d.

If the demand in vertex u is d_u, and the new supply in vertex u is s'_u, then the minimum of the two is the actual demand served (in vertex u). Note that if the two are not identical then there are either unhappy passengers (without service) or unhappy taxicabs (with no passengers to service). Formally,

Definition 2.1.1. we define the *demand served*, as follows:

- The *demand served* in vertex u, $\mathrm{ds}(s'_u, d_u)$, is the minimum of s'_u and d_u, *i.e.*, $\mathrm{ds}(s'_u, d_u) = \min(s'_u, d_u)$.
- Given a demand vector d and a supply vector s', the total demand served is $\mathrm{ds}(s', d) = \sum_{u \in V} \mathrm{ds}(s'_u, d_u) = \sum_{u \in V} \min(s'_u, d_u)$.

Switching supply from s to s' is implemented via a flow f. A flow from s to s' is a function $f(u, v) : V \times V \mapsto \mathbb{R}^{\geq 0}$ that has the following properties:

- For all $u, v \in V$, $f(u, v) \geq 0$.
- For all $v \in V$, $\sum_{u \in V} f(u, v) = s'_v$.

- For all $u \in V$, $\sum_{v\in V} f(u,v) = s_u$.

We define the earthmover distance between supply vectors,

Definition 2.1.2. The cost of flow f is $\mathrm{em}(f) = \sum_{u,v,\in V} f(u,v)\ell_{u,v}$. The earthmover distance from supply vector s to supply vector s' is

$$\mathrm{em}(s,s') = \min_{\text{flows } f \text{ from } s \text{ to } s'} \mathrm{em}(f).$$

We assume that switching supply from s to s' is implemented via a flow f of minimal cost. Note that there may be multiple flows with the same minimal cost — see Figures 3.2 and 3.3.

In order to incentivize our strategic taxicabs to move to a new supply vector, we use surge pricing in vertices.

Definition 2.1.3. Surge pricing is a vector, $r \in \mathbb{R}^{\geq 0}$, where r_v is the payment to a taxicab that serves demand in vertex $v \in V$.

We define the utility for an infinitesimal taxicab, given surge pricing r, as follows.

Definition 2.1.4. Given supply s, new supply s', surge prices r, demand d, and a min cost flow f from s to s', the utility for a taxicab that switches from vertex u to vertex v is

$$\mu(u \mapsto v|s',r,d) = r_v \cdot \left(\frac{\mathrm{ds}(s'_v,d_v)}{s'_v}\right) - \ell_{u,v}.$$

To motivate the above definition of utility $\mu(u \mapsto v|s',r,d)$, of switching from u to v, consider the following:

- The probability of serving a passenger in vertex v is $\frac{\mathrm{ds}(s'_v,d_v)}{s'_v}$. This follows since:
  - If passengers outnumber taxicabs in vertex v then any such taxicab will surely serve a passenger.

  - Alternately, if taxicabs outnumber passengers in vertex v then the choice of which taxicabs serve passengers is a random subset of the taxicabs.

- The profit from serving a passenger in vertex v is equal to the surge price for that vertex, r_v.

- The cost of serving a passenger in vertex v, given that the taxicab was previously in vertex u, is $\ell_{u,v}$.

Finally, we define the notion of a passenger-taxicab equilibrium, where no infinitesimal taxicab can benefit from deviations.

Definition 2.1.5. Given a demand vector d, current supply vectors s, and new supply s', we say that a surge pricing r is in *passenger-taxicab equilibrium*, if for every min cost flow f from s to s', for every $u, v \in V$ such that $f(u, v) > 0$ we have that

$$\mu(u \mapsto v|s', r, d) = \max_{w \in V} \mu(u \mapsto w|s', r, d). \tag{2.1}$$

I.e., every infinitesimal taxicab is choosing a best response. Such a passenger-taxicab equilibrium is said to *induce supply s'.*

Our goal in the continuous setting is to set surge prices so that the new supply $s' = d$ is a passenger-taxicab equilibrium.

In this continuous setting we take demand d to be insensitive to the surge prices. In the next section we describe the discrete setting where both the demand and the supply are sensitive to the prices. One could define a continuous passenger-taxicab setting where every location has an associated density function for passenger valuations. Then, we could convert this continuous setting to an instance of the discrete passenger-taxicab setting with $1/\epsilon$ taxicabs/passengers. Under appropriate conditions, this will give a good approximation to a continuous passenger-taxicab setting where both demand and supply are sensitive to surge pricing.

2.2 *The Discrete Passenger-Taxicab Setting*

As above, we model the network as a finite metric space $G = (V, E)$, and the cost to a taxicab to switch between vertices u and v is the distance between them, $\ell_{u,v}$. Unlike the continuous case, there is an integral number of taxicabs and passengers at every vertex.

Let $B = \{b_1, \dots, b_m\}$ be a set of m passengers and $T = \{t_1, \dots, t_n\}$ be a set of n taxicabs. Every passenger $b_i \in B$ has a value $\text{value}(b_i) \geq 0$ for service. A supply s is a vector $s = \langle s_v \rangle_{v \in V}$ where $s_v \subseteq T$ for all $v \in V$, $\cup_{v \in V} s_v = T$, and $s_v \cap s_u = \emptyset$ for all $u, v \in V$, $u \neq v$.

A profile P is a partition of the passengers B, where for each $u \in V$ the set $P_u \subseteq B$ is the set of passengers at u. A demand is a function of a vertex and a surge price at the vertex. We define the function d_v as follows:

$$d_v(r_v) = \{b_i \in P_v | \text{value}(b_i) \geq r_v\}.$$

Ergo, $d_v(r_v)$ is the set of passengers at vertex v that are interested in service given that the price is r_v, *i.e.*, those passengers whose value is at least r_v. Note that $d_v(0) = P_v$.

For ease of notation, we denote a collection of entities x_v for each vertex $v \in V$, by $x = \langle x_v \rangle_{v \in V}$. For example, $s = \langle s_v \rangle_{v \in V}$, $d = \langle d_v \rangle$, and $r = \langle r_v \rangle_{v \in V}$.

Define a flow f from supply s to supply s' as follows. The flow $f(x, y) : V \times V \mapsto \mathbb{Z}^+$ has the following properties:

- For all $u, v \in V$, $f(u, v) \in \mathbb{Z}^+$.
- For all $u \in V$, $\sum_{v \in V} f(u, v) = |s_u|$.
- For all $v \in V$, $\sum_{u \in V} f(u, v) = |s'_v|$.

The flow from a vertex u is equal to the number of taxicabs at u under supply s, *i.e.*, $|s_u|$. The flow into a vertex v is equal to the number of taxicabs at v under supply s', *i.e.*, $|s'_v|$. The cost of a flow in the discrete setting is the same as the cost of a flow in the continuous setting (Definition 2.1.2), *i.e.*, $\sum_{u,v\in V} f(u,v)\ell_{u,v}$.

We now define the demand served at a vertex u,

Definition 2.2.1. For a vertex v, given a supply s'_v, a surge price r_v, and a demand $d_v(r_v)$, we define the *demand served*, $\mathrm{ds}_v(s'_v, d_v, r_v) \subseteq P_v$, as follows:

- If $|d_v(r_v)| \leq |s'_v|$ then $\mathrm{ds}_v(s'_v, d_v, r_v) = d_v(r_v)$.
- If $|s'_v| < |d_v(r_v)|$ then $\mathrm{ds}_v(s'_v, d_v, r_v)$ is the set of the $|s'_v|$ highest valued passengers from $d_v(r_v)$, breaking ties arbitrarily.

Given demand functions d, surge prices r, and new supply s', the total demand served $\mathrm{ds}(s', d, r)$ and its value $\mathrm{dsv}(s', d, r)$ is given by

$$\begin{aligned} \mathrm{ds}(s', d, r) &= \cup_{v\in V}\mathrm{ds}_v(s'_v, d_v, r_v); \\ \mathrm{dsv}(s', d, r) &= \sum_{b_i \in \mathrm{ds}(s',d,r)} \mathrm{value}(b_i). \end{aligned}$$

Definition 2.2.2. The social welfare is the difference between the sum of the values of the passengers served and the cost of the min cost flow, which is the sum of the distances traveled by the taxis. Namely, for current supply s, new supply s', demand functions d, and surge prices r, the social welfare is

$$SW(s, s', r, d) = \mathrm{dsv}(s', d, r) - em(s, s'). \tag{2.2}$$

Remark: we did not define social welfare in the continuous passenger-taxicab setting where the passengers are price insensitive. However, one can view the social welfare in the price-insensitive demand setting as a special case of the responsive demand setting when all passenger valuations are very high.

Like the definitions for utility and passenger-taxicab equilibria in the continuous case, one can define them for the discrete case: The utility of a taxicab $t_j \in s_u$ moving from u to v, given new supply s', surge prices r and demand functions d, is

$$\mu_{t_j}(u \mapsto v | s', r, d) = \frac{\min(|d_v(r_v)|, |s'_v|)}{|s'_v|} \cdot r_v - \ell_{u,v}.$$

Definition 2.2.3. Given demand d , current supply s and new supply s', surge prices r are said to be in *passenger-taxicab* equilibrium if for every min cost flow f from s to s' and for any u, v such that $f(u, v) > 0$ we have that

- Taxicabs are choosing a best response: $\mu_{t_j}(u \mapsto v | s', r, d) = \max_{w \in V}(\mu_{t_j}(u \mapsto w | s', r, d))$.
- All passengers $b \in B$ with value$(b) > r_{\text{loc}(b)}$ are served. No passengers $b \in B$ with value$(b) < r_{\text{loc}(b)}$ are served.

2.3 Online Setting

In the online setting we inherit the continuous model setting, adding a function of time. Time progresses in discrete time steps $1, 2, \ldots, T$. At time t the demand vector $d^t = (d_1^t, d_2^t, \ldots, d_k^t)$ associates each vertex $v \in V$ with some demand $d_v^t \geq 0$, and we assume that the total demand $\sum_i d_i^t = 1$. One should not think of a time step as being instantaneous, but rather as a period of time during which the demands remain steady.

Every time step t also has an associated supply vector $s^t = (s_1^t, s_2^t, \ldots, s_k^t)$, where $s_i^t \geq 0$ and $\sum_i s_i^t = 1$ for all t. The supply at time t is a "reshuffle" of the supply at time $t-1$, by having infintestimally small suppliers moving about the network. In our model, the time required for suppliers to adjust supply from s^{t-1} to s^t is small relative to the period of time during which demand d^t is valid.

If the demand in vertex i at time t is d_i^t, and the supply in vertex i at time t is s_i^t, then the minimum of the two is the actual demand served (in vertex i at time

t). Note that if the two are not identical then there are either unhappy customers (without service) or unhappy suppliers (with no customer to service). Formally, we define the benefit derived during each time period, the *demand served*, as in the continuous model.

We define the social welfare as follows:

Definition 2.3.1. Given a demand sequence $d = (d^1, \dots, d^T)$ and a supply sequence $s = (s^1, \dots, s^T)$ we define the social welfare

$$\mathrm{sw}(s, d) = \mathrm{ds}(s, d) - \mathrm{em}(s) = \sum_{t=1}^{T} \mathrm{ds}(s^t, d^t) - \sum_{t=2}^{T} \mathrm{em}(s^{t-1}, s^t).$$

An online algorithm for social welfare follows the following structure. At time $t = 1, 2, \dots, T$:

1. A new demand vector d^t appears.

2. The online algorithm determines what the supply vector s^t should be. (Indirectly, by computing and posting surge prices so that the resulting passenger-taxicab-equilibrium induces supply s^t).

The goal of the online algorithm is to maximize the social welfare as given in Definition 2.3.1: Compute a supply sequence s, so as to maximize $\mathrm{sw}(s, d)$. The supply vector s^t is a function of the demand vectors $d^1, \dots, d^t$ but not of any demand vector d^τ, for $\tau > t$. Implicitly, we assume that the passenger-taxicab equilibrium is attained quickly relative to the rate at which demand changes.

The competitive ratio of such an online algorithm, Alg, is the worst case ratio between the numerator: the social welfare resulting from the demand sequence d and the online supply $\mathrm{Alg}(d)$, and the denominator: the optimal social welfare for the same demand sequence, *i.e.*,

$$\min_d \frac{\mathrm{sw}(\mathrm{Alg}(d), d)}{\max_s \mathrm{sw}(s, d)}.$$

III The Continuous Passenger-Taxicab Setting

In this section we deal with the continuous passenger-taxicab setting. Given current supply s, demand d and new supply $s' = d$, we show how to set surge prices r such that they are in passenger-taxicab equilibria. Moreover, for these s, d, and r, the only possible s' which results in a passenger-taxicab equilibria is $s' = d$. (Similar techniques give surge prices that induce [almost] arbitrary supply vectors, $\tilde{s}$, see below).

Proof overview: Given some min cost flow f^* from supply s to demand d, we construct a unit demand market, with bidders and items. For every x, y such that $f^*(x, y) > 0$ we construct a bidder and an item. We also define bidder valuations for all items. This unit demand market has Walrasian clearing prices that maximize social welfare (Lemma 3.0.2). We show how we can convert the Walrasian prices on items to surge pricing (Lemma 3.0.4).

We then show and that the resulting surge pricing has a passenger-taxicab equilibrium which induces supply equals demand (Lemma 3.0.5) and it is the case with all all passenger-taxicab equilibria (Lemma 3.0.7). Lemma 3.0.6 shows that the incentive requirements in Equation (2.1) also hold for any min cost flow $f \neq f^*$, from s to d. This proves *Theorem 3.0.8.*

As a running example, consider the road network in Figure 3.1. Also, assume that the supply vector $s^{t-1} = \langle \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0, 0, 0 \rangle$ and demand vector $d^t = \langle 0, 0, \frac{1}{8}, \frac{3}{8}, \frac{3}{8}, \frac{1}{8} \rangle$. Two minimum cost flows are given in Figures 3.2 and 3.3. Both these flows have cost 1.

Given a minimum cost flow f^*, we define a unit demand market setting as follows:

- Items M^{f^*}, and unit demand bidders B^{f^*}, both of which are indexed by pairs

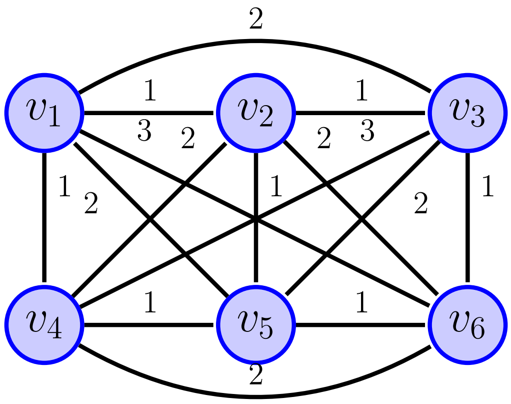

Figure 3.1: Example road network, with costs along edges.

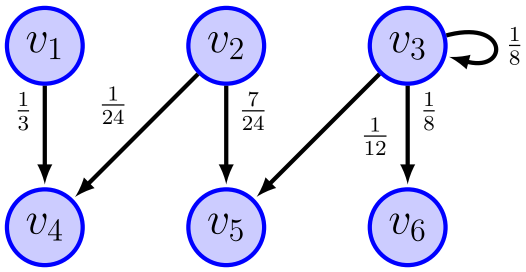

Figure 3.2: A min earthmover cost flow from supply vector $s^{t-1} = \langle \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0, 0, 0 \rangle$ **to demand vector** $d^t = \langle 0, 0, \frac{1}{8}, \frac{3}{8}, \frac{3}{8}, \frac{1}{8} \rangle$.

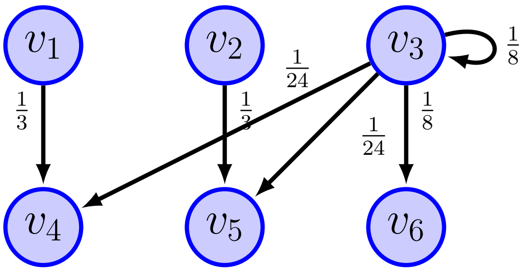

Figure 3.3: Another min earthmover cost flow from from supply vector $s^{t-1} = \langle \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0, 0, 0 \rangle$ **to demand vector** $d^t = \langle 0, 0, \frac{1}{8}, \frac{3}{8}, \frac{3}{8}, \frac{1}{8} \rangle$.

	m_{14}	m_{24}	m_{25}	m_{33}	m_{35}	m_{36}
b_{14}	3	3	2	2	2	1
b_{24}	2	2	3	3	3	2
b_{25}	2	2	3	3	3	2
b_{33}	1	1	2	4	2	3
b_{35}	1	1	2	4	2	3
b_{36}	1	1	2	4	2	3
p_{ij}	0	0	1	3	1	2

Figure 3.4: Item valuations for bidders $B^f = \{b_{14}, b_{24}, b_{25}, b_{33}, b_{35}, b_{36}\}$, **items** $M^f = \{m_{14}, m_{24}, m_{25}, m_{33}, m_{35}, m_{36}\}$, **where** f **is the min earthmover flow given in Figure 3.2. Note that in Figure 3.1 we have** $\max_{ij}(\ell_{ij}) = 3$ **and thus** $C = 4$. **The last row gives Walrasian market clearing prices for items** m_{ij}. **Note that** $p_{ij} = p_{ij'}$ **for all** $b_{ij}, b_{ij'} \in B^f$.

of vertices, where

$$M^{f^*} = \{m_{xy} | x, y \in V, f^*(x,y) > 0\} \qquad B^{f^*} = \{b_{wz} | w, z \in V, f^*(w,z) > 0\}.$$

- We set the value of item $m_{xy} \in M^{f^*}$ to bidder $b_{wz} \in B^{f^*}$ to be,

$$\zeta_{b_{wz}}(m_{xy}) = C - \ell_{w,y}, \qquad \text{where } C = \max_{i,j} \ell_{i,j} + 1.$$

- The utilities of bidders are unit demand and quasi-linear, *i.e.*, the utility $\eta_{b_{wz}}$ of bidder $b_{wz} \in B^{f^*}$ for item set S and price p is

$$\eta_{b_{wz}}(S) = \max_{m_{xy} \in S} \zeta_{b_{wz}}(m_{xy}) - p.$$

As an example, let f^* be the minimum cost flow of Figure 3.2. The market induced by f^* is illustrated in Figure 3.4.

Given a flow f^*, bidders B^{f^*} and items M^{f^*} we define the following weighted bipartite graph $G(B^{f^*}, M^{f^*}, E)$, where between bidder $b_{wz} \in B^{f^*}$ and item $m_{xy} \in M^{f^*}$ there is an edge of weight $C - \ell_{w,y} \geq 1$.

Definition 3.0.1. Given a flow f^*, a *matching* between bidders B^{f^*} and items M^{f^*} is a function $\pi : B^{f^*} \mapsto M^{f^*} \cup \{\emptyset\}$, where bidder $b \in B^{f^*}$ is matched to item $\pi(b) \in M^{f^*}$ or unmatched (if $\pi(b) = \emptyset$), such that no two bidders $b_1, b_2 \in B^{f^*}$ are matched to the same item $m \in M^{f^*}$.

As there is an edge between every bidder b_{wz} and every item m_{xy} with weight $C - \ell_{w,y} \geq 1$, the maximum weight matching is a perfect matching between bidders and items and the mapping π never assigns $\emptyset$ to a bidder.

Lemma 3.0.2. *The matching g where $g(b_{wz}) = m_{wz}$, maximizes social welfare. In addition, there exist Walrasian prices for which g is a competitive market equilibrium.*

	v_1	v_2	v_3	v_4	v_5	v_6
r_i^t	1	1	1	4	3	2

Figure 3.5: Surge prices resulting in a flow-equilibrium with $s^t = d^t$. These surge price for v_j is $C - p_{ij}$ if there exists some bidder $b_{ij} \in B^f$ and 1 otherwise. The Walrasian prices p_{ij} appear in Figure 3.4.

Proof. The proof is via contradiction. Assume there exists some matching $\tilde{g} : B^{f^*} \mapsto M^{f^*}$ with strictly greater social welfare than the matching g. For a bidder $b \in B^{f^*}$, define $\tilde{h}(b) = z$, iff $\tilde{g}(b) = m_{wz}$ for some $w \in V$, and $h(b) = z$, iff $g(b) = m_{wz}$ for some $w \in V$. Note that $h(b_{wu}) = u$ so for a given w and z we have $|\{u|h(b_{wu}) = z\}| = 1$ if $f(w, z) > 0$ and zero otherwise.

Choose ϵ to be the minimum non-zero flow in f^*, *i.e.*, $\epsilon = \min\{f^*(w, z)|f^*(w, z) > 0\}$. We now define a flow f', which is a slight perturbation of flow f^*. In flow f', the flow from w to z is:

$$f'(w, z) = f^*(w, z) + \epsilon\Big(\Big|\Big\{u|\tilde{h}(b_{wu}) = z\Big\}\Big| - |\{u|h(b_{wu}) = z\}|\Big).$$

We first prove that f' is a valid flow, and later we show that it has a lower cost than f^*, in contradiction to the minimality of f^*.

Lemma 3.0.3. *Flow f' is a valid flow from supply vector s^{t-1} to demand vector d^t.*

Proof. Consider the requirements that f' be a valid flow:

- For all $x, y \in V$, $f'(x, y) \geq 0$: By definition of f' if $f^*(w, z) = 0$ then $f'(w, z) \geq 0$ and if $f^*(w, z) > 0$ then $f'(w, z) \geq f^*(w, z) - \min\{f^*(w, z)|f(w, z) > 0\} \geq 0$.

- For all $x \in V$, $\sum_y f'(x, y) = s_x^{t-1}$: By definition of f^* we have $\sum_y f^*(x, y) =$

s_x^{t-1}. Thus,

$$\begin{aligned}
\sum_y f'(x,y) &= \sum_y \left(f^*(x,y) + \left(\left| \left\{ u | \tilde{h}(b_{xu}) = y \right\} \right| - \left| \{ u | h(b_{xu}) = y \} \right| \right) \cdot \epsilon \right) \\
&= s_x^{t-1} + \left(\sum_y \left| \left\{ u | \tilde{h}(b_{xu}) = y \right\} \right| - \sum_y \left| \{ u | h(b_{xu}) = y \} \right| \right) \cdot \epsilon \\
&= s_x^{t-1} + \left(\left| \left\{ u | b_{xu} \in B^f \right\} \right| - \left| \left\{ u | b_{xu} \in B^f \right\} \right| \right) \cdot \epsilon \\
&= s_x^{t-1}.
\end{aligned}$$

- For all $y \in V$, $\sum_x f'(x,y) = d_y^t$: By definition of f^* we have $\sum_x f^*(x,y) = d_y^t$. Thus,

$$\begin{aligned}
\sum_x f'(x,y) &= \sum_x \left(f^*(x,y) + \left(\left| \left\{ u | \tilde{h}(b_{uy}) = x \right\} \right| - \left| \{ u | h(b_{uy}) = x \} \right| \right) \cdot \epsilon \right) \\
&= d_y^t + \left(\sum_x \left| \left\{ u | \tilde{h}(b_{uy}) = x \right\} \right| - \sum_x \left| \{ u | h(b_{uy}) = x \} \right| \right) \cdot \epsilon \\
&= d_y^t + \left(\left| \left\{ u | b_{uy} \in B^f \right\} \right| - \left| \left\{ u | b_{uy} \in B^f \right\} \right| \right) \cdot \epsilon \\
&= d_y^t.
\end{aligned}$$

$\square$

From the fact that $\tilde{g}$ has a higher social welfare we get,

$$\sum_{w,z: b_{wz} \in B^{f^*}} \zeta_{b_{wz}}(\tilde{g}(b_{wz})) > \sum_{w,z: b_{wz} \in B^{f^*}} \zeta_{b_{wz}}(g(b_{wz})) .$$

Using the definition of the valuations we have,

$$\sum_{w,z: b_{wz} \in B^{f^*}} C - \ell_{w,\tilde{h}(b_{wz})} > \sum_{w,z: b_{wz} \in B^{f^*}} C - \ell_{w,h(b_{wz})} .$$

This implies that

$$\sum_{w,z: b_{wz} \in B^{f^*}} \ell_{w,h(b_{wz})} > \sum_{w,z: b_{wz} \in B^{f^*}} \ell_{w,\tilde{h}(b_{wz})} .$$

Using this last inequality, it follows that the cost of f^* (Definition 2.1.2) satisfies

$$
\begin{aligned}
\mathrm{em}(f^*) &= \sum_{x,y} f^*(x,y) \cdot \ell_{x,y} \\
&> \sum_{x,y} f^*(x,y) \cdot \ell_{x,y} + \left(\sum_{w,z: b_{wz} \in B^f} \ell_{w,\tilde{h}(b_{wz})} - \sum_{w,z: b_{wz} \in B^f} \ell_{w,h(b_{wz})} \right) \cdot \epsilon = \mathrm{em}(f'),
\end{aligned}
$$

which contradicts the fact that flow f^* is a minimum cost flow.

The fact that there exist Walrasian prices for g that are in competitive market equilibrium follows from [14]. This concludes the proof of Lemma 3.0.2. $\square$

Let the Walrasian price of m_{xy} be p_{xy} as guaranteed by the lemma above. We first show that any two prices which correspond to the same vertex must have the same price.

Lemma 3.0.4. *For any two items m_{xy} and $m_{x'y}$ we have $p_{xy} = p_{x'y}$.*

Proof. For contradiction assume that $p_{xy} > p_{x'y}$. Let b_{wz} be the bidder assigned m_{xy}. Thus, for item m_{xy} bidder b_{wz} has utility $\eta_{b_{wz}}(m_{xy}) = C - \ell_{w,y} - p_{xy} < C - \ell_{w,y} - p_{x'y} = \eta_{b_{wz}}(m_{x'y})$ which implies that m_{xy} is not in the demand set for bidder b_{wz}. A contradiction to the fact that p are Walrasian prices. $\square$

For any $y \in V$ such that there exist items of the form m_{xy} for some $x \in V$, let p_y denote the Walrasian price for such items (By Lemma 3.0.4 all those Walrasian prices are identical). If no items of the form m_{xy} exist, this implies that demand at vertex y, $d_y^t = 0$, and we can set $p_y = 0$. Define surge prices, $r_y^t = C - p_y$, for all $y \in V$.

Lemma 3.0.5. *Given current supply s^{t-1} and demand d^t, surge prices $r_y^t = C - p_y$, new support $s' = d$, and $x, y, w \in V$ such that $f^*(x,y) > 0$ then*

$$\mu^t(x \to y | s', r, d) \geq \mu^t(x \to w | s', r, d).$$

Proof. Let x, y be such that $f^*(x, y) > 0$.Then,

$$\begin{aligned}
\mu^t(x \to y|s', r, d) &= r^t_y \cdot \min\left(1, \frac{d^t_y}{s^t_y}\right) - \ell_{x,y} && (3.1)\\
&= r^t_y - \ell_{x,y} && (3.2)\\
&= C - p_y - \ell_{x,y} && (3.3)\\
&= \eta_{b_{xy}}(m_{xy}) && (3.4)\\
&\geq \eta_{b_{xy}}(m_{zw}) \qquad \forall m_{zw} \in M^{f^*} \Leftrightarrow \forall m_{zw} : s'_w > 0 && (3.5)\\
&= C - p_w - \ell_{x,w} && (3.6)\\
&= r^t_w - \ell_{x,w} && (3.7)\\
&\geq r^t_w \cdot \min\left(1, \frac{d^t_w}{s^t_w}\right) - \ell_{x,w} && (3.8)\\
&= \mu^t(x \to w|s', r, d) && (3.9)
\end{aligned}$$

Equations (3.1),(3.9) follow from the definition of the utility in the continuous passenger-taxicab setting, definition 2.1.4.

Equations (3.2),(3.8) follows from considering the passenger-taxicab equilibrium where $s^t = d^t$ resulting in $\frac{d^t_y}{s^t_y} = 1$ for all y.

Equations (3.3),(3.7) follow from the definition of the surge prices.

Equations (3.4),(3.6) follow from the definition of the utility in the market setting.

Equation (3.5) follows from the market equilibrium.

So, we have that

$$\mu^t(x \to y|s', r, d) \geq \mu^t(x \to w|s', r, d) \qquad \forall w : s'_w > 0.$$

It remains to consider $\mu^t(x \to w|s', r, d)$ for w such that $s'_w = 0$. In this case the surge price at w is zero, so the utility $\mu^t(x \to w|s', r, d) \leq 0$. □

The following lemma shows that the incentive requirements of Equation (2.1) hold, not only for the flow f^*, but also for *any* min cost flow from s to d.

Lemma 3.0.6. *Fix current supply s^{t-1} and demand d^t, surge prices $r^t_y = C - p_y$,* *and new support $s' = d$. Let f' be an arbitrary min cost flow from s to $s' = d$, then, for any $x, y, w \in V$ such that $f'(x, y) > 0$ we have that*

$$\mu^t(x \to y|s', r, d) \geq \mu^t(x \to w|s', r, d).$$

Proof. Define

$$\Gamma(f) = \sum_{u \in V} \sum_{v \in V} f(u, v) \cdot (r_v - \ell_{u,v}) = \sum_{v \in V} s'_v \cdot r_v - \mathrm{em}(s, s').$$

As f' and f^* are both min cost flows from s to s' we have that $\Gamma(f^*) = \Gamma(f')$.

For contradiction assume there exist some u, v such that $f'(u, v) > 0$ and $\mu(u \mapsto v|s', r, d) < \max_{w \in V} \mu(u \mapsto w|s', r, d)$.

$$\begin{aligned}
\Gamma(f^*) &= \sum_{u \in V} \sum_{v \in V} f^*(u, v) \cdot (r_v - \ell_{u,v}) && (3.10)\\
&= \sum_{u \in V} \sum_{v \in V} f^*(u, v) \cdot \max_{v \in V}(r_v - \ell_{u,v}) && (3.11)\\
&= \sum_{u \in V} s_u \cdot \max_{v \in V}(r_v - \ell_{u,v}) && (3.12)\\
&= \sum_{u \in V} \sum_{v \in V} f'(u, v) \cdot \max_{v \in V}(r_v - \ell_{u,v}) && (3.13)\\
&> \sum_{u \in V} \sum_{v \in V} f'(u, v) \cdot (r_v - \ell_{u,v}) = \Gamma(f') \,. && (3.14)
\end{aligned}$$

Eq. (3.10) follows from the definition of Γ. Eq. (3.11) follows from Lemma 3.0.5. Eq. (3.12) and (3.13) follows from the definition of a flow, since for any flow f from s we have that $s_u = \sum_{v \in V} f(u, v)$. Eq. (3.14) holds since we assumed, for contradiction, that there exist some u, v such that $f'(u, v) > 0$ and $\mu(u \mapsto v|s', r, d) < \max_{w \in V} \mu(u \mapsto w|s', r, d)$. Hence, we reached a contradiction to the assumption that f' is a min cost flow. □

If follows from the Lemma above that computing surge prices r via flow f^* ensures that taxicab routing using any other min cost flow f' is also a best response under surge prices r.

Next we show that all relevant passenger-taxicab equilibria all have new supply $s^t = d^t$.

Lemma 3.0.7. *Given current supply s, demand d, and surge prices $r^t_y = C - p_y$, all passenger-taxicab equilibria induce $s^t = d$.*

Proof. Let f be a min cost flow from s to d. For contradiction, assume that that there exists some $\bar{s} \neq d$ such that $s, s' = \bar{s}, d, r$ are in a passenger-taxicab equilibrium. Let f' be some min cost flow from s to $\bar{s}$.

Consider $H = \left\{y | \sum_x f'(x, y) > d^t_y\right\}$ (*i.e.*, the set of all vertices for which the flow f' results in strictly more supply than demand). Since $s' \neq d$ and they both sum to 1, we have that $H \neq \emptyset$. Let $H' = \{x | \exists y \in H \,\text{s.t.}\, f'(x, y) > 0\}$ (*i.e.*, the set of all vertices from which supply flows to H). As $H \neq \emptyset$ it follows that $H' \neq \emptyset$.

We claim that there exists some $w \in H'$, $y \notin H$, such that $f(w, y) > 0$. For contradiction assume that all the flow in f from vertices in H' is to vertices in H. By definition of flows: $\sum_{y \in V} f(x, y) = s^{t-1}_x = \sum_{y \in V} f'(x, y)$. We now have,

$$\begin{aligned}
\sum_{x \in H'} \sum_{y \in V} f(x, y) &= \sum_{x \in H'} \sum_{y \in H} f(x, y) \\
&= \sum_{x \in H'} \sum_{y \in H} f'(x, y) \\
&= \sum_{y \in H} \sum_{x \in H'} f'(x, y) \\
&> \sum_{y \in H} d^t_y \\
&= \sum_{x \in H'} \sum_{y \in H} f(x, y),
\end{aligned}$$

which is a contradiction.

This implies that there exist $w \in H'$, $x \notin H$ such that $f(w, x) > 0$. Since $w \in H'$ there also exists some $y \in H$ such that $f'(w, y) > 0$.

We have shown that $s, s' = d, d, r$ is in passenger-taxicab equilibrium, this implies that

$$\mu^t(w \mapsto x | s', r, d) = r^t_x - \ell_{w,x} \geq r^t_y - \ell_{w,y} = \mu^t(w \mapsto y | s', r, d).$$

Since $y \in H$ we have we have that $\sum_u f'(u,y) > d_y^t$ resulting in the utility

$$\mu^t(w \mapsto y|s', r, d) = \left(r_y^t - \ell_{w,y}\right) \cdot \min\left(1, \frac{d_y^t}{s_y^t}\right) < r_y^t - \ell_{w,y} \leq r_x^t - \ell_{w,x} = \mu^t(w \mapsto x|s', r, d),$$

where $x \notin H$ implies the last equality. This is in contradiction to the $s, s' = \bar{s}, d, r$ being a passenger-taxicab equilibrium. $\square$

Theorem 3.0.8 follows from Lemma 3.0.5, Lemma 3.0.6, and Lemma 3.0.7.

Theorem 3.0.8. *Given distances $\ell_{i,j}$ and an arbitrary supply vector, $s^{t-1} = \langle s_1^{t-1}, \dots, s_k^{t-1} \rangle$. Let the demand vector be $d^t = \langle d_1^t, \dots, d_k^t \rangle$. Then, there exists a surge price vector $r^t = \langle r_1^t, \dots, r_k^t \rangle$ that results in a passenger-texicab equilibrium which induces a supply $s^t = d^t$. Moreover, any passenger-taxicab equilibrium of r^t induces supply $s^t = d^t$, and the surge prices r^t can be computed in polynomial time.*

We can extend the result from equating supply and demand to modifying the supply vector s^{t-1} to any supply s^t, with the restriction that if $s_i^t > 0$ then $d_i^t > 0$. The new surge prices are computed as follows. First we compute, as before, the surge prices r^t from s^{t-1} to d^t. Then, we set $\bar{r^t}_i = \max\{1, \frac{s_i^t}{d_i^t}\} r_i^t$ and the resulting surge prices are $\bar{r^t}$. In a similar way we can establish,

Theorem 3.0.9. *Let $d^t = \langle d_1^t, \dots, d_k^t \rangle$ and let $\alpha = \langle \alpha_1, \dots, \alpha_k \rangle$ be the target supply vector, subject to the restriction that if $\alpha_i > 0$ then $d_i^t > 0$. Then there exist surge prices $\bar{r^t}$ for which some passenger-taxicab equilibrium induces supply α.*

IV The Discrete Passenger-Taxicab Setting

In this section we consider a more realistic scenario where both demand and supply are sensitive to the surge pricing. All else being equal, higher surge prices mean less demand and more supply.

We define social welfare to be the sum of valuations of passengers served minus the sum of the distances traversed by the taxicabs to serve these passengers (Definition 2.2.2). Given current supply s and a passenger profile P, we give an algorithm for computing surge prices r that creates a passenger-taxicab equilibrium that maximizes social welfare.

The location of a passenger b_i and taxi t_j is denoted by $\text{loc}(b_i)$ and $\text{loc}(t_j)$ respectively (*i.e.*, $b_i \in P_{\text{loc}(b_i)}$ $t_j \in s(\text{loc}(t_j)))$. For brevity, we use the notation $\overline{\ell}_{i,j} = \ell_{\text{loc}(b_i),\text{loc}(t_j)}$.

4.1 Maximizing Social Welfare

As in the continuous case, we reduce the problem of computing surge prices to computing market clearing prices in a unit demand market. Given a set of passengers B and taxicabs T, we construct a unit demand market $M(B,T)$, where B is the set of buyers and T is the set of items. For the unit demand market, $M(B,T)$, we set the value of buyer $b_i \in B$ for item $t_j \in T$ to be $\zeta_{b_i}(t_j) = \text{value}(b_i) - \overline{\ell}_{i,j}$.

Let the allocation where item t_j is given to $\text{buyer}(t_j) = b_i$ be a social welfare maximizing allocation in the unit demand market $M(B,T)$. Also, let $\text{buyer}(t_j) = \emptyset$ if item t_j is unallocated.

This social welfare maximizing allocation in $M(B,T)$ translates into a flow f^* for the discrete passenger-taxicab problem where t_j moves from $\text{loc}(t_j)$ to $\text{loc}(b_i)$ if buyer$(t_j) = b_i$. Ergo,

$$f^*(u,v) = \begin{cases} |\{(i,j)|b_i \in P_v, t_j \in s_u, \text{buyer}(t_j) = b_i\}| & \text{if } u \neq v, \\ |\{(i,j)|b_i \in P_v, t_j \in s_u, \text{buyer}(t_j) = b_i\}| + |\{j|t_j \in s_u, \text{buyer}(t_j) = \emptyset\}|, & \text{if } u = v. \end{cases}$$

Let s' be such that $s'_v = \sum_u f^*(u,v)$ for all $v \in V$. We say that the new supply s' is *induced* by f^*. We now show that

Lemma 4.1.1. *The flow f^* is a min cost flow from s to s'.*

Proof. Assume that f' is a flow from s to s' of strictly lower cost. As f' is an integral flow it can be decomposed into a union of unit flows. This can be interpreted as an alternative allocation in the $M(B,T)$ unit demand market, with strictly higher social welfare. This is in contradiction to our construction. □

Choose the minimal Walrasian prices to clear the unit demand market $M(B,T)$. Such prices are also VCG prices [17]. Let the Walrasian price for item t_j be p_{t_j}. We now define surge prices r_v, $v \in V$, for the discrete passenger-taxicab problem. Specifically, for all $v \in V$, set

$$r_v = \min_{t_j \in T} (\ell_{\text{loc}(t_j),v} + p_{t_j}). \tag{4.1}$$

Lemma 4.1.2. *Assigning t_j to serve passenger* buyer(t_j) *is a social welfare maximizing allocation.*

Proof. First, we show that for any allocation of taxicabs to passengers in the taxicab-passenger setting there exists an allocation of items to buyers in the unit demand market $M(B,T)$ such that the social welfare is the same. Then, we show that for the allocation of items to buyers that maximizes the social welfare in the unit demand market there exists an allocation of taxicabs to passengers with the same social welfare.

Fix an allocation of passengers to taxicabs, *i.e.*, $\Phi : B \to T \cup \{\emptyset\}$ is a matching. Given the matching Φ we define an allocation $\Pi : B \to T \cup \{\emptyset\}$ in the unit demand market where $\Phi(b) = \Pi(b)$ for all $b \in B$.

The social welfare of Φ in the taxicab-passenger setting is $\sum_{b \in B}(\text{value}(b) - \ell_{\text{loc}(b),\text{loc}(\Phi(b))})I_{\Phi(b) \neq \emptyset}$. Similarly, the social welfare of Π in the unit demand market setting is $\sum_{b \in B} \zeta_b(\Pi(b)) = \sum_{b \in B}(\text{value}(b) - \ell_{\text{loc}(b),\text{loc}(\Pi(b))})I_{\Pi(b) \neq \emptyset}$. Since $\Phi(b) = \Pi(b)$ it follows that any allocation in the taxicab-passenger setting has a corresponding allocation in the unit demand market with the same social welfare.

We now show that an allocation in the unit demand market that maximizes social welfare has a corresponding allocation in the passenger-taxicab setting that also maximizes social welfare. Denote the maximal allocation in the unit demand market by $\Pi_{\max} : B \to T \cup \{\emptyset\}$. Define the corresponding matching of passengers to taxicabs by $\Phi_{\max} : B \to T \cup \{\emptyset\}$, where $\Phi_{\max}(b) = \Pi_{\max}(b)$ for all $b \in B$ ($\Phi_{\max}$ is a matching since $\Pi_{\max}$ is a valid allocation in a unit demand market).

Moreover, we need to show that higher valued passengers have priority over lower valued passengers at the same location. *I.e.*, we need to show that for any two passengers, $b_1, b_2 \in B$, such that $\text{loc}(b_1) = \text{loc}(b_2)$ and $\Phi_{\max}(b_1) \neq \emptyset$, $\Phi_{\max}(b_2) = \emptyset$ we have that $\text{value}(b_1) \geq \text{value}(b_2)$.

Contrariwise, assume for some $b_1, b_2 \in B$ we have that $\text{loc}(b_1) = \text{loc}(b_2)$, $\Phi_{\max}(b_1) \neq \emptyset$, $\Phi_{\max}(b_2) = \emptyset$, but $\text{value}(b_1) < \text{value}(b_2)$. Define in the unit demand market then $\Pi' : B \to T \cup \{\emptyset\}$ such that $\Pi'(b_1) = \emptyset$, $\Pi'(b_2) = \Phi_{\max}(b_1)$ and $\Pi'(b) = \Pi_{\max}(b)$ for all $b \notin \{b_1, b_2\}$.

We now show that the social welfare under Π is strictly greater than the social welfare under Π':

$$\begin{aligned}
\sum_{b\in B}(\zeta_b(\Pi'(b))) &= \sum_{b\in B, b\neq b_1,b_2}(\zeta_b(\Pi'(b))) + \zeta_{b_2}(\Phi_{\max}(b_1)) \\
&= \sum_{b\in B, b\neq b_1,b_2}(\zeta_b(\Pi_{\max}(b))) + \mathrm{value}(b_2) - \mathrm{dist}(\mathrm{loc}(b_2), \mathrm{loc}(\Phi_{\max}(b_1))) \\
&> \sum_{b\in B, b\neq b_1,b_2}(\zeta_b(\Pi_{\max}(b))) + \mathrm{value}(b_1) - \mathrm{dist}(\mathrm{loc}(b_1), \mathrm{loc}(\Phi_{\max}(b_1))) \\
&= \sum_{b\in B}(\zeta_b(\Pi_{\max}(b))).
\end{aligned}$$

Thus, Π' has strictly higher social welfare than $\Pi_{\max}$ in unit demand setting in contradiction to $\Pi_{\max}$ maximizing social welfare. Thus, $\Phi_{\max}$ is a valid allocation in the taxicab-passenger setting which maximizes the social welfare. □

Lemma 4.1.3. *For any passenger b_i such that $b_i = \mathrm{buyer}(t_j)$ we have that*

$$\bar{\ell}_{i,j} + p_{t_j} = \min_{t_z\in T}(\bar{\ell}_{i,z} + p_{t_z}) = r_{\mathrm{loc}(b_i)}.$$

Proof. Since $b_i = \mathrm{buyer}(t_j)$, and p are Walrasian prices, we have that buyer b_i maximizes its utility η_{b_i}. Ergo,

$$\begin{aligned}
\eta_{b_i}(t_j) &= \max_{t_x\in T}(\eta_{b_i}(t_x)) \\
&= \max_{t_x\in T}(\mathrm{value}(b_i) - \bar{\ell}_{i,x} - p_{t_x}) \\
&= \mathrm{value}(b_i) - \min_{t_x\in T}(\bar{\ell}_{i,x} + p_{t_x}) \\
&= \mathrm{value}(b_i) - r_{\mathrm{loc}(b_i)}.
\end{aligned}$$

As

$$\eta_{b_i}(t_j) = \mathrm{value}(b_i) - \bar{\ell}_{i,j} - p_{t_j} = \mathrm{value}(b_i) - r_{\mathrm{loc}(b_i)}$$

it follows that

$$\bar{\ell}_{i,j} + p_{t_j} = \min_{t_x\in T}(\bar{\ell}_{i,x} + p_{t_x}).$$

□

Lemma 4.1.4. *Any passenger b_i that is not served is not interested in being served (or is indifferent), i.e., then* $\mathrm{value}(b_i) \leq r_{\mathrm{loc}(b_i)}$. *Any passenger b_i that is served has* $\mathrm{value}(b_i) \geq r_{\mathrm{loc}(b_i)}$.

Proof. Let b_i be some buyer allocated no item in the social welfare maximizing allocation for $M(B,T)$, then it must be that $\max_{t_x \in T} \eta_{b_i}(t_x) \leq 0$. It follows that

$$\max_{t_x \in T} (\mathrm{value}(b_i) - \bar{\ell}_{i,x} - p_{t_x}) \leq 0,$$

and thus

$$\mathrm{value}(b_i) \leq \min_{t_x \in T} (\bar{\ell}_{i,x} + p_{t_x}) = r_{\mathrm{loc}(b_i)}.$$

Consider some buyer b_i that was allocated an item, t_j, in the social welfare maximizing allocation for $M(B,T)$. It follows that $\max_{t_x \in T} \eta_{b_i}(t_x) \geq 0$. Thus,

$$\max_{t_x \in T} (\mathrm{value}(b_i) - \bar{\ell}_{i,x} - p_{t_x}) \geq 0,$$

and

$$\mathrm{value}(b_i) \geq \min_{t_x \in T} (\bar{\ell}_{i,x} + p_{t_x}) = r_{\mathrm{loc}(b_i)}.$$

□

Lemma 4.1.5. *For supply s, demand d, surge prices r, and new supply s' as defined above. A taxicab t_j that serves passenger* $\mathrm{buyer}(t_j)$ *is doing a best response.*

Proof. Consider the following cases:

1. Item t_j is not allocated, *i.e.*, $\mathrm{buyer}(t_j) = \emptyset$. It follows that the Walrasian pricing for item t_j is zero: $p_{t_j} = 0$. Now, for any $w \in V$ we have that

$$r_w = \min_{t_x \in T} (\ell_{w,\mathrm{loc}(t_x)} + p_{t_x}) \leq \ell_{w,\mathrm{loc}(t_j)} + p_{t_j} = \ell_{w,\mathrm{loc}(t_j)},$$

hence, $r_w - \ell_{w,\mathrm{loc}(t_j)} \leq 0$. Ergo, not serving any passenger is a best response for t_j.

2. Item t_j is allocated to some buyer b_i. From Lemma 4.1.3 we know that $\bar{\ell}_{i,j} + p_{t_j} = \min_{t_x \in T}(\bar{\ell}_{i,x} + p_{t_x}) = r_{\text{loc}(b_i)}$ and thus t_j gains a utility of p_{t_j} from serving b_i. If taxicab t_j could serve a passenger at location $w \in V$, it will gain a utility of

$$r_w - \ell_{w,\text{loc}(t_j)} = \min_{t_x \in T}(\ell_{w,\text{loc}(t_x)} + p_{t_x}) - \ell_{w,\text{loc}(t_j)} \le \ell_{w,\text{loc}(t_j)} + p_{t_j} - \ell_{w,\text{loc}(t_j)} = p_{t_j}.$$

Implying that serving passenger b_i is a best response for taxicab t_j.

□

Lemma 4.1.6. *It is a dominant strategy for the passengers to reveal their true valuations given that surge prices are computed via the algorithm above.*

Proof. The utilities of the bidders for the minimal Walrasian prices in a unit demand market coincide with VCG payments [17]. This implies that buyers truthfully reveal their valuations for the items. In our setting the utility for a passenger b_i is exactly equal to the utility for the corresponding bidder b_i. Ergo, misreporting passenger valuation implies misreporting bidder valuations. As misreporting item valuations in the unit demand market setting cannot benefit buyers (and thus passengers) we conclude it is a dominant strategy for passengers to report true valuations. □

To summarize, our main result in this section, Theorem 4.1.7, follows from Lemma 4.1.2, Lemma 4.1.4, Lemma 4.1.5, and Lemma 4.1.6.

Theorem 4.1.7. *For any Profile P and supply s there exist surge prices r, demand $d(r)$ and new supply s' such that*

- *Supply s, new supply s', demand $d(r)$, and surge prices r are in passenger-taxicab equilibrium.*
- *s' is social welfare maximizing with respect to supply s, profile P, and demand d.*

- *The surge prices r can be computed in polynomial time.*
- *It is a dominant strategy for passengers to report their true valuations to the surge-price computation.*

V Optimal Competitive Online Algorithms for Social Welfare

In this section we give online algorithms that determine supply (using surge prices) so as to maximize social welfare as given in Definition 2.3.1. *I.e.*, striking a balance between maximizing the quality of service *vs.* the costs associated with shifting resources about.

The results[1] in this section can be obtained by online algorithms that set the supply to be one of the following:

1. Set supply at time t equal demand at time t, *i.e.*, set $s^t = d^t$.
2. Set supply at time t equal to the supply at time $t-1$, *i.e.*, set $s^t = s^{t-1}$.

It follows from Theorem 3.0.8 that using appropriate surge prices we can determine that $s^t = d^t$ as the unique passenger-taxicab equilibrium. It is easy to leave the supply unchanged by choosing $r_i^t = 1$ for all i. It follows that the resulting passenger-taxicab equilibrium has no positive flow from i to $j \neq i$, as $\ell_{ij} \geq 1$ for all $j \neq i$ — ergo $s^t = s^{t-1}$.

Given a demand sequence d we define ρ as the inverse of the maximum demand at any vertex and time, *i.e.*, $1/\rho = \max_{i,t} d_i^t$. Note that $\rho \leq k$ since at any time t there is a vertex i such that $d_i^t \geq 1/k$. Moreover, $\rho \geq 1$ since $d_i^t \leq 1$, for any time t and vertex i.

Consider the following online algorithms:

[1] These are randomized online algorithms. Alternately, one could give deterministic online algorithms with the same guarantees by using the passenger-taxicab equilibria and surge prices derived from Theorem 3.0.9, with the disadvantages that the equilibria is no longer unique and that this requires some additional technical assumptions.

`rand`(p) — With probability p set surge prices such that supply equals demand at all vertices. *I.e.*, at time $t = 1$ set $s^1 = d^1$; for all $t > 1$ with probability p set $s^t = d^t$ and with probability $1 - p$ set $s^t = s^{t-1}$.

`stay` — Split the supply equally over all vertices. *I.e.*, at time $t = 1$ set $s^1 = \langle \frac{1}{k}, \frac{1}{k}, \ldots, \frac{1}{k} \rangle$ and for all $t > 1$ set $s^t = s^{t-1}$.

`match` — Always set supply equal demand, *i.e.*, set $s^t = d^t$ for all $t \geq 1$. Note that `match` and `rand`(1) are identical.

`composite`(p) — Toss a fair coin, if heads run `stay` otherwise run `rand`(p). The expected social welfare of `composite`(p) satisfies $\mathbb{E}[\texttt{composite}(p)] = \mathbb{E}[\texttt{stay}]/2 + \mathbb{E}[\texttt{rand}(p)]/2$.

In different scenarios different algorithms are useful. We later discuss how to switch between different online algorithms in changing circumstances, varying over time.

Like many other online problems, we first show that the optimal solution can be assumed to be "lazy", never move supply about unnecessarily (Section 5.1). Section 5.2 gives our main technical result. In this setting the cost of moving from one vertex to another always equals 1, i.e., $\ell_{ij} = 1$ for $i \neq j$. In this scenario we show that `composite`$(\sqrt{1/k})$ achieves [an optimal] $\Theta(1/\sqrt{k})$ fraction of the optimal social welfare. More generally, the competitive ratio improves as a function of the maximal demand in a single vertex (a $1/\rho$ fraction of the total demand) — in this setting `composite`$(\sqrt{\rho/k})$ achieves [an optimal] $\Theta(\sqrt{\rho/k})$ fraction of the optimal social welfare. The positive result appears in Theorem 5.2, whereas optimality follows from Lemma 5.3.1.

In Section 5.4 we consider several other scenarios:

- Clearly, even for completely arbitrary costs ℓ_{ij} (to move supply from i to j), algorithm `stay` is trivially ρ/k competitive. In Section 5.4.1 we prove that this cannot be improved. This shows that it is critical that $\ell_{ij} = 1$ to obtain a non-trivial bound, without other assumptions on the input sequence.

- In Section 5.4.2 we consider inputs where the total drift (average total variation distance between successive demand vectors) is small. In such settings the `match` algorithm approaches the optimal social welfare, for sufficiently small drift. Moreover, essentially the same bounds are tight.

5.1 The Optimal Supply Sequence is Lazy

We define lazy sequences and show that without loss of generality the optimal supply sequence is a lazy sequence. We have two types of "non-lazy" actions: increasing supply in a location with supply greater than demand (over supply), or reducing supply in a location while creating over demand. Both actions can be avoided, without loss in social welfare. We start by defining a lazy sequence.

Definition 5.1.1. A supply sequence is *lazy* if for any time t and any $u, v \in V, u \neq v$ such that $f^t(u,v) > 0$ then both (1) $s_v^t \leq d_v^t$ and (2) $s_u^{t-1} > d_u^t$.

We show that for any supply sequence there exists a lazy supply sequence whose social welfare is at least the social welfare of the original sequence.

Lemma 5.1.2. *Fix a demand sequence d. Given an arbitrary supply sequence s, there exists a lazy supply sequence $\bar{s}$ such that* $\mathrm{sw}(\bar{s}) \geq \mathrm{sw}(s)$.

Proof. For contradiction, assume there is a sequence s for which for any lazy sequence $\bar{s}$ we have $\mathrm{sw}(s) > \mathrm{sw}(\bar{s})$. Note that essentially we are saying that there is an optimal sequence s for which no lazy sequence has the same social welfare. This implies that for any optimal sequence s there is a time t such that $f^t(u,v) > 0$ and either (1) $s_v^t > d_v^t$ or (2) $s_u^{t-1} < d_u^t$. Out of all the optimal sequences, consider the optimal sequence s with the largest such time t and largest pair (u,v) (given some full order on the pairs $V \times V$).

We create a new flow $\bar{f}$ depending on the type of violation. Assume that we have $f^t(u,v) > 0$ and $s_v^t > d_v^t$. At time t set $\bar{f}^t(u,v) = f^t(u,v) - \epsilon$ and $\bar{f}^t(u,u) =$

$f^t(u,u)+\epsilon$, where $\epsilon = \min\{s_v^t - d_v^t, f^t(u,v)\}$. The rest of the flow remains unchanged, i.e., $\bar{f}^t(u',v') = f^t(u',v')$ for $(u',v') \neq (u,v)$ or $(u',v') \neq (u,u)$.

At time $t+1$ we adjust the flow to correspond to the original supply. Namely, for all $w \in V$ such that $f^{t+1}(v,w) > 0$, we set $\bar{f}^{t+1}(v,w) = f^{t+1}(v,w)\frac{s_v^t - \epsilon}{s_v^t}$ and $\bar{f}^{t+1}(u,w) = f^{t+1}(u,w) + f^{t+1}(v,w)\frac{\epsilon}{s_v^t}$, and all the remaining flows remain unchanged. It is straightforward to verify that $\bar{f}$ is a valid flow, and we set $s_v^{t+1} = \bar{s}_v^{t+1} = \sum_u \bar{f}^{t+1}(u,v)$.

Note that the only influence on the social welfare are in times t and $t+1$. Comparing the movement cost of $\bar{s}$ to s, at time t it decreased by ϵ and in time $t+1$ increased by at most ϵ. The demand served in $\bar{s}$ and s at time t and $t+1$ in unchanged (since the ϵ flow that was modified did not serve any demand in time t and at time $t+1$ the supplies are identical). This implies that the social welfare of $\bar{s}$ is at least that of s. Therefore we have a contradiction to our selection of t and (u,v).

The case that we have $f^t(u,v) > 0$ and $s_u^{t-1} < d_u^t$ is similar and omitted. $\square$

We derive the following immediate corollary.

Corollary 5.1.3. *Without loss of generality the optimal supply sequence is lazy.*

5.2 Online Algorithms for Social Welfare Maximization when $\ell_{ij} = 1$

We now analyse the lazy optimal supply sequence. We first introduce some notation. Given an optimal lazy supply sequence s, define $h_i^t = \min\{s_i^{t-1}, d_i^t\}$. Let $n \geq 0$ be an integer parameter, and define[2]

$$z_i^t = \max\{0, h_i^t - g_i^t\}, \text{ where } g_i^t = \max_{\tau \in [\max(1,t-n),t-1]} d_i^\tau.$$

Note that the definitions depend on s, but we use a fixed optimal lazy sequence s. Note too that n is yet undetermined.

[2] For notational convenience we define $d_i^t = 0$ and $s_i^t = s_i^1$ for all $t \leq 0$.

Lemma 5.2.1. *Fix a demand sequence d and an optimal lazy supply sequence s for d. The resulting social welfare*

$$\mathrm{opt} = \mathrm{sw}(s,d) = \sum_{t,i} h_i^t \leq \sum_{t,i} z_i^t + \sum_{t,i} g_i^t.$$

Proof. Note that when $\ell_{ij} = 1$ for all i,j we get that $\mathrm{em}(s) = \sum_t \frac{1}{2}\|s^t - s^{t-1}\|_1$. This means that for an optimal lazy sequence we have

$$\mathrm{opt} = \mathrm{sw}(s,d) = \mathrm{ds}(s,d) - \mathrm{em}(s) = \sum_t \sum_i \min(s_i^t, d_i^t) - \sum_t \sum_{i:s_i^t \geq s_i^{t-1}} \left(s_i^t - s_i^{t-1}\right).$$

First consider $s_i^t > s_i^{t-1}$. Since the sequence is lazy and $s_i^t > s_i^{t-1}$ this implies that $s_i^t \leq d_i^t$. Hence, $\min(s_i^t, d_i^t) = s_i^t$ and $\min(s_i^{t-1}, d_i^t) = s_i^{t-1}$. It follows that the identity $\min(s_i^t, d_i^t) - (s_i^t - s_i^{t-1}) = \min(s_i^{t-1}, d_i^t)$ holds.

Next consider $s_i^t < s_i^{t-1}$. Since the sequence is lazy and $s_i^t < s_i^{t-1}$ implies that $s_i^t \geq d_i^t$ and that $\min(s_i^t, d_i^t) = d_i^t = \min(s_i^{t-1}, d_i^t)$. It follows yet again that the identity $\min(s_i^t, d_i^t) = \min(s_i^{t-1}, d_i^t)$ holds.

Combining both identities we have

$$\mathrm{opt} = \mathrm{sw}(s,d) = \sum_t \sum_i \min(s_i^{t-1}, d_i^t) = \sum_t \sum_i h_i^t,$$

by the definition of h_i^t. Since, $h_i^t \leq z_i^t + g_i^t$ the lemma follows. □

Our next goal is to bound the sum of z_i^t and relate it to the social welfare of the algorithm `stay`. We first prove the following properties of the optimal lazy supply sequence.

Lemma 5.2.2. *Fix an optimal lazy sequence s and a parameter $n \geq 1$. If for some i,t we have $s_i^{t-1} \geq \max_{\tau \in [t-n,t)} d_i^\tau$ then we have $\min_{\tau \in [t-n,t)} s_i^\tau \geq s_i^{t-1}$.*

Proof. For contradiction assume there exists some maximal $\tau \in [t-n, t)$ such that $s_i^\tau < s_i^{t-1}$. Then, $\tau \neq t-1$ and thus $\tau + 1 \in [t-n+1, t)$ which by the assumption of the lemma implies that $s_i^{t-1} \geq d_i^{\tau+1}$. Also, because this is the maximal such τ we have that $s_i^{\tau+1} \geq s_i^{t-1}$. Thus, we have $s_i^{\tau} < s_i^{\tau+1}$ and $d_i^{\tau+1} < s_i^{\tau+1}$. This contradicts the assumption that s is an optimal lazy sequence, since there is a flow to i at time $\tau + 1$ which strictly exceeds the demand. $\square$

We derive the following immediate corollary:

Corollary 5.2.3. *Fix an optimal lazy sequence s and a parameter $n \geq 1$. If for some i, t we have $s_i^{t-1} \geq \max_{\tau \in [t-n,t)} d_i^{\tau}$ then for any $\tau \in [t-n+1, t)$ we have $s_i^{\tau-1} \geq s_i^{\tau}$.*

Proof. From Lemma 5.2.2, for any $\tau \in [t-n, t)$ we have that $s_i^{\tau} \geq s_i^{t-1} \geq \max_{\tau' \in [t-n,t)} d_i^{\tau'}$. Therefore, $s_i^{\tau} \geq \max_{\tau' \in [\tau-n',\tau)} d_i^{\tau'}$, where $n' = \tau - (t-n) > 0$. Now applying Lemma 5.2.2 again we obtain the corollary. $\square$

Lemma 5.2.4. *Fix an optimal lazy sequence s and a parameter $n \geq 1$. Then, $\sum_i \sum_{\tau \in [t-n,t)} z_i^{\tau} \leq 1$.*

Proof. Clearly we care only about $z_i^{\tau} > 0$. Fix a location i and let $\tau_1, \ldots, \tau_m$ be all the times $\tau \in [t-n, t)$ for which $z_i^{\tau} > 0$. Clearly, $\sum_{\tau \in [t-n,t)} z_i^{\tau} = \sum_{j=1}^{m} z_i^{\tau_j}$.

First, if $s_i^{t-1} \leq \max_{\tau \in [t-n,t)} d_i^{\tau} = g_i^t$, since $h_i^t \leq s_i^{t-1}$ then $z_i^t = 0$. Therefore, at any time τ_j we have $s_i^{\tau_j - 1} > \max_{\hat{\tau} \in [t-n,t)} d_i^{\hat{\tau}}$, which implies that we can apply Corollary 5.2.3 at the times τ_j.

We claim that $s_i^{\tau_j - 1} > d_i^{\tau_j}$ for $1 \leq j \leq m-1$. For contradiction assume that $s_i^{\tau_j - 1} \leq d_i^{\tau_j}$. We have

$$h_i^{\tau_m} \leq s_i^{\tau_m - 1} \leq s_i^{\tau_j - 1} \leq d_i^{\tau_j} \leq g_i^{\tau_m},$$

where the first inequality is from the definition of h, the second follows from Corollary 5.2.3, the third from our assumption, and the fourth from the definition of g. This implies that $z_i^{\tau_m} = \max\{0, h_i^{\tau_m} - g_i^{\tau_m}\} = 0$. In contradiction to our construction that $z_i^{\tau_m} > 0$. Therefore, $s_i^{\tau_j - 1} > d_i^{\tau_j}$, which implies that $h_i^{\tau_j} = d_i^{\tau_j}$.[3]

[3] This applies only to $j \leq m-1$ since $g_i^{\tau_m}$ does not include $d_i^{\tau_m}$ but does include all previous $d_i^{\tau_j}$.

Since $z_i^{\tau_j} > 0$, we have that $z_i^{\tau_j} = h_i^{\tau_j} - g_i^{\tau_j}$. We showed that $h_i^{\tau_j} = d_i^{\tau_j}$ and $g_i^{\tau_j} \geq d_i^{\tau_{j-1}}$, hence, $z_i^{\tau_j} \leq d_i^{\tau_j} - d_i^{\tau_{j-1}}$, for $2 \leq j \leq m-1$.

Summing over all τ_j we have

$$\begin{aligned}
\sum_{\hat{\tau} \in [t-n,t)} z_i^{\hat{\tau}} &= \sum_{j=1}^{m} z_i^{\tau_j} \\
&= z_i^{\tau_m} + z_i^{\tau_1} + \sum_{j=2}^{m-1} z_i^{\tau_j} \\
&\leq z_i^{\tau_m} + z_i^{\tau_1} + \sum_{j=2}^{m-1} d_i^{\tau_j} - d_i^{\tau_{j-1}} \\
&\leq z_i^{\tau_m} + z_i^{\tau_1} + d_i^{\tau_{m-1}} - d_i^{\tau_1} \\
&\leq h_i^{\tau_m} - (g_i^{\tau_m} - d_i^{\tau_{m-1}}) + (h_i^{\tau_1} - g_i^{\tau_1} - d_i^{\tau_1}) \\
&\leq h_i^{\tau_m}
\end{aligned}$$

For the last inequality note that $g_i^{\tau_m} \geq d_i^{\tau_{m-1}}$ and that $h_i^{\tau_1} \leq d_i^{\tau_1}$.

Summing over all locations i we have

$$\sum_i \sum_{\hat{\tau} \in [t-n,t)} z_i^{\hat{\tau}} \leq \sum_i h_i^{\tau_{m_i}} \leq \sum_i s_i^{\tau_{m_i}} \leq \sum_i s_i^{t-n} = 1$$

where the last inequality uses again Corollary 5.2.3. □

We now analyze `stay` for arbitrary relocation costs ℓ_{ij}.

Lemma 5.2.5. *At all times t, the demand served by* `stay` *is at least ρ/k of the total demand.*

Proof. Recall that $\mathrm{ds}(s^t, d^t) = \sum_i \min(s_i^t, d_i^t) = \sum_i \min\left(\frac{1}{k}, d_i^t\right)$. Denote $S = \left\{i | s_i^t \geq \frac{1}{k}\right\}$. If we have $|S| \geq \rho$ then $\mathrm{ds}(s^t, d^t) \geq \frac{1}{k} \cdot |S| \geq \frac{\rho}{k}$. Otherwise, since $\frac{1}{\rho} \geq \frac{1}{k}$ the total demand not in S is at least $1 - \frac{|S|}{\rho}$ and it is completely served by `stay`. Therefore,

$$\mathrm{ds}\left(s^t, d^t\right) \geq |S| \cdot \frac{1}{k} + 1 - \frac{|S|}{\rho} = \frac{k\rho + |S|\rho - |S|k}{k\rho} = \frac{k\rho - |S|(k-\rho)}{k\rho} \geq \frac{k\rho + \rho^2 - k\rho}{k\rho} = \frac{\rho}{k}.$$

$\square$

Now we analyze `rand`(p) and relate it to g_i^t.

Lemma 5.2.6. *Let $\widehat{s}_i^t$ be the random variable representing the supply of* `rand`(p) *time t in vertex i. Then, $\mathbb{E}[\widehat{s}_i^t] \geq g_i^t p(1-p)^n$. In addition, the expected social welfare of* `rand`(p) *is at least $p(1-p)^n \sum_{i,t} g_i^t$.*

Proof. Let $\tau = \arg\max_{\hat{\tau} \in [t-n,t)} d_i^{\hat{\tau}}$, i.e., $d_i^\tau = g_i^t$. We lower bound the expectation of $\widehat{s}_i^t$ by the probability that `rand`(p) sets $s^\tau = d^\tau$ and keeps the supply until time t, i.e., $s^t = s^\tau$. The probability that we have $s^\tau = d^\tau$ is at least p. The probability that $s^t = s^\tau$ is at least $(1-p)^n$. Therefore, $\mathbb{E}[\widehat{s}_i^t] \geq g_i^t p(1-p)^n$, which implies that the expected social welfare of `rand`(p) is at least $p(1-p)^n \sum_{i,t} g_i^t$. $\square$

Theorem 5.2.7. *The algorithm* `composite`$(\sqrt{\rho/k}) = \frac{1}{2}$`stay` $+ \frac{1}{2}$`rand`$(\sqrt{\rho/k})$ *is $(\frac{1}{2e}\sqrt{\frac{\rho}{k}})$-competitive.*

Proof. By Lemma 5.2.1 we have that $OPT = \sum_{t,i} h_i^t \leq \sum_{t,i} z_i^t + g_i^t$. We bound separately $\sum_{t,i} z_i^t$ and $\sum_{t,i} g_i^t$.

By Lemma 5.2.4 we can partition the time to $\frac{T}{n}$ blocks of size n each, and in each the sum is at most 1, therefore $\sum_{t,i} z_i^t \leq \frac{T}{n}$. On the other hand, `stay` guarantees a social welfare of at least $\rho \cdot \frac{T}{k}$.

We have that,

$$OPT \leq \frac{T}{n} + \sum_{i,t} g_i^t.$$

Using Lemma 5.2.5 and Lemma 5.2.6, we have

$$\frac{1}{2}\texttt{stay} + \frac{1}{2}\texttt{rand}(p) \geq \frac{\rho}{2k}T + \frac{1}{2}p(1-p)^n \sum_{i,t} g_i^t$$

For $p = \sqrt{\frac{\rho}{k}}$ and $n = \frac{1}{p}$ we bound the competitive ratio as follows:

$$\frac{\rho\frac{T}{2k} + \frac{1}{2}p(1-p)^n \sum_{i,t} g_i^t}{\frac{T}{n} + \sum_{i,t} g_i^t} = \frac{\frac{1}{2}\sqrt{\frac{\rho}{k}}T\sqrt{\frac{\rho}{k}} + \frac{1}{2e}\sqrt{\frac{\rho}{k}}\sum_{i,t} g_i^t}{T\sqrt{\frac{\rho}{k}} + \sum_{i,t} g_i^t} \geq \frac{1}{2e}\sqrt{\frac{\rho}{k}}.$$

$\square$

5.3 *Social Welfare Maximization when* $\ell_{ij} = 1$*: Impossibility Results*

We show that no online algorithm can hope to achieve a competitive ratio better (greater) than $O\left(\sqrt{\frac{\rho}{k}}\right)$. Recall, that Section 5.2 describes an online algorithm, `composite`$(\sqrt{\rho/k})$, that achieves this bound on the competitive ratio. Ergo, `composite`$(\sqrt{\rho/k})$ achieves the optimal competitive ratio, up to a constant factor.

Theorem 5.3.1. *Fix the metric* $\ell_{ij} = 1$*. No online algorithm can achieve a competitive ratio better (greater) than* $O\left(\sqrt{\frac{\rho}{k}}\right)$.

Proof. We first describe the proof for $\rho = 1$ and then extend it to arbitrary ρ.

Consider the following stochastic demand sequence. At time t we select at random a vertex $c^t \in V$, and assign all the demand to it, *i.e.*, $d^t_{c^t} = 1$ and $d^t_i = 0$ for $i \neq c^t$. Clearly any online algorithm has an expected social welfare of T/k.

Essentially, for the optimal offline we use the birthday paradox to show that its social welfare is $\Theta(T/\sqrt{k})$. Consider the following offline strategy. Partition the time to intervals of size of $2\sqrt{k}$. We show that in any such interval the offline can increase social welfare by at least 1 with constant probability.

Fix such a time interval. We claim that with constant probability some vertex appears twice in the interval. If in the first $\sqrt{k}$ times there is a vertex i that appears twice, we are done. Otherwise, we have $\sqrt{k}$ distinct vertices. The probability that we resample one of those vertices in the next $\sqrt{k}$ time steps is at least $1/e$. Now, if vertex i appears twice in the interval then the offline algorithm can move at the start of the interval to vertex i and increase social welfare by at least 1. This implies that the expected social welfare of this offline strategy is $\Theta(T/\sqrt{k})$, which lower bounds the expected social welfare of the optimal offline strategy.

Since the online algorithm has expected social welfare of T/k and the optimal offline algorithm has expected social welfare of $\Theta(T/\sqrt{k})$, the competitive ratio, for $\rho = 1$, is bounded by $O(\sqrt{1/k})$.

We now sketch how the proof extends to a general $\rho \geq 1$. In this case we partition the k vertices into $N = \lfloor k/\lceil \rho \rceil \rfloor$ disjoint subsets, each of size $M = \lceil \rho \rceil$. (Note, that $N \cdot M \leq k$.) The N subsets replace the vertices V and each time we select a subset, we give a uniform demand over the subset. (note that the demand per vertex is $1/M \leq 1/\rho$.)

As before, any online algorithm has expected social welfare of $\Theta(T/N) = \Theta(T\rho/k)$. Similar to before, there is an offline strategy that guarantees an expected social welfare of $\Theta(T\sqrt{\rho/k})$. This implies that the competitive ratio is at most $\Theta(\sqrt{\rho/k})$. □

5.4 Extensions

In (Section 5.4.1 we show that the assumption that $\ell_{ij} = 1$ was critical to achieve the non-trivial competitive ratio of Section 5.2 unless ρ (the fraction of demand at any single vertex) was sufficiently small. We also consider restricting the demand sequences by bounding the average variability in demand. In Section 5.4.2 we show that the online algorithm that greedily matches supply and demand works well, the average drift is sufficiently small.

5.4.1 Arbitrary Metric Spaces

We can apply the online algorithm `stay` and guarantee a competitive ratio of ρ/k as shown in Lemma 5.2.5. The following theorem establishes an impossibility result when the costs are different than 1 (even if they are still identical).

Theorem 5.4.1. *Fix some $1 > \epsilon > 0$, and consider costs $\ell_{ij} = 1 + \epsilon$ for $i \neq j$. No online algorithm has a competitive ratio better (greater) than $\frac{(1+\epsilon)^2}{\epsilon} \cdot \frac{1}{k}$ for this metric.*

Proof. The idea is the following: we generate a demand sequence that at every time step demand is concentrated in a single vertex. We generate a random sequence of vertices, such that no two successive positions are identical. We then duplicate every position for a random duration. The duration, the number of successive demands at that position, is geometrically distributed. We set the parameters such that no online algorithm can benefit by switching between vertices. On the other hand, given a sufficiently long duration of repeated demands for the same vertex, the optimal schedule switches to this vertex.

We now describe the stochastic demand sequence generation. We first generate a sequence of locations c. We set $c_1 = i \in V$ uniformly at random. For c_τ we set $c_\tau = j$ where $j \in V \setminus \{c_{\tau-1}\}$ uniformly. In addition we generate a sequence of duration b distributed geometrically with parameter $p = \frac{1}{1+\epsilon}$. Namely, $b_\tau = j$ with probability $p^{j-1}p$, for $j \geq 1$. We are now ready to generate the demand sequence d. For each $c_\tau = i$ we associate a unit vector e_i which has $e_{i,i} = 1$ and $e_{i,j} = 0$ for $j \neq i$. We duplicate e_{c_τ} exactly b_τ times. We truncate the sequence at time T, and this is the demand sequence d.

First consider an arbitrary online algorithm. We claim that it does not gain (in expectation) any social welfare by moving supply, and hence it's expected social welfare is T/k. The argument is that the cost of moving δ supply to a new location is $(1+\epsilon)\delta$. On the other hand, the expected duration in the new location is only $1+\epsilon$, so in expectation there is no benefit. For an online algorithm that does not move any supply the expected social welfare is T/k.

We now analyze the social welfare attained by an optimal offline algorithm. The main benefit of an offline algorithm is that it has access to the realized $b = \tau$. It is simple to see that if $b_\tau \geq 2$ then the offline algorithm has a benefit of $b_\tau - (1+\epsilon) > 0$.

$$\mathbb{E}[b_\tau - (1+\epsilon) | b_\tau \geq 2] \Pr[b_\tau \geq 2] = \sum_{i=2}^{\infty} (\frac{\epsilon}{\epsilon+1})^{i-1} \cdot \frac{i-1-\epsilon}{1+\epsilon} = \frac{\epsilon}{1+\epsilon}.$$

We now would like to sum over τ however the numbers summands in the sum is a random variable. Since we have a random sums of random variables we need to use Wald's identity. Since the expected number of summands is $\frac{T}{1+\epsilon}$ and the expectation of each is $\frac{\epsilon}{1+\epsilon}$ we have that the optimal offline algorithm has an expected social welfare of at least $\frac{\epsilon}{(1+\epsilon)^2}T$.

This implies that no algorithm has a competitive ratio better than $\frac{(1+\epsilon)^2}{\epsilon}\frac{1}{k}$. □

5.4.2 *Restricted Drift*

For any demand sequence d let $\delta \leq 1$ be the average drift, *i.e.*, $\sum_t \|d^t - d^{t-1}\|_{tv} = (1/2)\sum_t \|d^t - d^{t-1}\|_1 = \delta T$.

Theorem 5.4.2. *For the case where costs $\ell_{ij} = 1$ for all $i \neq j$, setting demand and supply equal (the* `match` *algorithm) gives social welfare of $(1-\delta)T$, and is $(1-\delta)$-competitive.*

For arbitrary ℓ_{ij}, where $\ell_{ij} \leq \ell_{\max}$, the `match` *algorithm has social welfare of at least $(1 - \delta\ell_{\max})T$, and is $(1 - \delta\ell_{\max})$-competitive.*

Proof. Since for $\ell_{ij} = 1$ the earthmover distance metric coincides with the total variation metric, we have that at time t the social welfare of `match` is $1 - \|d^t - s^{t-1}\|_{tv} = 1 - \|d^t - d^{t-1}\|_{tv}$ since `match` sets $s^{t-1} = d^{t-1}$. Summing over all time steps we get that the social welfare of `match` is $T - \delta T$. Since the social welfare of opt is at most T we have that `match` is $(1-\delta)$-competitive.

For a general metric, note that $\mathrm{em}(d^t, d^{t-1}) \leq \ell_{\max}\|d^t - d^{t-1}\|_{tv}$. This implies that the social welfare of `match` is at least $(1 - \ell_{\max}\delta)T$, and hence it is $(1 - \ell_{\max}\delta)$-competitive. □

Theorem 5.4.3. *For the metric $\ell_{ij} = 1$, no online algorithm has a competitive ratio better (greater) than $1 - \delta/4$.*

Proof. Consider the following demand sequence. The demand sequence uses only the first two locations, *i.e.*, for all locations $i \neq 1, 2$ and times t we have $d_i^t = 0$. For each time t we select the demand randomly from the following distribution.

$$d^t = \begin{cases} d_1^t = 1, d_2^t = 0 & \text{With probability } \frac{1}{2} \\ d_1^t = 1 - 2\delta, d_2^t = 2\delta & \text{With probability } \frac{1}{2} \end{cases}.$$

The generated sequence has an expected drift of δT. Any online algorithm ALG has, in expectation, social welfare of $(1-\delta)T$. The main point is that opt has a strictly better expected social welfare.

Consider the online algorithm `match` as a starting point. Partition the time to $T/2$ pairs of time slots, $[2m-1, 2m]$. Consider the event that $d^{2m-2} = d^{2m} \neq d^{2m-1}$. This event occurs with probability 1/4. In such an event we can modify `match` and at time $2m-1$ set $s^{2m-1} = d^m$. (This requires knowing the future, but we are interested in opt so it is fine.) Such a modification increases the social welfare by 2δ (lowering the serviced demand by 2δ and lowering the movement costs by 4δ). Therefore, the expected social welfare is improved by $(1/4)(2\delta)(T/2)$. This implies that the expected social welfare of opt is at least $(1-(3/4)\delta)T$.

This means that no algorithm is more than $\frac{1-\delta}{1-(3/4)\delta}$-competitive. This implies that no online algorithm can have a competitive ratio better than $(1-\delta/4)T$.

$\square$

VI Discussion

Social welfare in our setting depends on the taxicabs and their locations (the supply s), passengers, their locations and values (the profile P), and distances between taxicabs and passengers. In this thesis we introduce passenger-taxicab equilibria, prove their existence and give poly time algorithms for computing surge prices so as to maximize social welfare.

We have shown that although time series are a critical part of the social welfare gains of any taxicab provider, no algorithm can hope to achieve significant worst-case ratios. Thus, in the future different relaxations to the problem might be considered in order to allow for more adaptive algorithms.

When computing the surge prices above, we have implicitly assumed that taxicab locations are known (e.g., via GPS). Contrawise, passengers have no incentive to misreport their location (trivially) and valuation (as proved above). An interesting variation on our models would be to consider taxicabs declaring their own distances to passengers. Those would not be physical distances but rather a personalized cost for service at a given location.

If such personalized costs are verifiable, and social welfare is redefined as the sum of passenger values served minus the personalized service costs, then the surge prices computed in this thesis maximize this new social welfare. This allows for more robust pricing mechanisms which allow us to incorporate issues such as "start up costs" which are a bonus for drivers to get out of bed.

Taxicab personalized costs are private to the taxicab. Thus, any surge price computation would have to contend with private values of the taxicabs as well as private values for the passengers. It is easy to see that without Bayesian assumptions on the private values, little can be done. Just consider a passenger and a taxicab at the same location, they need to agree upon a price. In the Bayesian setting this is called the bilateral trading problem and there is a rich literature on the topic.

Bibliography

[1] R. Arnott. Taxi travel should be subsidized. *Journal of Urban Economics*, 40(3):316 – 333, 1996.

[2] S. Banerjee, C. Riquelme, and R. Johari. Pricing in ride-share platforms: A queueing-theoretic approach. 2015.

[3] T. Berger, C. Chen, and C. B. Frey. Drivers of disruption? estimating the uber effect. Technical report, 2017.

[4] A. Borodin, N. Linial, and M. E. Saks. An optimal on-line algorithm for metrical task system. *Journal of the ACM*, 39(4):745–763, 1992.

[5] A. Braverman, J. Dai, X. Liu, and L. Ying. Empty-car routing in ridesharing systems. *arXiv preprint arXiv:1609.07219*, 2016.

[6] C. Camerer, L. Babcock, G. Loewenstein, and R. Thaler. Labor supply of new york city cabdrivers: One day at a time. *The Quarterly Journal of Economics*, 112(2):407–441, 1997.

[7] J. C. Castillo, D. Knoepfle, and G. Weyl. Surge pricing solves the wild goose chase. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, pages 241–242, New York, NY, USA, 2017. ACM.

[8] M. K. Chen. Dynamic pricing in a labor market: Surge pricing and flexible work on the uber platform. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 455–455, New York, NY, USA, 2016. ACM.

[9] M. K. Chen. Dynamic pricing in a labor market: Surge pricing and flexible work on the uber platform. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 455–455, New York, NY, USA, 2016. ACM.

[10] I. R. Cohen, A. Eden, A. Fiat, and L. Jez. Pricing online decisions: Beyond auctions. In P. Indyk, editor, *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 73–91. SIAM, 2015.

[11] P. Cohen, R. Hahn, J. Hall, S. Levitt, and R. Metcalfe. Using big data to estimate consumer surplus: The case of uber. Working Paper 22627, National Bureau of Economic Research, September 2016.

[12] J. Cramer and A. B. Krueger. Disruptive change in the taxi business: The case of uber. *American Economic Review*, 106(5):177–82, May 2016.

[13] A. Fiat, Y. Rabani, and Y. Ravid. Competitive k-server algorithms. In *Proceedings [1990] 31st Annual Symposium on Foundations of Computer Science*, pages 454–463 vol.2, Oct 1990.

[14] F. Gul and E. Stacchetti. Walrasian Equilibrium with Gross Substitutes. *Journal of Economic Theory*, 87(1):95–124, July 1999.

[15] A. R. Karlin, M. S. Manasse, L. Rudolph, and D. D. Sleator. Competitive snoopy caching. *Algorithmica*, 3(1):79–119, 1988.

[16] C. T. Lam and M. Liu. Demand and consumer surplus in the on-demand economy: The case of ride sharing. 2017.

[17] H. B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461–79, 1983.

[18] H. Ma, F. Fang, and D. C. Parkes. Spatio-temporal pricing for ridesharing platforms. *arXiv preprint arXiv:1801.04015*, 2018.

[19] M. S. Manasse, L. A. McGeoch, and D. D. Sleator. Competitive algorithms for server problems. *Journal of Algorithms*, 11(2):208 – 230, 1990.

[20] C. J. Martin. The sharing economy: A pathway to sustainability or a nightmarish form of neoliberal capitalism? *Ecological Economics*, 121:149 – 159, 2016.

[21] L. Richardson. Performing the sharing economy. *Geoforum*, 67:121 – 129, 2015.

[22] D. D. Sleator and R. E. Tarjan. Amortized efficiency of list update and paging rules. *Communications of the ACM*, 28(2):202–208, 1985.

[23] D. D. Sleator and R. E. Tarjan. Self-adjusting binary search trees. *Journal of the ACM*, 32(3):652–686, 1985.

[24] J. G. Wardrop and J. I. Whitehead. Correspondence. some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers*, 1(5):767–768, 1952.

[25] C.-L. Xin and W.-M. Ma. Scheduling for on-line taxi problem on a real line and competitive algorithms. In *Proceedings of 2004 International Conference on Machine Learning and Cybernetics (IEEE Cat. No.04EX826)*, volume 5, pages 3078–3083 vol.5, Aug 2004.

[26] H. Yang, S. Wong, and K. Wong. Demand–supply equilibrium of taxi services in a network under competition and regulation. *Transportation Research Part B: Methodological*, 36(9):799 – 819, 2002.

תקציר

שירותים שיתופיים נמצאים בעליה משמעותית בשנים האחרונות ברחבי העולם – שירותים כגון uber ו-lyft מאפשרים לנהגים להסיע נוסעים ממקום למקום. היתרון של השירותים האלה נעוץ בגמישות שהוא מספק הן לנהגים והן לנוסעים. הנהגים מסוגלים לעבוד בשעות כרצונם ובכך מאפשר להם גמישות רבה. הנוסעים לעומתם מקבלים שירות ברמה גבוהה ומסוגלים לבקש נסיעה בכל זמן ולקבלה בסיכוי גבוה.

על מנת לאזן את כמות הנוסעים והנהגים בזמנים ומקומות שונים השירותים משתמשים במחירים דינאמיים – בשעות שבהן יש יותר נוסעים מנהגים השירות יעלה את המחירים באיזור (בלעז surge pricing) ובכך יוריד את כמות הנוסעים שיחפצו בנסיעה וכן ימשוך יותר נהגים לאיזור. אך כיצד יחליט השירות מה יהיה המחיר בכל מקום ורגע נתונים?

נתעניין בשני מודלים שונים לייצוג הבעיה (בשני המודלים המוניות משלמות את המרחק ביניהן לבין הנוסע על מנת להגיע אליו ולאפשר את הנסיעה) – אחד רציף אך מתעלם מהשערוך של הנוסעים עבור הנסיעה והשני בדיד ובו לכל נוסע יש שערוך לנסיעה. כמו כן, נבחן את יעילותם של אלגוריתמים מקוונים (online) עבור המודל הרציף.

עבודה זו מציגה תוצאות עבור כל אחת מהבעיות –

- במודל הרציף – נראה אלגוריתם ביעילות פולינומיאלית על מנת למצוא מחירים כך שכל הנוסעים יקבלו שירות. זאת תחת ההנחה שכל מונית תפעל בצורה שתשיג עבורה את הרווחים המקסימליים תחת אותם מחירים. בנוסף, אנחנו מראים שזהו המצב המאוזן היחיד האפשרי עבור המחירים שאנו מחשבים.
- במודל הבדיד – נראה אלגוריתם ביעילות פולינומיאלית על מנת למצוא מחירים כך שהמצב שבו המערכת נמצאת בתועלת המקסימלית הוא מצב יציב. התועלת של המערכת היא סכום השערוכים של הנוסעים אשר קיבלו שירות פחות סכום המרחקים (התשלום של המוניות). כמו כן, אנחנו נראה שתחת המחירים כל אחת מהמוניות מגיבה בצורה האופטימלית עבורה וכן האסטרטגיה האופטימלית עבור כל נוסע היא להציג למערכת את השערוך האמיתי שלו.
- במודל המקוון – נראה שגם בהינתן היכולת לקבוע מחירים, לא נוכל להשיג יותר מ- $O(\frac{1}{k})$ מהתועלת של האלגוריתם האופטימלי ללא הגבלות על המרחקים בין המיקומים, כאשר k הוא מספר המיקומים במודל ו-$\frac{1}{\rho}$ הוא המספר המקסימלי של נוסעים בכל מיקום. נראה אלגוריתם טרוויאלי שמשיג חסם זה. בנוסף, כאשר כל המרחקים שווים ל-1 לא נוכל להשיג יותר מ- $O\left(\sqrt{\frac{\rho}{k}}\right)$ תועלת מאשר האלגוריתם האופטימלי. נראה אלגוריתם שמשיג את החסם הזה.

אוניברסיטת תל-אביב

הפקולטה למדעים מדויקים ע"ש ריימונד ובברלי סאקלר

בית הספר למדעי המחשב ע"ש בלבטניק

איזוני נוסעים-מוניות מקוונים באמצעות מחירים דינאמיים

חיבור זה הוגש כעבודת מחקר לקראת התואר "מוסמך אוניברסיטה" במדעי המחשב

על ידי

ליאור שולץ

בהנחייתם של פרופ' עמוס פיאט ופרופ' ישי מנצור

ניסן התשע"ח