

Pricing Online Decisions: Beyond Auctions

Ilan Reuven Cohen¹, Alon Eden¹, Amos Fiat¹ and Łukasz Jeż^{1,2}

¹Blavatnik School of Computer Science, Tel-Aviv University*

²University of Wrocław, Institute of Computer Science[†]

{alonarden, ilanrcohen}@gmail.com, fiat@tau.ac.il, lje@cs.uni.wroc.pl

October 8, 2014

“You pays your money and you takes your choice.”

— Mark Twain, *Huckleberry Finn*

Abstract

We consider dynamic pricing schemes in online settings where selfish agents generate online events. Previous work on online mechanisms has dealt almost entirely with the goal of maximizing social welfare or revenue in an auction settings. This paper deals with quite general settings and minimizing social costs. We show that appropriately computed posted prices allow one to achieve essentially the same performance as the best online algorithm. This holds in a wide variety of settings. Unlike online algorithms that learn about the event, and then make enforceable decisions, prices are posted without knowing the future events or even the current event, and are thus inherently dominant strategy incentive compatible.

In particular we show that one can give efficient posted price mechanisms for metrical task systems, some instances of the k -server problem, and metrical matching problems. We give both deterministic and randomized algorithms. Such posted price mechanisms decrease the social cost dramatically over selfish behavior where no decision incurs a charge. One alluring application of this is reducing the social cost of free parking exponentially.

1 Introduction

We consider scenarios where agents arrive and make decisions which influence the global environment. One such setting is that of online auctions, where goods or combinations of goods are sold over time. In this paper we

consider a more general setting that captures more general online social choice issues.

We suggest using dynamic posted prices that dynamically assign a surcharge to every decision. Under the assumption that agents are rational, this allows one to prove approximation bounds with respect to optimal target function.

The competitive analysis of online algorithms typically deals with problems cast as follows:

- A sequence of events $\sigma_1, \sigma_2, \dots$ occurs over time.
- Subsequent to observing event σ_i , a centralized online algorithm takes some decision $s_i \in S_i$, where S_i is the set of possible decisions that can be taken in response to event σ_i .
- Centralized online algorithms choose their decisions for event i without knowing what the future holds so as to maximize (or minimize), some function $f(\bar{\sigma}, \bar{s})$ of the sequence of events $\bar{\sigma} = \sigma_1, \dots, \sigma_n$ and associated decisions $\bar{s} = s_1, \dots, s_n$.

In this paper we reinterpret the framework above by having strategic agents initiate events. So rather than an online sequence of events we have an online sequence of strategic agents, and rather than a centralized online algorithms making decisions the strategic agents make the decisions themselves.

The i th agent, has some (private) type $t_i \in T_i$. Conditioned upon previous events $\sigma_1, \dots, \sigma_{i-1}$, and previous decisions $s_1 \in S_1, \dots, s_{i-1} \in S_{i-1}$, an agent type $t \in T_i$ is a function $t : S_i \mapsto \mathbb{R}$. An agent of type $t \in T_i$ has value $t(s)$ for decisions $s \in S_i$. Ergo, a rational agent i will choose decision

$$s_i \in \arg \max_{s \in S_i} t_i(s).$$

We remark that t may take either positive or negative values. In the translation of online problems to the strategic setting, costs translate to negative value, benefits

*Partially supported by the Israeli Centers of Research Excellence (I-CORE) program, (Center No.4/11) and the Google Inter-University Center.

[†]Partially supported by the Foundation for Polish Science (FNP) START Scholarship and the Polish National Science Center (NCN) Grant DEC-2013/09/B/ST6/01538.

translate to positive values. Note that the target function f remains as defined in the non-strategic online setting.

To motivate these strategic agents to “do the right thing” (from the perspective of the designer), we introduce *dynamic posted pricing schemes* that set a surcharge on the set of possible decisions. A dynamic posted pricing scheme P generates a sequence of functions $\langle P_1, P_2, \dots \rangle$. Dynamic posted pricing schemes determine the i 'th function, P_i , before event σ_i arrives.

The function $P_i : S_i \mapsto \mathbb{R}$ sets posted prices (a surcharge) for every possible decision in S_i . We assume quasilinearity, *i.e.*, the utility of agent i of type $t \in T_i$ — should she decide upon $s \in S_i$ — is

$$u_i(s) = t(s) - P_i(s).$$

Rational agents always seek to maximize their utility, or minimize their disutility which is the negation of their utility.

Dynamic posted prices of particular interest are varying tolls on express lanes and the SFPark system (SFPark.org) which sets parking prices in San Francisco as a function of parking congestion [22, 6]. Interestingly, our dynamic posted pricing scheme for metric matching on the line described below can be viewed as a refinement of the SFPark system.

A very simple example for dynamic pricing, where prices are set once and for all (all P_i 's are the same), is the problem of online matching on unweighted bipartite graphs [13]. In this problem, there is a stationary right side, vertices (agents) on the left arrive online along with edges to all adjacent vertices on the right. The target is to maximize the number of matches made.

The KVV algorithm [13] performs a random permutation of the vertices on the right and assigns an incoming vertex to the first available vertex on the right. Translating this into the strategic agent setting, agent types give a utility of one for a match with one of the adjacent vertices. Strategic agents can be induced to emulate the KVV algorithm by a pricing scheme assigning a surcharge of $\varepsilon, 2\varepsilon, \dots$ to the vertices on the right in some random order. Note that the surcharge for a vertex never changes so there is no “dynamic” aspect to this pricing scheme.

Posted prices are inherently a dominant strategy incentive compatible mechanism. Given a dynamic posted pricing scheme P that produces functions $\langle P_1, P_2, \dots \rangle$, and a sequence of agent types $\bar{t} = t_1, t_2, \dots, t_n$, the set of decision sequences in equilibria is

$$\text{eq}(P, \bar{t}) = \left\{ \bar{s} \mid \bar{s} = s_1, \dots, s_n, \quad s_i \in \arg \max_{s \in S_i} u_i(s) \right\}.$$

Because of ties it may be that $|\text{eq}(P, \bar{t})| > 1$.

In the context of mechanism design it is often assumed that if a mechanism is dominant strategy truthful,

the agents may as well reveal their true type to the mechanism. Moreover, the revelation principle states that without loss of generality, one need consider only such direct revelation mechanisms.

In many contexts, revealing the agent type is problematic or absurd. When going to the supermarket one may observe what people are buying, it is hardly plausible to interview customers at length as to what their true preferences are. Moreover, revealing the true agent type may be a grave breach of privacy, see [20].

One could, in principle, combine a dynamic posted pricing scheme yet insist that the agents reveal their true type. Were we to allow this, the functions P_i could depend on the true types of previous agents: t_1, \dots, t_{i-1} , and on previous decisions taken by these agents: s_1, \dots, s_{i-1} (may be relevant if agents use randomization).

To avoid the unnatural aspects of type revelation, we like to restrict the information available to the dynamic pricing scheme. This should be information that is (relatively) easy to collect by simple observation of the consequence of agent decisions without having to understand the causes for such decisions. *E.g.*, one can observe what customers purchase, but not what they plan to cook for dinner or whom they are having over. We would like to allow dynamic pricing schemes¹ to observe (some) consequences of decisions made by (earlier) agents, but not necessarily allow access to the true agent types.

The less one assumes about observable history, the more robust the concept. *E.g.*, in the context of dynamic posted pricing for parking, it seems reasonable to place sensors at parking spaces so as to learn what parking spaces are free but it seems much less plausible to require a host of miniaturized spy drones to follow everyone about (so as to learn the true agent type — her true destination). See discussion and definitions in Section 6.

In the context of strategic agents, we use the term price of anarchy [17] instead of the term competitive ratio. *I.e.*, under the assumption that agents are rational, the price of anarchy is the ratio between the value of the target function f when making the best possible decisions and the value of f when decisions are the worst case equilibria.

We distinguish between benefit problems, where the goal is to maximize f , and cost problems where the goal is to minimize f . In particular, social welfare maximization is to maximize the function $\sum_{i=1}^{|\sigma|} t_i(s_i)$, and social cost minimization is to minimize the social cost function $-\sum_{i=1}^{|\sigma|} t_i(s_i)$. Clearly the goals are equivalent.

Given a dynamic pricing scheme \bar{P} , target function f , the price of anarchy (for benefit problems) is

$$(1.1) \quad \min_{\bar{\sigma}, \bar{s} \in \text{eq}(\bar{P}, \bar{\sigma})} \frac{f(\bar{\sigma}, \bar{s})}{f(\bar{\sigma}, \text{opt}(\bar{\sigma}))},$$

¹Or mechanisms in general.

where

$$\overline{\text{opt}(\sigma)} \in \arg \max_{\bar{s}: s_i \in S_i \text{ for all } i} f(\bar{\sigma}, \bar{s}).$$

Given a dynamic pricing scheme \bar{P} , and target function f , the price of anarchy (for cost problems) is

$$(1.2) \quad \max_{\bar{\sigma}, \bar{s} \in \text{eq}(\bar{P}, \bar{\sigma})} \frac{f(\bar{\sigma}, \bar{s})}{f(\bar{\sigma}, \overline{\text{opt}(\sigma)})},$$

where

$$\overline{\text{opt}(\sigma)} \in \arg \min_{\bar{s}: s_i \in S_i \text{ for all } i} f(\bar{\sigma}, \bar{s}).$$

It is common to consider the competitive ratio in terms of an adversary, and one can do the same in the context of the price of anarchy. With this interpretation, one may assume that the adversary knows the true types of the entire sequence in advance.

To clarify the differences between online algorithms, online mechanisms, and online posted prices (a special case of online mechanisms) the following summary may be helpful:

1. Online algorithms: events appear and are dealt with in a centralized fashion. It is assumed that events are “true”. Online algorithms enforce their decisions. There is no sense in which an event says “I don’t like what you’ve decided in my case.”
2. Online mechanisms: events appear and decisions are made, prizes and fees set, in a centralized fashion. Just as with competitive algorithms, online mechanisms can (magically?) enforce their decisions. There is no sense in which an event can “revolt” against a decision that she does not like.
3. Dynamic posted pricing (this paper) sets a posted price for every possible decision. Incoming strategic events make their own decisions based upon their own self interest. “You pays your money and you makes your choice”. Moreover, different dynamic posted pricing schemes can be compared both with respect to the competitive ratio they achieve and with respect to the quality of the observable history they require.

1.1 Problems and Results We consider the following problems

Metrical task systems [3] Metrical task systems are described by a symmetric non-negative $m \times m$ matrix, $(d_{st})_{1 \leq s, t \leq m}$, where $d_{ss} = 0$ and the triangle inequality holds: $d_{st} \leq d_{sk} + d_{kt}$ for all $1 \leq s, t, k \leq m$. The set $S = \{1, \dots, m\}$ is the set of states. A task is a vector of m non negative real values,

$$w = (w_1, w_2, \dots, w_m).$$

In response to the arrival of a new task, w^i , a centralized online algorithm decides upon some state $s \in S$. The cost of a sequence of decisions s_1, \dots, s_n for a sequence of tasks w^1, \dots, w^n is

$$\sum d_{s_{i-1}s_i} + w_{s_i}^i.$$

Strategic metrical task systems The [dis]utility of an agent is the sum of transition cost (to go from state to state), the cost of executing the task in the state chosen, and the surcharge for the decision taken. It is very easy to see that, without introducing a surcharge, the price of anarchy is unbounded. Strategic agents will never take expensive transitions. If one were to consider an online algorithm as a mechanism without money, agents will lie shamelessly. Every state other than their greedy preference would cost infinity.

Here we give a dynamic posted pricing scheme with a competitive ratio of $O(m)$, within $O(1)$ of the best deterministic competitive ratio possible. The observable history consists of the actual state transitions and the cost to serve the tasks in the states chosen. The observable history *does not* include the cost of the tasks in states that were not chosen.

The k -server problem [19, 16, 15] The k server problem has k servers located in some metric space. Online requests arrive following which some server is to be moved to the request location. *I.e.*, one needs to decide upon which server is to be moved. The goal is to minimize the sum of distances traveled. A competitive ratio of k is known for some metric spaces, whereas $2k - 1$ is an upper bound for general metric spaces [16], *vs.* a lower bound of k . Closing the gap for deterministic algorithms is an infamous open problem.

The strategic k -server problem The disutility of an agent is the distance traversed by the server. Thus, without a dynamic posted pricing scheme, agents will always choose the closest server to move, giving an unbounded price of anarchy. We give a dynamic posted pricing scheme for any metric space with 2 servers with a price of anarchy 10. For k -servers on the line, we give the optimal price of anarchy k . In this case, the observable history (the server positions) allows one to infer the true type of previous agents.

Metric Matching [14, 12] The metric matching problem is described on some metric space. Events are points in the metric space, who are to be matched to some unmatched point, with the goal of minimizing the sum of distances between the incoming points and their match. One application of this problem is parking, where one seeks to assign homeowners a parking spot close to their

home. The problem on the line is another infamous open problem, conjectured to have a constant competitive ratio in [12] but still open.

Strategic Metric Matching The disutility of an agent is the distance between the agent destination and the point assigned to the agent. Ergo, strategic agents will always choose the closest possible point.

It follows from [14] that without surcharges the price of anarchy may be exponential in the size of the metric space.

We give a randomized dynamic posted pricing scheme with a competitive ratio logarithmic in the ratio between the largest and smallest distances between points. In particular, this implies an exponential improvement in the price of anarchy over the price of anarchy without surcharges. Moreover, given an estimate on the cost of the optimal solution, we can attain a doubly exponential improvement over the price of anarchy without surcharges.

2 Related Work

Lavi and Nisan [18] initiated the study of online auctions. They consider a model of indivisible goods where bidders have a privately known valuation for different quantities of the goods. Agents learn their valuation at some time and must bid immediately. The auction must decide immediately how many units to allocate and at what price. They show that the only dominant strategy incentive compatible mechanisms are those that offer a price menu for different quantities.

Awerbuch, Azar, and Myerson [1] give a general scheme that produces posted prices for general combinatorial auctions, with a competitive ratio equal to the logarithm of the ratio between highest and lowest prices, times the underlying competitive ratio for the combinatorial auction (which is terrible, in general).

In [7] Friedman and Parkes introduced a model of online mechanism design, they seek strategyproof mechanisms where agents may manipulate their arrival time and values. They give a variant of the VCG mechanism for some such problems. One of the issues from this paper and subsequent work has been misrepresentation of arrival times (or departure times, or both).

In the main, we do not consider misrepresentation with respect to time or position in the sequence. See discussion in Section 6

In Parkes [21] a general formalism for dynamic environments and online mechanism design is given. In this general model a mechanism makes (and enforces) a sequence of decisions. The decisions may depend on interaction with agents as well as changes to the environment. In the online mechanism design model, agent types include arrival times, departure times, and a valuation function of the decisions made by the algorithm.

Unfortunately, the positive results in the general setting of [21] have been limited to so-called “single valued preferences” where an agent gets a fixed value r iff one of a set of “interesting decisions” is made between her arrival and departure. (The positive results also require limited misrepresentation of the agent types).

It is important to note that none of the problems we consider are single valued. For task systems, the k -server problem, and the parking problem, different decisions will generally give different value to the agent. Also note that we do not allow misrepresentation of arrival time, and insist that departure time be equal to arrival time.

There has been a large body of work on online auctions in the setting where the order of events is a uniformly random permutation. Hajiaghayi, Kleinberg and Parkes [9] discuss limited supply online auctions, and Babaioff, Immorlica, Kempe, and Kleinberg [2] introduced the knapsack and matroid secretary problems, about which there has been much subsequent work.

3 Pricing States for Metrical Task Systems

3.1 Problem Definition A metrical task system [3] is described by a symmetric non-negative $m \times m$ matrix, $(d_{st})_{1 \leq s, t \leq m}$, where $d_{ss} = 0$ and the triangle inequality holds: $d_{st} \leq d_{sk} + d_{kt}$ for all $1 \leq s, t, k \leq m$. The set $S = \{1, \dots, m\}$ is the set of states.

The strategic metrical task system is a variant of the metrical task system where tasks are associated with agents. Agents arrive online, the i th agent to arrive has some task to perform, represented by a vector of m non negative real values,

$$w^i = (w_1^i, w_2^i, \dots, w_m^i).$$

The task w^i is private to agent i .

The initial state of the system is s_0 , the state of the system may change over time. Let s_{i-1} be the state of the system upon the arrival of agent i . The i th agent to arrive may decide to change the state of the system upon her arrival to some other state (*i.e.*, $s_i \neq s_{i-1}$) or not (*i.e.*, $s_i = s_{i-1}$).

Given states s , t , and a task $w = (w_1, w_2, \dots, w_m)$, define

$$C(s, t, w) = w_t + d_{st},$$

i.e. the cost to switch states from s to t and process task w in state t .

A dynamic pricing scheme for strategic metrical task systems sets prices (surcharges) on the states of the system. Formally, a dynamic pricing scheme $P = (P_1, P_2, \dots)$ is a sequence of price functions P_1, P_2, \dots , where the i th function

$$P_i : \langle (s_1, w_{s_1}^1), (s_2, w_{s_2}^2), \dots, (s_{i-1}, w_{s_{i-1}}^{i-1}) \rangle \times S \mapsto \mathbb{R}^+.$$

The information available to the pricing function P_i is the state s_j chosen made by agent j and the work

$w_{s_j}^j$ expended on the j th agent's task w^j in that state, for all $j < i$. This is captured by the sequence $\langle (s_1, w_{s_1}^1), (s_2, w_{s_2}^2), \dots, (s_{i-1}, w_{s_{i-1}}^{i-1}) \rangle$. The prices computed by pricing function P_i are relevant for agent i .

For notational convenience, we use P_i as a pricing function over the set of states after $i - 1$ tasks are processed, where the "history" of the system is implicit. We stress that the pricing is set without any information regarding future events.

Note that pricing function P_i cannot depend on tasks w^j , $j \geq i$, (this is distinct from an online algorithm that may depend on w^j). Moreover, it knows very little about w^1, w^2, \dots, w^{i-1} . It only knows the decisions made by the agents and the actual work done, i.e., s_j and $w_{s_j}^j$ for $1 \leq j \leq i - 1$. In particular, this means that the pricing scheme does not know the optimal cost to process tasks w^1, \dots, w^{i-1} and cannot compute the work function for a state.

The goal of a strategic agent is to minimize her disutility. This is the sum of the following components:

1. The work required to do the task in the state selected by the agent. For agent i this is $w_{s_i}^i$.
2. The transition cost to change states. For agent i this is $d_{s_{i-1}s_i}$.
3. The surcharge associated with choosing (or remaining in) some state, this is determined by the pricing scheme (see below). For agent i this is $P_i(s_i)$.

Therefore, the agent will choose:

$$s_i \in \arg \min_{s \in S_i} \{w_s^i + d_{s_{i-1}s} + P_i(s)\}$$

As the goal in metrical task systems is cost minimization, the target function f is the social cost,

$$f(\bar{\sigma}, \bar{s}) = \sum_{i=1}^{|\bar{\sigma}|} (w_{s_i}^i + d_{s_{i-1}s_i}),$$

i.e., the sum of all state transitions and task processing costs. The price of anarchy is as defined in Equation (1.2).

3.2 The Fractional Traversal Algorithm for Metrical Task Systems We now describe the fractional traversal algorithm for metrical task systems from [3]. The fractional traversal algorithm solves a relaxation of the original problem in that it allows states to perform fractions of a task, possibly switching states repeatedly until the task is completed.

A traversal sequence, $\tau = \tau_1, \tau_2, \dots$, is an infinite sequence of states of the task system, where every state of the task system may appear several times in the sequence. The fractional traversal algorithm given here as

Algorithm 1 is very simple, it executes a fraction of task until the work done in the current state, since moving to this state, equals the distance to move to the next state in the sequence.

Borodin et al. showed the following with regard to the traversal algorithm:

THEOREM 3.1. (BORODIN ET AL. [3]) • *There exists a traversal sequence such that the fractional traversal algorithm has a competitive ratio of $8(m - 1)$. (Versus the optimal competitive ratio of $2m - 1$ achieved by the work function algorithm).*

- *Given any online algorithm that schedules tasks fractionally on multiple states with competitive ratio c , there exists another online algorithm that executes the task in a single state with competitive ratio $\leq c$.*

The following variables are used in the traversal algorithm (Algorithm 1):

- The current position in the traversal sequence τ is indicated by the variable j .
- Task w^i is (fractionally) executed in states

$$\tau_{i-1}, \tau_{i-1+1}, \dots, \tau_i.$$

I.e., during the processing of task w^i the index j takes values t_{i-1}, \dots, t_i .

- The fraction of task w^i that is executed in state τ_j is denoted by λ_j^i . Note that $\lambda_\ell^i = 0$ for all $\ell \notin t_{i-1}, \dots, t_i$. Note too that there may be $\ell, \ell' \in t_{i-1}, \dots, t_i$, $|\ell - \ell'| \geq 2$, such that $\tau_\ell = \tau_{\ell'}$.
- The variable ρ_j when dealing with task i represents the work expended thus far in the j th position of τ . I.e.,

$$\rho_j = \sum_{i': t_{i'} \leq j \leq t_{i'+1}, i' \leq i} \lambda_j^{i'} w_{\tau_j}^{i'}.$$

An illustration of the execution of Algorithm 1 is given in Appendix A.1.

Given a traversal sequence $\tau = \tau_1, \tau_2, \dots$, define the traversal distance between indices ℓ and ℓ' to be

$$\delta_{\ell, \ell'} = \delta_{\ell', \ell} = \sum_{j=\min(\ell, \ell')}^{\max(\ell, \ell')-1} d_{\tau_j, \tau_{j+1}}.$$

Note that

- The total work done by the fractional traversal algorithm after completing n tasks is exactly $\delta_{0, t_n} + \rho_{t_n}$. This follows as the algorithm advances to the next state in the traversal sequence when the amount of work done in the current state is equal to the distance to the next state.

Algorithm 1 The Fractional Traversal Algorithm.

Input: A traversal sequence τ , a set of states S , the function d , and an online sequence of tasks w_1, w_2, \dots . For each task i :

$j \leftarrow t_{i-1}$
 while $(\sum_{k=t_{i-1}}^j \lambda_k^i < 1)$

- $\lambda_j^i \leftarrow \min \left\{ \frac{d_{j,j+1} - \rho_j}{w_j^i}, 1 - \sum_{k=t_{i-1}}^{j-1} \lambda_k^i \right\}$
- $\rho_j \leftarrow \rho_j + \lambda_j^i \cdot w_{\tau_j}^i$
- if $\rho_j = d_{\tau_j, \tau_{j+1}}$ then $j \leftarrow j + 1, \rho_j \leftarrow 0$

$t_i \leftarrow j$

- The total traversal cost to the fractional traversal algorithm for n tasks is exactly δ_{t_0, t_n} , every transition from τ_j to τ_{j+1} has transition cost $d_{j, j+1}$.

We now show that a Traversal Algorithm's performance is monotone with respects to the tasks it receives.

LEMMA 3.1. *Let τ be some traversal sequence, and let $\mathbf{w} = \langle w^1, \dots, w^n \rangle$ and $\tilde{\mathbf{w}} = \langle \tilde{w}^1, \dots, \tilde{w}^n \rangle$ be two sequences of tasks such that for every $i \in \{1, \dots, n\}$ and every $s \in S$, we have that $\tilde{w}_s^i \leq w_s^i$. The total work done by the traversal algorithm with respect to τ and \mathbf{w} is at least the amount of work done with respect to τ and $\tilde{\mathbf{w}}$.*

Proof. Let t_i be the position of the traversal algorithm after processing task i in the task sequence \mathbf{w} , and \tilde{t}_i the position of the traversal algorithm after processing task i in the task sequence $\tilde{\mathbf{w}}$. In addition, let ρ_{t_i} be the amount of work done by the algorithm since index t_i given sequence \mathbf{w} , and $\tilde{\rho}_{\tilde{t}_i}$ the corresponding amount of work since reaching \tilde{t}_i and given task sequence $\tilde{\mathbf{w}}$.

We show by an inductive argument that for every $i \in \{0, \dots, n\}$, $\delta_{t_0, t_i} + \rho_{t_i} \geq \delta_{\tilde{t}_0, \tilde{t}_i} + \tilde{\rho}_{\tilde{t}_i}$. This is clearly true for $i = 0$. Let the claim hold before the arrival of the i -th task. If the traversal algorithm that processes sequence $\tilde{\mathbf{w}}$ doesn't reach t_i when processing task \tilde{w}^i , ($\tilde{t}_i < t_i$) then

$$\begin{aligned} \delta_{t_0, \tilde{t}_i} + \rho_{\tilde{t}_i} &\leq \delta_{t_0, \tilde{t}_i+1} \\ &\leq \delta_{t_0, t_i} \\ &\leq \delta_{t_0, t_i} + \rho_{t_i}. \end{aligned}$$

where the 1st derivation above is by the definition of the algorithm, the 2nd is by assumption that $\tilde{t}_i < t_i$ and the 3rd is since $\rho_j \geq 0$ for all j .

Lets assume that $\tilde{t}_i \geq t_i$. We'll show that it cannot be that $\tilde{t}_i > t_i$, and when $\tilde{t}_i = t_i$ we show that $\tilde{\rho}_{\tilde{t}_i} \leq \rho_{t_i}$, thus concluding the inductive proof.

Because $w_s^i \geq \tilde{w}_s^i$ for every state s , we have that for every $j \in \{t_{i-1}, \dots, t_i - 1\}$,

$$(3.3) \quad \tilde{\lambda}_j^i \geq \lambda_j^i,$$

where $\tilde{\lambda}_j^i$ and λ_j^i are the fractions of tasks \tilde{w}_i and w^i done in τ_j . This follows from the definition of λ in the fractional traversal algorithm and the inductive hypothesis.

Therefore,

$$(3.4) \quad \tilde{\lambda}_{t_i}^i \leq 1 - \sum_{j=t_{i-1}}^{t_i-1} \tilde{\lambda}_j^i$$

$$(3.5) \quad \begin{aligned} &\leq 1 - \sum_{j=t_{i-1}}^{t_i-1} \lambda_j^i \\ &= \lambda_{t_i}^i. \end{aligned}$$

Inequality (3.4) follows since the fraction left over cannot be more that one minus the fractions already done. Inequality (3.5) follows from Equation (3.3), and the final equality is by definition.

Therefore,

$$\tilde{\rho}_{t_i} = \tilde{\lambda}_{t_i}^i \cdot \tilde{w}_i(\tau_{t_i}) \leq \lambda_{t_i}^i \cdot w_{\tau_{t_i}}^i = \rho_{t_i} < d_{t_i, t_i+1},$$

so $\tilde{t}_i = t_i$, and we can conclude the induction. □

Lemma 3.1 immediately implies that the total cost of the fractional traversal algorithm given task sequence $\tilde{\mathbf{w}}$ is smaller than the cost of the fractional traversal algorithm given task sequence \mathbf{w} . This is since $\delta_{\tilde{t}_0, \tilde{t}_i} \leq \delta_{t_0, t_i}$ for every $i \in \{0, \dots, n\}$.

3.3 "Follow the Traversal" Algorithm We would like to come up with a dynamic pricing scheme for metrical task systems for which agents driven by self interest follow the traversal sequence.

We now describe a new algorithm which can be simulated by a dynamic pricing scheme that encourages greedy agents to "follow" the fractional traversal algorithm. The algorithm, for a given metrical task system (a set of states S and a distance metric $d : S \times S \mapsto \mathbb{R}^+$) and traversal sequence τ , receives and processes a sequence of tasks w^1, w^2, \dots online.

Roughly speaking, after processing each task, the algorithm simulates the traversal algorithm presented in Section 3.2. When handling task i , the algorithm views all states between its current position in the sequence, and the fractional traversal algorithm's current position in the sequence, t_{i-1} , as a superstate where the cost of a task in a superstate is the min cost over its constituent states. The details of the algorithm are given in Algorithm 2.

Even though there may be several indices that minimize \tilde{c} , tie breaking can be done arbitrarily, since all the claims hold for every index which minimizes \tilde{c} .

Algorithm 2 “Follow the traversal” algorithm that allows going to all the states between its current state and the current state of the traversal algorithm “for free”.

Input: A traversal sequence τ , a set of states S , a distance metric d and an online sequence of tasks w^1, w^2, \dots

Let ℓ_{i-1} denote the position of the algorithm in the traversal sequence after completing tasks w^1, \dots, w^{i-1} and t_{i-1} be the position of the traversal algorithm after completing these tasks ($\ell_0 = t_0 = 1$).

Given task w^i :

1. Let $j_{\min} \leftarrow \min(\ell_{i-1}, t_{i-1})$
and $j_{\max} \leftarrow \max(\ell_{i-1}, t_{i-1})$.
2. Let $D_0 \leftarrow \{j_{\min}, \dots, j_{\max}\}$.
3. Define $\tilde{c}(j) = \begin{cases} w_{\tau_j}^i & j \in D_0 \\ w_{\tau_j}^i + \delta_{j_{\max}, j} & j > j_{\max} \\ \infty & \text{otherwise} \end{cases}$
4. Set $\ell_i \leftarrow \arg \min_{j \geq j_{\min}} \tilde{c}(j)$.

3.4 Proof of Approximation

THEOREM 3.2. For any traversal sequence τ , distance metric d , and sequence of tasks w^1, \dots, w^n , let $C^{A2} = w^{A2} + d^{A2}$ be the total cost of Algorithm 2, and $C^T = w^T + d^T$ the total cost of the traversal algorithm for that sequence of tasks (the processing cost plus the transition cost). Then $C^{A2} \leq 2 \cdot C^T$.

Proof. We use a potential argument. The potential function we use is $\Phi = \delta_{\ell_i, t_i}$, which is the traversal distance the current position of Algorithm 2 and the current position of the traversal algorithm. We show that for every incoming task i ,

$$C_i^{A2} + \Delta\Phi_i \leq 2 \cdot C_i^T,$$

where C_i^{A2} and C_i^T are the costs of Algorithm 2 and Algorithm 1 respectively when processing the i th task, and $\Delta\Phi_i$ is the change in potential after processing it.

Let $T_i = \{j : t_{i-1} \leq j \leq t_i\}$ be the set of positions that the traversal algorithm visits while processing the i -th task, and let λ_j^i be the amount of work done in τ_j in the traversal while processing the task. Since the traversal algorithm processes the whole task, $\sum_{j \in T_i} \lambda_j^i = 1$.

Note that

$$C_i^T = \delta_{t_{i-1}, t_i} + \sum_{j \in T_i} \lambda_j^i \cdot w_{\tau_j}^i,$$

and also that

$$(3.6) \quad C_i^{A2} \leq \delta_{\ell_{i-1}, \ell_i} + w_{\tau_{\ell_i}}^i,$$

since Algorithm 2 may use a direct edge from $\tau_{\ell_{i-1}}$ to τ_{ℓ_i} . By definition, the difference in the potential after processing the i th task is:

$$(3.7) \quad \Delta\Phi_i = -\delta_{\ell_{i-1}, t_{i-1}} + \delta_{\ell_i, t_i}$$

Let I_j be an indicator as to whether $\tau_j \notin D_0$. Algorithm 2 chooses an index that minimizes \tilde{c} , therefore $w_{\tau_{\ell_i}}^i + \delta_{j_{\max}, \ell_i} \cdot I_{\ell_i} \leq w_{\tau_j}^i + \delta_{j_{\max}, j} \cdot I_j$. Hence, for any $j \in \{t_{i-1}, \dots, t_i\}$ we have $w_{\tau_{\ell_i}}^i + \delta_{j_{\max}, \ell_i} \cdot I_{\ell+1} \leq w_{\tau_j}^i + \delta_{j_{\max}, t_i} \cdot I_{t_i}$. Since λ_j^i is a convex combination we have:

$$(3.8) \quad w_{\tau_{\ell_i}}^i + \delta_{j_{\max}, \ell_i} \cdot I_{\ell+1} \leq \sum_{j \in \{t_{i-1}, \dots, t_i\}} \lambda_j^i \cdot (w_{\tau_j}^i + \delta_{j_{\max}, t_i} \cdot I_{t_i}) = \sum_{j \in \{t_{i-1}, \dots, t_i\}} \lambda_j^i \cdot w_{\tau_j}^i + \delta_{j_{\max}, t_i} \cdot I_{t_i} \leq C_i^T.$$

Case 1: $\ell_{i-1} \leq t_{i-1}$ and $\ell_i \leq t_{i-1}$



In this case, $\Delta\Phi = -\delta_{\ell_{i-1}, \ell_i} + \delta_{t_{i-1}, t_i}$. Therefore,

$$\begin{aligned} C_i^{A2} + \Delta\Phi &\stackrel{(3.6)}{\leq} w_{\tau_{\ell_i}}^i + \delta_{\ell_{i-1}, \ell_i} + \Delta\Phi \stackrel{(3.7)}{=} w_{\tau_{\ell_i}}^i + \delta_{t_{i-1}, t_i} \\ &\stackrel{(3.8)}{\leq} C_i^T + \delta_{t_{i-1}, t_i} \leq 2C_i^T. \end{aligned}$$

Case 2: $\ell_{i-1} \leq t_{i-1}$ and $t_{i-1} < \ell_i \leq t_i$



We have,


$$\begin{aligned} C_i^{A2} + \Delta\Phi &\stackrel{(3.6)}{\leq} w_{\tau_{\ell_i}}^i + \delta_{\ell_{i-1}, t_{i-1}} + \delta_{t_{i-1}, \ell_i} + \Delta\Phi \\ &\stackrel{(3.7)}{=} w_{\tau_{\ell_i}}^i + \delta_{t_{i-1}, \ell_i} + \delta_{\ell_i, t_i} \\ &\stackrel{(3.8)}{\leq} C_i^T + \delta_{\ell_i, t_i} \\ &\leq 2C_i^T. \end{aligned}$$

Case 3: $\ell_{i-1} \leq t_{i-1}$ and $t_i < \ell_i$



$$\begin{aligned}
 C_i^{A2} + \Delta\Phi &\stackrel{(3.6)}{\leq} w_{\tau_{\ell_i}}^i + \delta_{\ell_{i-1}, t_{i-1}} + \delta_{\ell_{i-1}, \ell_i} + \Delta\Phi \\
 &\stackrel{(3.7)}{=} w_{\tau_{\ell_i}}^i + \delta_{\ell_{i-1}, \ell_i} + \delta_{\ell_i, \ell_i} \\
 &\stackrel{(3.8)}{\leq} C_i^T + \delta_{\ell_i, \ell_i} \\
 &\leq 2C_i^T.
 \end{aligned}$$


Case 4: $t_{i-1} < \ell_{i-1}$ and $t_i \leq \ell_i \leq \ell_{i-1}$



In this case, $\Delta\Phi = -\delta_{\ell_{i-1}, t_i} - \delta_{\ell_i, \ell_{i-1}}$. Hence,

$$C_i^{A2} + \Delta\Phi \stackrel{(3.6)}{\leq} w_{\tau_{\ell_i}}^i + \delta_{\ell_i, \ell_{i-1}} + \Delta\Phi \stackrel{(3.7)}{=} w_{\tau_{\ell_i}}^i - \delta_{\ell_{i-1}, t_i} \stackrel{(3.8)}{\leq} C_i^T.$$


Case 5: $t_{i-1} < \ell_{i-1}$ and $t_i \leq \ell_{i-1} < \ell_i$



In this case, $\Delta\Phi = -\delta_{\ell_{i-1}, t_i} + \delta_{\ell_{i-1}, \ell_i}$. Therefore,

$$C_i^{A2} + \Delta\Phi \stackrel{(3.6)}{\leq} w_{\tau_{\ell_i}}^i + \delta_{\ell_{i-1}, \ell_i} + \Delta\Phi \stackrel{(3.7)}{\leq} w_{\tau_{\ell_i}}^i + 2\delta_{\ell_{i-1}, \ell_i} \stackrel{(3.8)}{\leq} 2C_i^T.$$

Case 6: $t_{i-1} < \ell_{i-1}$ and $\ell_i < t_i \leq \ell_{i-1}$



In this case, $\Delta\Phi = -\delta_{\ell_{i-1}, \ell_i} - \delta_{\ell_i, \ell_{i-1}}$. Therefore,

$$\begin{aligned}
 C_i^{A2} + \Delta\Phi &\stackrel{(3.6)}{\leq} w_{\tau_{\ell_i}}^i + \delta_{\ell_i, \ell_{i-1}} + \Delta\Phi \\
 &= w_{\tau_{\ell_i}}^i + \delta_{\ell_i, t_i} + \delta_{t_i, \ell_{i-1}} + \Delta\Phi \\
 &\stackrel{(3.7)}{\leq} w_{\tau_{\ell_i}}^i + \delta_{\ell_i, t_i} \\
 &\stackrel{(3.8)}{\leq} 2C_i^T.
 \end{aligned}$$

Case 7: $t_{i-1} < \ell_{i-1}$ and $\ell_i < \ell_{i-1} < t_i$



In this case, $\Delta\Phi = -\delta_{\ell_i, \ell_{i-1}} - \delta_{\ell_{i-1}, \ell_i} + \delta_{\ell_i, t_i} \leq -\delta_{\ell_i, \ell_{i-1}} + \delta_{\ell_{i-1}, t_i}$. Therefore,

$$\begin{aligned}
 C_i^{A2} + \Delta\Phi &\stackrel{(3.6)}{\leq} w_{\tau_{\ell_i}}^i + \delta_{\ell_i, \ell_{i-1}} + \Delta\Phi \stackrel{(3.7)}{\leq} w_{\tau_{\ell_i}}^i + \delta_{\ell_{i-1}, t_i} \\
 &\stackrel{(3.8)}{\leq} 2C_i^T.
 \end{aligned}$$

Case 8: $t_{i-1} < \ell_{i-1}$ and $\ell_{i-1} \leq \ell_i < t_i$



In this case, $\Delta\Phi = -\delta_{\ell_{i-1}, \ell_{i-1}} + \delta_{\ell_i, t_i} \leq \delta_{\ell_{i-1}, t_i}$. Therefore,

$$\begin{aligned}
 C_i^{A2} + \Delta\Phi &\stackrel{(3.6)}{\leq} w_{\tau_{\ell_i}}^i + \delta_{\ell_{i-1}, \ell_i} + \Delta\Phi \\
 &\stackrel{(3.7)}{\leq} w_{\tau_{\ell_i}}^i + \delta_{\ell_{i-1}, \ell_i} + \delta_{\ell_{i-1}, t_i} \\
 &\stackrel{(3.8)}{\leq} 2C_i^T.
 \end{aligned}$$

Case 9: $t_{i-1} < \ell_{i-1}$ and $\ell_{i-1} \leq t_i \leq \ell_i$



In this case, $\Delta\Phi = -\delta_{\ell_{i-1}, \ell_{i-1}} + \delta_{\ell_i, t_i} \leq \delta_{\ell_{i-1}, \ell_i}$. Therefore,

$$\begin{aligned}
 C_i^{A2} + \Delta\Phi &\stackrel{(3.6)}{\leq} w_{\tau_{\ell_i}}^i + \delta_{\ell_{i-1}, \ell_i} + \Delta\Phi \\
 &\stackrel{(3.7)}{\leq} w_{\tau_{\ell_i}}^i + 2\delta_{\ell_{i-1}, \ell_i} \\
 &\stackrel{(3.8)}{\leq} 2C_i^T.
 \end{aligned}$$

□

3.5 Pricing Scheme In this section, we define a dynamic pricing scheme that observes decisions made by greedy agents and sets a surcharge for the next agent to arrive. We want to set prices so that states chosen by (greedy) agents are the same as the ones chosen by Algorithm 2 (up to tie breaking) given the sequence of tasks \mathbf{w} .

The problem is that dynamic pricing schemes, as we've defined them, can only observe the greedy behavior of agents. the state chosen by the agents, and the amount of work done in that state while processing the task; Namely, after agent i has selfishly used the system, the mechanism observes s_i , the state chosen by agent i , and $w_{s_i}^i$, the amount of work done in s_i while processing the task, but the other coordinates of w^i ($w_1^i, \dots, w_{s_i-1}^i, w_{s_i+1}^i, \dots, w_m^i$) are not known to it.

We resolve this issue by giving the pricing scheme a different (imaginary) sequence of tasks $\tilde{\mathbf{w}}$ with the following *consistency* properties:

1. For every i , $w_{s_i}^i$ is equal to $\tilde{w}_{s_i}^i$.
2. For every i and for every $j \in S \setminus \{s_i\}$, $w_j^i \geq \tilde{w}_j^i$.
3. For every i , given the pricing function P_i over the state of the system after observing events $1, \dots, i-1$, the computed vector \tilde{w}^i is consistent with the agent selfishly choosing to process her task in state s_i ; I.e., for all j

$$(3.9) \quad \tilde{w}_j^i + d_{s_{i-1},j} + P_i(j) \geq w_{s_i}^i + d_{s_{i-1},s_i} + P_i(s_i) .$$

Note that running Algorithm 2 on the sequence of tasks $\tilde{\mathbf{w}}$ yields a 2-approximation to the cost of the Fractional Traversal Algorithm on the sequence of tasks \mathbf{w} : by Lemma 3.1, the cost of the Fractional Traversal Algorithm on $\tilde{\mathbf{w}}$ is no larger than its cost on \mathbf{w} , and by Theorem 3.2, Algorithm 2 gives a 2-approximation to the cost of the Fractional Traversal Algorithm (on any sequence, in particular $\tilde{\mathbf{w}}$). Therefore, it suffices to find a pricing scheme that makes the agents follow Algorithm 2 on sequence $\tilde{\mathbf{w}}$.

We show how to find $\tilde{\mathbf{w}}$ that satisfies the consistency properties in Section 3.6. From now on, we assume such \tilde{w}^{i-1} can be computed after observing agent i 's greedy behavior. We show how to devise a pricing scheme that given $\tilde{\mathbf{w}}$, makes incoming agents act *as though* they are being guided by Algorithm 2 given input $\tilde{\mathbf{w}}$. We describe how to do this in Pricing Scheme 1. We note that although Pricing Scheme 1 may price states using negative number, we can increase all prices by a large enough constant without affecting the decisions of the agents.

THEOREM 3.3. *Pricing Scheme 1 makes greedy agents simulate Algorithm 2 with task sequence $\tilde{\mathbf{w}}$ given as input.*

Proof. For every state $s' \in S$, let $m(s')$ be as defined in Pricing Scheme 1 and let \tilde{c} be as defined in Algorithm 2. Since $m(s') \in \arg \min_{j: \tau_j = s'} \tilde{c}(j)$, we may assume that Algorithm 2 always prefers $m(s')$ amongst the indices for which $\tau_j = s'$.

Let s be the state chosen by the agent. If $m(s) \in D_0$, then the disutility of the agent is $w_s^i + d_{s_{i-1},s} + P_i(s) = w_s^i + d_{s_{i-1},s} - d_{s_{i-1},s} + \tilde{c}(m(s))$. If $m(s) \notin D_0$, then now the disutility is $w_s^i + d_{s_{i-1},s} + P_i(s) = w_s^i + d_{s_{i-1},s} - d_{s_{i-1},s} + \delta_{j_{\max},m(s)} = \tilde{c}(m(s))$.

By the way \tilde{w}^i is computed, we get that if the algorithm would have received \tilde{w}^i as input, then given $\tilde{s} \in S$ that was not chosen by agent i , we have that,

Pricing Scheme 1 A dynamic pricing scheme that makes incoming agents “follow” Algorithm 2.

Input: A traversal sequence τ , a set of states S , a distance metric d and an observed sequence of decisions $(s_1, w_{s_1}^1), (s_2, w_{s_2}^2), \dots, (s_{i-1}, w_{s_{i-1}}^i)$, made by agents $1, 2, \dots, i-1$.

The variables ℓ_{i-1} and t_{i-1} represent two indices in the traversal sequence, set after the $i-1$ th agent's decision. For $i-1=0$ take $\ell_0 \leftarrow 1$ and $t_0 \leftarrow 1$.

Setting prices for agent i :

1. Let $j_{\min} \leftarrow \min(\ell_{i-1}, t_{i-1})$
and $j_{\max} \leftarrow \max(\ell_{i-1}, t_{i-1})$.
2. Let $D_0 \leftarrow \{j_{\min}, \dots, j_{\max}\}$.
3. For a state s let $m(s) \leftarrow \min\{j \geq j_{\min} \mid \tau_j = s\}$.
4. Define $P_i(s) \equiv \begin{cases} -d_{s_{i-1},s} & m(s) \in D_0, \\ -d_{s_{i-1},s} + \delta_{j_{\max},m(s)} & \text{otherwise.} \end{cases}$

Determining ℓ_i and t_i after observing $(s_i, w_{s_i}^i)$:

1. Set $\ell_i \leftarrow m(s_i)$.
 2. Compute \tilde{w}^i using the set of equations (3.10) as described in Section 3.6.
 3. Update t_i as done by the Fractional Traversal Algorithm (Algorithm 1) using \tilde{w}^i as input.
-

- if $m(\tilde{s}) \in D_0$, then

$$\begin{aligned} \tilde{c}(m(\tilde{s})) &= \tilde{w}_{\tilde{s}}^i + d_{s_{i-1},\tilde{s}} - d_{s_{i-1},\tilde{s}} \\ &= \tilde{w}_{\tilde{s}}^i + d_{s_{i-1},\tilde{s}} + P_i(\tilde{s}) \\ &\geq w_s^i + d_{s_{i-1},s} + P_i(s) \\ &= \tilde{c}(m(s)), \end{aligned}$$

where the third inequality is derived from (3.9), and

- if $m(\tilde{s}) \notin D_0$, then

$$\begin{aligned} \tilde{c}(m(\tilde{s})) &= \tilde{w}_{\tilde{s}}^i + d_{s_{i-1},\tilde{s}} - d_{s_{i-1},\tilde{s}} + \delta_{j_{\max},m(\tilde{s})} \\ &= \tilde{w}_{\tilde{s}}^i + d_{s_{i-1},\tilde{s}} + P_i(\tilde{s}) \\ &\geq w_s^i + d_{s_{i-1},s} + P_i(s) \\ &= \tilde{c}(m(s)), \end{aligned}$$

where the third inequality is derived from (3.9).

Therefore, the agent picks a state s such that $m(s)$ minimizes \tilde{c} given task \tilde{w}^i . Algorithm 2 also chooses an index that minimizes \tilde{c} . Since the tie breaking of Algorithm 2 can be done arbitrarily, we can assume that

the algorithm chooses index $m(s)$ as well when given task \bar{w}^i . \square

COROLLARY 3.1. *Pricing Scheme 1 yields a price of anarchy of $16(m-1)$.*

3.6 Computing “Imaginary” Tasks As mentioned in Section 3.5, Pricing Scheme 1 cannot observe the real sequence of task vectors \mathbf{w} . This is problematic, since the pricing scheme runs a simulation of the fractional traversal algorithm, which takes the task vectors as input.

To resolve this issue, we need to compute imaginary task vectors with the three *consistency* properties, as described in Section 3.5. We are given a pricing function $P_i : S \mapsto \mathbb{R}^+$ that is set on states of the task system after the arrival of the $i-1$ -th agent, a state s_i chosen by the i -th agent, and the amount of work done in that state while processing the tasks (which is exactly $w_{s_i}^i$). Let $w^i = (w_1^i, w_2^i, \dots, w_m^i)$ be the current incoming task (which we cannot observe). We claim that the vector \bar{w}^i defined below is coordinate-wise dominated by w^i and consistent with the agent greedily choosing to process the task in s_i while the states are priced using P_i :

$$(3.10) \quad \bar{w}_j^i := \max\{0, w_{s_i}^i + d_{s_{i-1}, s_i} + P_i(s_i) - d_{s_{i-1}, j} - P_i(j)\}$$

for all $j \in S$. Note that $\bar{w}_{s_i}^i = w_{s_i}^i$.

By definition, $\bar{w}_j^i \geq 0$, and since the agent chooses s_i , we have that $w_j^i + d_{s_{i-1}, j} + P_i(j) \geq w_{s_i}^i + d_{s_{i-1}, s_i} + P_i(s_i)$, i.e., $w_j^i \geq w_{s_i}^i + d_{s_{i-1}, s_i} + P_i(s_i) - d_{s_{i-1}, j} - P_i(j)$. Therefore, it holds that $w_j^i \geq \bar{w}_j^i$. In addition, since $\bar{w}_{s_i}^i \geq w_{s_i}^i + d_{s_{i-1}, s_i} + P_i(s_i) - d_{s_{i-1}, s_i} - P_i(s_i)$, we have that $\bar{w}_j^i + d_{s_{i-1}, j} + P_i(j) \geq w_{s_i}^i + d_{s_{i-1}, s_i} + P_i(s_i)$ as required in (3.9). We conclude that we are able to compute a task vector which satisfies the desired *consistency* properties, thus enabling us to devise pricing over states that insures a nearly optimal price of anarchy.

4 Pricing Servers on a Line

The k -server problem, introduced by Manasse et al. [19], is defined as follows. Let (V, d) be a fixed metric space. There are k servers, initially located in specified points $S^0 = \{s_1^0, s_2^0, \dots, s_k^0\}$ s.t. for $1 \leq j \leq k$, $s_j^0 \in V$, that can be moved by the algorithm. Moving a server from x to y costs $d(x, y)$. The goal is to serve a sequence of arriving requests with minimum possible cost. Each request r_i is a point in V , and to serve it, the algorithm has to move (at least) one server to that point.

In our setting, we assume that the requests are made by successive agents. Moreover, we assume that each agent is greedy, and will therefore ask that their request is served with the server that minimizes the sum of its

distance to the request and the price we set for the server. Therefore, before the arrival of the i th agent, the algorithm sets a payment function for the servers: $P_i : S^{i-1} \mapsto \mathbb{R}$, and agent i would choose to move server:

$$\ell_i \in \arg \min_j \{d(r_i, s_j^{i-1}) + P_i(s_j^{i-1})\},$$

and the next state would be

$$S^i = S^{i-1} \setminus \{s_{\ell_i}^{i-1}\} \cup \{r_i\}.$$

The social cost is defined as:

$$f(\bar{\sigma}, \bar{s}) = \sum_{i=1}^{|\bar{\sigma}|} d(s_{\ell_i}^{i-1}, r_i).$$

On a high level, we want to find out if any algorithm can be simulated by appropriate pricing, i.e., whether given an algorithm A and its current state (locations of servers), we can give a pricing of servers that will make an arriving greedy agent with a request r , serve r in the same way that A would have, wherever the request occurs.

For some algorithms, designing a payment scheme is straightforward. For instance, Irani and Rubinfeld proved that the following 2-server algorithm *Balance2* is 10-competitive on any metric space [10]. The algorithm keeps track of the total distance that each of the two servers has travelled, and serves each request with the server that minimizes the sum of its distance to the request and *half* the total distance it has traveled before. Clearly, setting the price of each server to the second summand, i.e., half the total distance the server has moved already, will make the agents follow the algorithm. The only catch is that we have no control over tie-breaking since it is done by the agents. However, this can only affect the total cost by an additive constant.

In general, however, simulating any algorithm A is not as easy. For instance, A could move many servers (in a non-lazy fashion) upon a request. This is the case with some known algorithms, for example with the *Double Coverage* (DC) algorithm [4] that is k -competitive for the line metric. (This algorithm has been generalized to the case of all tree metrics, retaining the optimal ratio [5].) As an agent is only allowed to pick a server that will serve the request, such complex algorithms cannot be simulated directly.

In this work, we restrict our attention to k -server on the line metric, and the *Double Coverage* algorithm in particular. We observe that the standard transformation that makes it a *lazy* algorithm, if done carefully, preserves its other properties, which we call *locality* and *monotonicity* (all defined later), and that combined these enable simulation by a pricing scheme. We remark that, in general, such transformation employs keeping track of “virtual positions” of servers and basing decisions on those. This

makes designing a pricing scheme to simulate the algorithm challenging, since an agent's decision is based on the "real positions" of the servers (and their prices).

4.1 Notation Given an algorithm A , let $S_A = \{s_1, \dots, s_k\}$ denote its *state*, i.e., the set of points where A 's servers are currently located. We will associate the servers with the points they are in. Similarly, we will associate a request r with the point in which it occurs. Moreover, we associate every point on the real line with a real number. Consequently, we will write $v \leq w$ to denote that v lies to the left of w (or coincides with it) and denote the distance between v and w by $|v - w|$. We use the notation $[v, w]$ to denote the interval spanned by the points v and w , even if $v > w$. Given a point r and a finite set of points X , we call the (at most) two points of X , $\max\{x \in X \mid x \leq r\}$ and $\min\{x \in X \mid x \geq r\}$ *adjacent* to r ; note that the first one does not exist if $r < \min X$, and the second does not exist if $r > \max X$, and they coincide if $r \in X$.

A deterministic online algorithm is *lazy* if the only moves it makes consist in moving a single server to the current request. An algorithm is *local* if it always serves a request with a server that is adjacent to it. Furthermore, an algorithm A is *monotone* if it satisfies the following property: given S_A and a request (at) r , if A would serve a request at r with $s \in S_A$, then it would also serve every request in $[r, s]$ with s .

4.2 Matchings Later, we give a transformation that takes an algorithm A and returns an "equivalent" algorithm A' that is lazy. This transformation relies on minimum cost (perfect, bipartite) matching of S_A to $S_{A'}$, where the costs are specified by the metric distance d ; in our case, the cost of matching $s \in S_A$ to $s' \in S_{A'}$ is $d(s, s') = |s - s'|$. Given two sets of points X and Y of equal cardinality, we denote any such matching by $M(X, Y)$ and its cost by $d(X, Y)$. In our case, when the metric space is a line, we focus on a particular minimum cost perfect matching. Namely, given $X = \{x_1, x_2, \dots, x_k\}$ and $Y = \{y_1, y_2, \dots, y_k\}$ that are sorted from left to right, i.e., $x_1 \leq x_2 \leq \dots \leq x_k$ and $y_1 \leq y_2 \leq \dots \leq y_k$, we call a matching in which x_i is matched to y_i for all i a *canonical matching*, and denote it $\mathcal{M}(X, Y)$.

LEMMA 4.1. *The canonical matching is a minimum cost perfect matching.*

Proof. Let $M = M(X, Y)$ be any minimum cost perfect matching of X to Y that minimizes the total distance. We show how M can be transformed to the canonical matching without increasing the cost. As long as M is not canonical, find the minimum i such that x_i is not matched to y_i . Then x_i is matched to some y_j and y_i is matched to some x_h . Moreover, $j > i$ and $h > i$ since i is minimal.

We assume wlog that $x_i \leq y_i$ (if not, swap the sets). This implies that $y_j \geq y_i \geq x_i$. Modify M as follows: match x_i with y_i and x_h with y_j . Matching x_i to y_i decreases the total distance by $y_j - y_i$. By the triangle inequality, $|x_h - y_j| \leq |x_h - y_i| + |y_i - y_j| = |x_h - y_i| + y_j - y_i$, i.e., matching x_h to y_j increases the total distance by at most $y_j - y_i$. Hence, the cost does not increase throughout the transformation, and is thus minimized by the canonical matching $\mathcal{M}(X, Y)$ we obtain at the end. \square

Given X and Y and a point $r \in X$, we call a matching of X to Y *r-local* if r is matched to a point in Y that is adjacent to it.

LEMMA 4.2. *For any two sets of points X and Y on a line such that $|X| = |Y|$ and a given point $r \in X$, there exists a minimum cost perfect matching of X to Y that is r -local.*

Proof. Let $M = M(X, Y)$ be any minimum cost perfect matching of X to Y . If M is not r -local, do the following. Let y be the point in Y that r is matched to, and y' the point in $[r, y] \cap Y$ that is closest to r , i.e., minimizes $|r - y'|$. Let $x \in X$ be the point that y' is matched to in $M(X, Y)$. Then an r -local matching M' is obtained from M by matching r to y' and x to y . To see that the cost of M' is no larger than the cost of M , consider the change in distances of r and x to the points they are matched to. Since $y' \in [y, r]$, we have $|y' - r| = |y - r| - |y' - y|$, i.e., the distance for r decreases by $|y' - y|$. As $|x - y| \leq |x - y'| + |y' - y|$, the distance for x increases by at most the same amount. \square

We denote the r -local minimum cost matching of X to Y by $M_r(X, Y)$. In particular, we call the r -local matching that is the result of the procedure described in the proof of Lemma 4.2 applied to $\mathcal{M}(X, Y)$ an *r-canonical matching*, and denote it by $\mathcal{M}_r(X, Y)$.

4.3 Algorithm transformation Algorithm 3 describes how to transform any algorithm A into an equivalent one that is lazy and never incurs larger cost than A . Later, we observe that (wlog) the output algorithm is also local. Finally, we note that for DC the transformation also preserves monotonicity.

In general, the transformation only guarantees that its output is a lazy algorithm of no larger cost than A . The additional properties rely on particular structure of the input algorithm and the matching that is used upon requests.

LEMMA 4.3. *Given an algorithm A , the transformation produces a lazy algorithm A' such that for every sequence of requests $R = r_1, \dots, r_i$, it holds that $C(A', R) \leq C(A, R)$.*

Proof. It follows directly from the transformation that A' is lazy. In order to bound its cost, denote the costs of

Algorithm 3 The algorithm transformation.

Input: An algorithm A and an online sequence of requests r_1, \dots, r_i .

Output: A lazy algorithm A' .

1. Let S_A be the state of A after serving r_i and $S_{A'}$ the state of A' right before serving r_i ; note that $r_i \in S_A$.
 2. Find $M = M(S_A, S_{A'})$ and serve r_i with the server matched to it in M .
-

A and A' upon serving request r_i by $C_i(A)$ and $C_i(A')$ respectively. We will prove that

$$(4.11) \quad C_i(A') + \Delta\Phi \leq C_i(A),$$

where $\Phi = d(S_A, S_{A'})$ is a non-negative potential function.

To prove (4.11), it suffices to consider the moves of A and A' independently, in this order. Fix $M = M(S_A, S_{A'})$.

We keep M fixed as A moves its servers. Clearly, when A moves a server s by distance d , the cost of M does not increase by more than d . Hence, the same holds for the perfect matching of minimum cost. Thus Φ increases by at most d , and (4.11) holds.

Once A is done with its moves, we analyze the move of A' . Note that at this point $r_i \in S_A$, i.e., A has one of its servers at r_i . Let $M = M(S_A, S_{A'})$ be a minimum cost matching of the algorithms' servers and $s \in S_{A'}$ the server of A' that is matched to r_i . Upon the move of s to r_i , Φ decreases by at least $d(s, r_i)$ since this is the case if we keep M fixed. (It is actually easy to see that there is no cheaper matching). Thus, $C_i(A') + \Delta\Phi \leq 0$ holds. \square

REMARK 4.1. *The transformation produces a lazy local algorithm if, upon each request r_i , an r_i -local matching is used.*

Note that we have proved that the algorithm that our transformation produces is local but not that it is monotone. Nevertheless, for the case of line and the *Double Coverage* (DC) algorithm [4, 5], we can prove monotonicity.

4.4 Transforming DC on a line The *Double Coverage* (DC) algorithm has been originally designed for the line metric [4], and later generalized to all tree metrics [5]. For these metrics, it attains the optimal competitive ratio of k . For completeness, we describe DC for the case of the line metric in Algorithm 4.

To prove that DC on a line can be transformed to a lazy local monotone algorithm, we will rely on the fact that DC itself is local and monotone, as well as on properties of canonical matchings on a line.

Algorithm 4 Double Coverage on a line.

Input: A request r on the line.

1. Let s_L and s_R be the servers of DC immediately to the left of r and right of r respectively. (Note that one of them may not exist; in such case ignore it.)
 2. Move both s_L and s_R towards r by $\min\{d(s_L, r), d(s_R, r)\}$.
 3. Serve r with the server that reached it. (Note that both s_L and s_R could have.)
-

LEMMA 4.4. *Applied to DC on a line, the transformation produces a lazy local monotone algorithm if, upon each request r_i , an r_i -canonical matching is used.*

Proof. We focus on monotonicity since the other properties follows from Lemma 4.3 and Remark 4.1. To this end, let DC' be the output of the transformation and $S_{DC'} = \{s_1, \dots, s_k\}$, where $s_1 \leq s_2 \leq \dots \leq s_k$, its current state. Since for each request r , we use the r -local canonical matching in the transformation, r will be served by a server of DC' that is local for it. In particular, DC' will serve a request $r \leq s_1$ with s_1 and a request $r \geq s_k$ with s_k . Similarly, DC' will serve a request at position of one its servers s with s .

Now consider $s_i < r < s_{i+1}$, i.e., a request between two adjacent servers of DC' . Suppose that DC serves r with its j -th leftmost server. Recall that an r -local canonical matching between the state of DC after serving r to $S_{DC'}$ would be used in the transformation for a request (at) r . In the canonical matching r would be matched to s_j but as this matching is made local, r is matched to the last server of DC' on the path from s_j to r . In other words, r is matched to s_i if $j \leq i$ and to s_{i+1} otherwise, and it is served by the server it is matched to. Since DC itself is local monotone, j is a non-decreasing function of r , which implies that DC' is monotone. \square

4.5 Payment scheme We define a payment scheme as $P: S_A \mapsto \mathbb{R}$. The disutility of an agent while serving a request r using server s is defined as $d(r, s) + P(s)$. We assume that at each iteration the request r is served using the server s that minimizes $d(r, s) + P(s)$, where r is the request the incoming agent wishes to serve.

THEOREM 4.1. *For every lazy local monotone deterministic k -server algorithm A on a line, there exists a payment scheme P that makes incoming agents weakly follow A and incur movement cost no larger than that of A .*

See the remarks after the proof for explanation what “weakly” means.

Proof. We obtain the desired pricing P by fixing the price of an arbitrary server in S_A and setting the remaining prices so that they satisfy a system of linear equations. The equations follow from the following consideration.

Let S_A be the state of the A . Since A is lazy local monotone, for every server $s \in S_A$, the set of all points that would be served by s if a request were made in them is an interval that contains s ; we denote such interval $I(s)$ and note that on both ends it could be either open or closed. However, these intervals form a partition of the line.

Let s and s' be a pair of adjacent servers in S_A . Then the intervals $I(s)$ and $I(s')$ are adjacent. Specifically, there is a *threshold point* $v(s, s')$ (which may be one of s, s') in $[s, s']$ such that A would serve all requests on $[s, v(s, s'))$ with s and all requests on $[s', v(s, s'))$ with s' ; a request at $v(s, s')$ would be served with either of s, s' . Having identified the threshold vertex $v(s, s')$, we introduce the following equation:

$$(4.12) \quad P(s) + |s - v(s, s')| = P(s') + |s' - v(s, s')| .$$

We claim that the system of equation(s) (4.12) taken for all pairs of adjacent servers has a unique solution once the price for an arbitrary server is fixed. This is easiest to see if it is one of the extreme servers, say the leftmost. Then, the pricing that is the unique solution to our system can be found by considering the servers from left to right, using (4.12) at each server past the leftmost one to determine its price.

To see that the pricing P makes the agents follow the algorithm, consider a request at point v that A would serve with s . Then for any $s' \in S_A$, consider the interval $[s', v]$, and let $s' = s_0, s_1, \dots, s_i = s$ be the sequence of servers whose intervals intersect $[s', v]$ as we move from s' to v . By (4.12), for each $j < i$, we have $P(s_j) = P(s_{j+1}) + |s_{j+1} - v(s_j, s_{j+1})| - |s_j - v(s_j, s_{j+1})| \geq P(s_{j+1}) - |s_{j+1}, s_j|$. Summing these inequalities up to $j - 1$, we get that for all $j \leq i$

$$(4.13) \quad P(s') = P(s_0) \geq P(s_j) - |s_0 - s_j| .$$

Therefore, if $s_i = s \in [s', v]$, then (4.13) yields that $P(s') + |s' - v| = P(s') + |s' - s| + |s - v| \geq P(s) + |s - v|$, i.e., that s minimizes the disutility among servers whose intervals intersect $[s', v]$. If s does not lie on this path, i.e., if v lies between s' and s , we note that by (4.13) for $j = i - 1$, $P(s') + |s' - v| = P(s') + |s' - s_{i-1}| + |s_{i-1} - v| \geq P(s_{i-1}) + |s_{i-1} - v|$. But as s_{i-1} and $s_i = s$ are adjacent, (4.12) guarantees that $P(s) + |s - v| \leq P(s) + |s - v(s, s_{i-1})| = P(s_{i-1}) + |s_{i-1} - v(s, s_{i-1})| \leq P(s_{i-1}) + |s_{i-1} - v|$. Thus s minimizes the disutility also in this case, and thus overall. \square

As noted before, we stress that tie-breaking is up to the agents, and therefore we cannot guarantee that the agents will follow an algorithm exactly. This is

what we mean by saying in Theorem 4.1 that the agents “weakly follow” such algorithm. However, tie-breaking affects neither Lemma 4.3 nor Remark 4.1, i.e., even if agents break ties differently than the algorithm, in particular, even if they move non-local servers, we can still price the servers in such a way that the cost of moving servers by the agents is no larger than the cost of moving them by the algorithm. We sketch why this is the case. If for a given request r , both adjacent servers minimize the disutility of serving r , then r is exactly at the threshold vertex by the pricing scheme’s design. Then it is easy to see that either server can be matched to r in a minimum cost matching, i.e., a move of either server could result from the transformation, depending on the matching that is chosen. To deal with non-adjacent servers, it is easiest to assume that for two successive servers s_i and s_{i+1} , the threshold vertex $v(s_i, s_{i+1})$ is always strictly between them. This can be done by moving the vertex by arbitrarily small $\epsilon > 0$ away from the server s it coincides with. Clearly, this may result in violating (4.11) only by 2ϵ , and we can make ϵ decrease exponentially for successive requests, so that the total additional additive cost is at most $c \cdot \epsilon$, where c can also be arbitrarily small.

COROLLARY 4.1. *There exists a payment scheme for the k -server problem on a line which yields a price of anarchy of k .*

5 Pricing Parking Slots on a City Block

5.1 Problem Definition We consider the following strategic parking problem; the underlying problem is known as the (online) metric matching but we use the parking terminology.

The “town” is modeled as an undirected graph $G = (V, E)$, $|V| = m$. Vertices of G represent empty parking slots, which are the possible destinations of drivers arriving in town. Every edge $e = (u, v) \in E$ has an associated weight w_e , corresponding to the distance between vertices v and u .

As for the cars:

- A total of n cars are going to park over time ($n \leq m$). Cars do not depart until every car has parked.
- Cars arrive online, one at a time, in order of their index, $1, 2, \dots, n$.
- Cars can only park at unoccupied parking spots.
- Cars behave selfishly and seek to minimize their disutility (defined below).

We denote the parking spot taken by car i by $x(i)$, and call the function $x : [n] \mapsto V$ that maps cars to parking slots *an assignment*.

A dynamic pricing scheme $P = (P_1, P_2, \dots, P_n)$, where

$$P_i : \{x(1), x(2), \dots, x(i-1)\} \times V \mapsto \mathbb{R},$$

sets a price for an unoccupied parking space $v \in V$ given that there are $i-1$ occupied parking spaces: $x(1), \dots, x(i-1)$. The functions P_i may use random bits. Set $P_0(v) = 0$ for all $v \in V$ (or any other constant). When we refer to $P_i(v)$, we assume that $i-1$ cars have already arrived and are parked at $x(1), \dots, x(i-1)$. Thus,

- Given G and $i \in 1, \dots, n$, and given that cars $j \in 1, \dots, i-1$ are already parked at locations $x(1), \dots, x(i-1)$, $P_i(v) \in \mathbb{R}$ gives a price to park at $v \in V \setminus \{x(1), \dots, x(i-1)\}$. This is a one-time parking fee, not (say) charged by the hour.
- The current parking fees, $P_i(v)$, $v \in V \setminus \{x(1), \dots, x(i-1)\}$ are known to car i before parking.

Incentives for cars are determined as follows:

- Let $g(i)$, $1 \leq i \leq n$ be the destination of car i , g is called the *goal function*.
- Given the graph G , goal function g , and assignment function x , the *non-monetary cost* to agent $i \in [n]$ with destination $g(i)$ and an assigned parking slot $x(i)$ is $d(g(i), x(i))$, the weight of a minimum weight path from $x(i)$ to $g(i)$ in G .
- The disutility of agent i (who chooses to park at unoccupied slot $x(i)$) is $d(g(i), x(i)) + P_i(x(i))$. *I.e.*, the disutility is the sum of the non-monetary cost to agent i plus the parking fee to park at $x(i)$.

Therefore, the agent will choose to park at:

$$s_i = \arg \min_{s \in S_i} \{d(g(i), s) + P_i(s)\},$$

where S_i is the set of unoccupied vertices of G at the time agent i 's arrival.

As the goal is cost minimization, the target function f is the social cost,

$$f(\bar{\sigma}, \bar{s}) = \sum_{i=1}^{|\bar{\sigma}|} d(g(i), s),$$

i.e., the sum of all the agents' distances from their goals.

Since the pricing schemes presented in this section are randomized, the price of anarchy, defined in Equation (1.2) should be taken in expectation, and is defined as

$$(5.14) \quad \max_{\bar{\sigma}, \bar{s} \in \text{eq}(\bar{P}, \bar{\sigma})} \frac{\mathbb{E}(f(\bar{\sigma}, \bar{s}))}{f(\bar{\sigma}, \text{opt}(\bar{\sigma}))},$$

where $\text{eq}(\bar{P}, \bar{\sigma})$ is the set of equilibria assignments of \bar{P} , which depends on the random bits used in \bar{P} .

We refer to the pricing scheme P that satisfies $P_i(v) = 0$ for all $i \in 1, \dots, n$, all $x(1), \dots, x(i-1)$ and all $v \in V \setminus \{x(1), \dots, x(i-1)\}$ as the *free parking*, and its price of anarchy as the *cost of free parking*.

5.2 Dynamic Harmonic Pricing for the Weighted Line Graph

In this section we describe the *harmonic pricing scheme* that guaranties an $O(\min\{\log R, n^2\})$ approximation to the minimal cost matching *in expectation*, where R is the aspect ratio of the underlying metric. The expectation is over the random bits of the mechanism. The coordination mechanisms sets fees to the vacant parking slots. The payments are devised without knowing the goals of the drivers. Once set, a car parks so as to minimize its disutility.

The improved approximation guaranties of this mechanism should be compared with the $\Omega(2^n)$ approximation lower bound for the greedy matching on the weighted line graph with free parking [11, 14]. On the same input, the *harmonic pricing scheme* incurs a cost of n in expectation.

The pricing is designed to emulate the Harmonic Algorithm of Gupta and Lewi [8] for a line metric. This simple algorithm works as follows: let d_l be the distance of the arriving car from the nearest free parking slot to the left of it, and d_r be its distance from the nearest free parking slot to the right. The algorithm assigns the car to the slot on the left with probability $\frac{d_r}{d_l + d_r}$, and to the slot on the right with the remaining probability. Gupta and Lewi proved the following.

THEOREM 5.1. ([8]) *Let d_{\max} be the maximal distance between two vertices, and let d_{\min} be the minimal one. The competitive ratio of the harmonic algorithm on a line metric is $O\left(\log \frac{d_{\max}}{d_{\min}}\right)$.*

Theorem 5.1 yields an $O(\log n)$ approximation algorithm to the min cost matching in the case where quantity d_{\max}/d_{\min} , called the *aspect ratio* of the metric, can be bounded by a polynomial of n . In fact, Gupta and Lewi [8] also show how this algorithm can be combined with the doubling technique so that effectively the aspect ratio is always bounded by a polynomial of n . This relies on continuous estimation of the optimum cost and modifying the metric distances based on that.

We provide two pricing schemes. The price of anarchy of the simpler one matches the competitive ratio guarantee of Harmonic as stated in Theorem 5.1. The more sophisticated pricing scheme matches the refined bound of $O(\log n)$ given an estimate of the (final) optimum social cost a priori.

5.2.1 Computing the posted parking fees We now show how to set prices to the vacant parking slots that effectively makes incoming [selfish] agents park *as though* they were obeying orders given by the Harmonic Algorithm. As the algorithm is randomized, so will be the pricing.

We say that a pricing scheme *ensures harmonic behavior* if, when an agent arrives:

- She chooses to park at her goal if vacant.
- Otherwise, she chooses between the two closest vacant slots to the left and right of her goal with harmonic probabilities.

A maximal contiguous run of occupied parking spaces is called a *block*.

Just before car i is set to park, let $V' = \{v_1, \dots, v_k\}$ the set of vacant parking spots and let $B_1, B_2, \dots, B_\ell \subseteq V$ the current set of “blocks”. We use $L(B_j)$ to denote the first available vertex to the left of block B_j , for some $j \in \{1, \dots, \ell\}$, and let $R(B_j)$ denote the first available vertex to the right of B_j . Let $d_j = d(L(B_j), R(B_j))$, we define a set of distributions D_1, D_2, \dots, D_ℓ , where for every $j \in \{1, \dots, \ell\}$, D_j is uniformly distributed on the real interval $[-d_j, d_j]$.

Let $P: V' \mapsto \mathbb{R}^+$ be a random function that sets prices to vacant slots. Previously, this function was denoted P_i , but we omit the index i for the remainder of this section.

We say that P satisfies the *Harmonic Payment Conditions* if:

1. For any $j \in \{1, \dots, \ell\}$, $P(L(B_j)) - P(R(B_j)) \sim D_j$.
2. For any $u, v \in V'$ s.t. there are no unoccupied parking slots between them: $-d(u, v) < P(u) - P(v) < d(u, v)$.

LEMMA 5.1. *Any pricing function P that satisfies the Harmonic Payment Conditions, ensures harmonic behaviour on the part of arriving cars, regardless of their goals.*

Proof. First, notice that the second condition of the harmonic payment conditions implies that for any $u, v \in V'$, $-d(u, v) < P(u) - P(v) < d(u, v)$, even if there are vacant slots between them. This implies that if an agent’s goal, g , is vacant, then for all $v \in V' \setminus \{g\}$, we have that $P(g) < P(v) + d(g, v)$. Ergo, if the agent’s goal is vacant, she parks there.

Now consider an agent whose goal v is blocked, i.e., $v \in B_j$ for some $j \in \{1, \dots, \ell\}$. Let v_L denote $L(B_j)$ and v_R denote $R(B_j)$. For every vacant slot $v' \in V'$ left of v_L we have $d(v, v_L) + P(v_L) < d(v, v_L) + d(v_L, v') + P(v') = d(v, v') + P(v')$. Symmetrically, for every vacant slot v' right of v_R , we have that $d(v, v_R) + P(v_R) < d(v, v') +$

$P(v')$. Therefore, the agent has only to choose between v_L and v_R .

Let F_j be the density function of distribution D_j . Notice that $F_j(x) = \frac{x+d_j}{2d_j}$. We have that:

$$\begin{aligned} & \Pr[\text{The agent takes } v_L] \\ &= \Pr[d(v, v_L) + P(v_L) \leq d(v, v_R) + P(v_R)] \\ &= \Pr[P(v_L) - P(v_R) < d(v, v_R) - d(v, v_L)] \\ &= \Pr[P(v_L) - P(v_R) < d_j - 2d(v, v_L)] \\ &= F_j(d_j - 2d(v, v_L)) \\ &= \frac{2d_j - 2d(v, v_L)}{2d_j} \\ &= \frac{d(v, v_R)}{d_j}, \end{aligned}$$

where the forth equality is due to the first harmonic payment condition. Therefore, we get that the agent follows the harmonic behavior. \square

Next, we show a specific pricing scheme that satisfies the harmonic payment conditions.

OBSERVATION 5.1. *Let $R(v)$ be the set $\{j \in \{1, \dots, \ell\} : v \text{ is left of } B_j\}$ and let $q_1 \sim D_1, q_2 \sim D_2, \dots, q_\ell \sim D_\ell$. For any constant c , setting $P(v) = c + \sum_{j \in R(v)} q_j$ satisfies the Harmonic Payment Conditions.*

We can use the constant c in order to obtain a pricing scheme that satisfies the harmonic payment conditions in which there are no negative transfers. One can ask for a pricing scheme that optimizes an additional objective: minimizes the sum of prices of all vacant parking slots, the maximal price of a slot, the difference in pricing of the slots over time, etc. We can use an LP to satisfy such objectives. For example, Linear Program 1 minimizes the sum of prices of slots given $q_j \sim D_j$ as prescribed in Observation 5.1. Note that Linear Program 1 is feasible since the payment scheme of Observation 5.1 satisfies all its constraints.

Since we have shown there exists a pricing scheme which makes incoming agents simulate the behavior of the harmonic algorithm, we can use this pricing scheme to coordinate incoming agents, and improve the PoA significantly.

In graphs where the aspect ratio is bounded by some polynomial in n , we get:

COROLLARY 5.1. *Any pricing scheme which satisfies the harmonic payment conditions gives an $O(\log n)$ approximation to the optimal assignment when the aspect ratio is bounded by $\text{poly}(n)$.*

For general weighted line graphs, without further assumptions, we get the following:

Linear Program 1 A linear program for pricing that satisfies the harmonic payment conditions and also minimizes the sum of parking fees over all parked cars. Note that this program actually minimizes the sum of the posted prices of vacant parking spaces, but — any feasible solution to this LP results in the same harmonic behavior so, by repeatedly computing such pricing, we minimize the sum of payments of all parked cars.

$$\min \sum_{v \in V'} p_v \text{ s.t.}$$

- $\forall j \in \{1, \dots, \ell\}, p_{L(B_j)} = p_{R(B_j)} + q_j.$
- $\forall u, v \in V'$ s.t. there is no vacant slot between them, $-d(u, v) < p_u - p_v < d(u, v).$
- $\forall v \in V', p_v \geq 0.$

THEOREM 5.2. Any pricing scheme which satisfies the harmonic payment conditions gives an $O\left(\min\{\log \frac{d_{\max}}{d_{\min}}, n^2\}\right)$ approximation to the optimal assignment in weighted line graphs in expectation.

Proof. The proof that the harmonic algorithm is $O\left(\log \frac{d_{\max}}{d_{\min}}\right)$ competitive is given in [8]. For proving that the competitive ratio is also bounded by $O(n^2)$, we use the same notation as in [8]².

Replacing the final bound on $E[\delta_{\max} | \delta_1]$ in the proof of Lemma 4 [8] by the following one gives us the desired result.

LEMMA 5.2. $E[\delta_{\max} | \delta_1] \leq n^2 \delta_1.$

Proof. First note that δ_{\max} is the distance between two servers on the line, therefore there are at most n^2 possible values for this distance. Let J be the set of possible distances. We get that

$$\begin{aligned} E_{\pi}[\delta_{\max} | \delta_1] &= \sum_{k \in J} k \Pr[\delta_{\max} = k | \delta_1] \leq \sum_{k \in J} k \Pr[\delta_{\max} \geq k | \delta_1] \\ &= \sum_{k \in J} k Q_{\delta_1, k} \leq \sum_{k \in J} k \frac{\delta_1}{k} = \sum_{k \in J} \delta_1 = |J| \delta_1 \\ &\leq n^2 \delta_1, \end{aligned}$$

where $Q_{\delta_1, k} \leq \frac{\delta_1}{k}$ by Lemma 4(ii) of [8]. □ □

5.2.2 Pricing with a Prior In this section we show how to improve the expected PoA via a dynamic pricing scheme when the city planner has an estimate of the cost of the optimal matching. This seems to be a reasonable

²A full version is available at <http://theory.stanford.edu/~klewi/papers/metric-matching-full.pdf> but its analysis of Harmonic is slightly different.

assumption in many scenarios, since the city planner only needs to approximate the optimal matching to within a poly(n) factor in order to devise a pricing scheme that achieves an $O(\log n)$ approximation. This is an improvement over the *harmonic pricing scheme* whenever the aspect ratio is superpolynomial in n .

Gupta and Lewi show that given $Z \in [\text{OPT}, c \cdot \text{OPT}]$ where $c > 1$, one can get an $O(\log cn)$ approximation using the following metric change:

- Disconnect every edge with weight $\geq Z.$
- Change every edge of weight smaller than $Z/(2c \cdot n^2)$ to be of weight *exactly* $Z/(2c \cdot n^2).$

Let d' denote the new metric after the transformation. First, note that the aspect ratio in d' is $2c \cdot n^3$, since the maximal distance is at most Zn , while the minimal one is at least $Z/(2c \cdot n^2)$. Another key observation is that since the optimal matching doesn't use any disconnected edges and the weight of each edge increased by at most $Z/(2c \cdot n^2)$, the cost of every matched server-request in the optimal matching increased by at most $Z/(2c \cdot n)$. Therefore, if OPT denotes both the optimal matching in (V, d) and its cost in (V, d) , then the cost of OPT in (V, d') is at most $\text{OPT} + Z/2c \leq 1.5\text{OPT}$. As the weights in d are no larger than in d' , running the harmonic algorithm in (V, d') yields an $O(\log cn)$ approximation to OPT in $(V, d).$

Thus, given Z , an estimate of the cost of the optimal matching, a dynamic pricing scheme needs to simulate the change in the metric. Specifically, the pricing scheme needs to:

1. Make sure agents don't "jump over" edges of weight bigger than $Z.$
2. Make incoming agents simulate the behavior of the harmonic algorithm according to the transformed metric by acting selfishly.

We say that an online matching algorithm is *monotone* if the following conditions hold:

1. If an incoming request is at a position of an unmatched server, the server will be matched to that request.
2. Let B_j be some block of matched servers, let $v_L = L(B_j)$ and $v_R = R(B_j)$. Consider any two points $a, b \in B_j$ and a request to be made at either of them. Then

$$\begin{aligned} d(a, v_L) \leq d(b, v_L) &\iff \\ \Pr[\text{request at } a \text{ would be matched to } v_L] &\geq \Pr[\text{request at } b \text{ would be matched to } v_L]. \end{aligned}$$

It is easy to see that running the harmonic algorithm on both the original and the transformed metrics yields a monotone algorithm. The following theorem shows that every *monotone* algorithm can be priced:

THEOREM 5.3. *Every monotone online algorithm for the min cost matching problem on a line can be transformed into a dynamic pricing scheme which makes incoming agents simulate the algorithm.*

Proof. Let A be a monotone algorithm for the online min cost matching problem. Let $B_1, B_2, \dots, B_\ell \subseteq V$ be the blocks that represent the current state of the graph, and let $V' = V \setminus \bigcup_{j=1}^\ell B_j$ the set of vacant parking slots. Let $v_L = L(B_j)$ and $v_R = R(B_j)$. Let $p_L^j : B_j \mapsto (0, 1]$ be a probability function that maps every vertex in the block into the probability of the vertex being matched by v_L in A . Let $V_j = \{v_j^1, v_j^2, \dots, v_j^t\} \subseteq B_j$ the set of vertices for which $p_L^j(v_j^1) > p_L^j(v_j^2) > \dots > p_L^j(v_j^t)$. We define a distribution D_j which determines the difference between $P(v_R)$ and $P(v_L)$. For every $i \in \{1, \dots, t\}$, let $\Delta_i = d(v_j^i, v_L) - d(v_j^i, v_R)$.

D_j is defined via the following cumulative density function F_j :

$$F_j(\Delta) = \begin{cases} 0 & \Delta \leq -d_j \\ 1 - p_L^j(v_j^i) & \Delta_i \geq \Delta > \Delta_{i-1} \\ 1 & \Delta > d_j, \end{cases}$$

where $\Delta_0 = -d_j$ and $\Delta_{t+1} = d_j$. From A 's monotonicity, it is clear that $-d_j \leq \Delta_1 \leq \Delta_2 \leq \dots \leq \Delta_t \leq d_j$, and therefore F_j is a valid density function.

We claim that a dynamic pricing scheme $P : V' \mapsto \mathbb{R}^+$ that satisfies the following two conditions makes incoming agents simulate algorithm A by acting selfishly:

1. For any $j \in \{1, \dots, \ell\}$, $P(R(B_j)) - P(L(B_j)) \sim D_j$.
2. For any $u, v \in V'$ s.t. there are no unoccupied slots between them: $-d(u, v) < P(u) - P(v) < d(u, v)$.

Just as in Lemma 5.1, the second condition yields that if an agent's goal vertex is vacant, she will occupy it. Let $v \in B_j$ be the incoming agent's goal for some $j \in \{1, \dots, \ell\}$, and let v_j^i be the first vertex in V_j (the one with minimal index i) such that v is *not right* of it. We have that $\Pr[A \text{ matches } v \text{ to } v_L] = p_L^j(v_j^i)$. The probability that the selfish agent will occupy v_L is:

$$\begin{aligned} \Pr[v \text{ occupies } v_L] &= \Pr[d_L + P(v_L) < d_R + P(v_R)] \\ &= \Pr[P(v_R) - P(v_L) > d_L - d_R] \\ &= \Pr[P(v_R) - P(v_L) > \Delta_i] \\ &= 1 - \Pr[P(v_R) - P(v_L) \leq \Delta_i] \\ &= 1 - F(\Delta_i) \\ &= p_L^j(v_j^i), \end{aligned}$$

Algorithm 5 A dynamic pricing scheme that makes selfish agents simulate algorithm A .

- Let $q_1 \sim D_1, q_2 \sim D_2, \dots, q_\ell \sim D_\ell$.
 - Let $L(v)$ be the set $\{j \in \{1, \dots, \ell\} : v \text{ is right of } B_j\}$.
 - $\forall v \in V'$, set $P(v) = c + \sum_{j \in L(v)} q_j$ for some constant c .
-

where the third equality stems from the fact that $\Delta_i \geq d_L - d_R > \Delta_{i-1}$.

The dynamic pricing scheme given in Algorithm 5 satisfies the conditions for simulating algorithm A . \square

We conclude:

COROLLARY 5.2. *Given an estimate of the cost of the optimal matching by a $\text{poly}(n)$ factor, there exists a payments scheme that gives an $O(\log n)$ approximation to the optimal assignment in weighted line graphs in expectation.*

Proof. Let Z be an estimate, and p be some polynomial such that $\text{OPT}/p(n) \leq Z \leq p(n) \cdot \text{OPT}$. First, apply the metric change using $Z' = Z \cdot p(n)$ as input. Afterwards, use a dynamic pricing scheme that simulates the harmonic algorithm according to the transformed metric. There exists such a pricing scheme according to Theorem 5.3. Since $Z' \in [\text{OPT}, p(n)^2 \text{OPT}]$, this yields an $O(\log(p(n)^2 \cdot n)) = O(\log n)$ approximation to the optimal assignment in expectation. \square

6 Discussion

Observable History For some arbitrary function h , we define the h -observable history after event $i - 1$ to be $h(\bar{\sigma}, \bar{s})$ where $\bar{\sigma} = \sigma_1, \dots, \sigma_{i-1}$ is the sequence of previous agent types and $\bar{s} = s_1, \dots, s_{i-1}$ the sequence of decisions made thus far.

If for some non-invertible function q we have that

$$g(\bar{\sigma}, \bar{s}) = q(g'(\bar{\sigma}, \bar{s})), \quad \forall \bar{\sigma}, \bar{s},$$

we say that g *hides more* than g' .

This suggests a tradeoff between the price of anarchy of a dynamic posted pricing scheme and the information available to it.

We remark that the study of h -observable histories also makes sense for arbitrary online problems, in the context of mechanism design (not restricted to posted prices) and in the non-strategic setting of online algorithms.

Pricing With posted prices, prices are determined before the event occurs.

For non-posted price mechanisms — one issue is “when will the price be determined?” Is the payment

determined promptly (when the agent arrives) or only later? In [7] this issue came up in the context of an online Groves mechanism.

Commitment to pricing, Rate of Price Change Obvious disadvantages of dynamic pricing include (a) prices are unknown, you plan to go downtown but don't know what it will cost you, (b) you may start circling the block waiting for the prices to change. Thus it would be highly advantageous to be able to (a) commit to prices in advance and (b) slow the rate of change.

One interesting model is to have the price function P_i depend, not on the events $\sigma_1, \dots, \sigma_{i-1}$ and decisions s_1, \dots, s_{i-1} , but on a shorter prefix of these sequences. This has the effect of committing to prices in advance. A natural question is how does this impact the price of anarchy?

Simulating Simple Online Algorithms It is not clear if or how one can convert an arbitrary online algorithm to a dynamic pricing scheme. The simulation technique that we've repeatedly used above seems to work when the algorithm is particularly simple. For example, it is not clear how to give dynamic posted prices with good price of anarchy for the problem of metrical matching on metric spaces other than the line. One desirable goal would be to find alternative online algorithms, possibly with higher competitive ratios, and try to find pricing schemes that simulate these algorithms.

Structural Relationships Given an online setting, it is clear that the following weakly increase:

- The competitive ratio of an online algorithm.
- The price of anarchy achievable by a dominant strategy incentive compatible mechanism, where the true agent types are revealed.
- The price of anarchy achievable by a dominant strategy incentive compatible mechanism, when restricted to the current agent type and some h -observable history of previous agents.
- The price of anarchy achievable by a dynamic pricing scheme restricted to h -observable history.

However, it is not clear if there is any stronger relation that holds in general. Moreover, these can be further refined, leading to additional questions, e.g.:

- How does promptness of payment determination impact the price of anarchy of mechanisms?
- Given that a mechanism or pricing schemes uses h -observable history, how does the price of anarchy depend on properties of h ?

References

- [1] Baruch Awerbuch, Yossi Azar, and Adam Meyerson. Reducing truth-telling online mechanisms to online optimization. In *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*, STOC '03, pages 503–510, New York, NY, USA, 2003. ACM.
- [2] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. Online auctions and generalized secretary problems. *SIGecom Exchanges*, 7(2), 2008.
- [3] Allan Borodin, Nathan Linial, and Michael E. Saks. An optimal on-line algorithm for metrical task system. *J. ACM*, 39(4):745–763, 1992.
- [4] Marek Chrobak, Howard J. Karloff, T. H. Payne, and Sundar Vishwanathan. New results on server problems. *SIAM J. Discrete Math.*, 4(2):172–181, 1991.
- [5] Marek Chrobak and Lawrence L. Larmore. An optimal on-line algorithm for k -servers on trees. *SIAM J. Comput.*, 20(1):144–148, 1991.
- [6] Michael Cooper. A meter so expensive, it creates parking spots. <http://nyti.ms/1dIvpA0>, March 2012.
- [7] Eric J. Friedman and David C. Parkes. Pricing wifi at starbucks: issues in online mechanism design. In *EC*, pages 240–241, 2003.
- [8] Anupam Gupta and Kevin Lewi. The online metric matching problem for doubling metrics. In *ICALP*, pages 424–435, 2012.
- [9] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and David C. Parkes. Adaptive limited-supply online auctions. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, EC '04, pages 71–80, New York, NY, USA, 2004. ACM.
- [10] Sandy Irani and Ronitt Rubinfeld. A competitive 2-server algorithm. *Inf. Process. Lett.*, 39(2):85–91, 1991.
- [11] Bala Kalyanasundaram and Kirk Pruhs. Online weighted matching. *J. Algorithms*, 14(3):478–488, 1993.
- [12] Bala Kalyanasundaram and Kirk Pruhs. On-line network optimization problems. In Amos Fiat and Gerhard J. Woeginger, editors, *Online Algorithms, The State of the Art (the book grow out of a Dagstuhl Seminar, June 1996)*, volume 1442 of *Lecture Notes in Computer Science*, pages 268–280. Springer, 1996.
- [13] Richard M. Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An optimal algorithm for on-line bipartite matching. In *STOC*, pages 352–358, 1990.
- [14] Samir Khuller, Stephen G. Mitchell, and Vijay V. Vazirani. On-line algorithms for weighted bipartite matching and stable marriages. *Theor. Comput. Sci.*, 127(2):255–267, 1994.
- [15] Elias Koutsoupias. The k -server problem. *Computer Science Review*, 3(2):105–118, 2009.
- [16] Elias Koutsoupias and Christos Papadimitriou. On the k -server conjecture. *Journal of the ACM*, 42:507–511, 1995.
- [17] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 404–413, 1999.

[18] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *Proceedings of the 2Nd ACM Conference on Electronic Commerce*, EC '00, pages 233–241, New York, NY, USA, 2000. ACM.

[19] Mark S. Manasse, Lyle A. McGeoch, and Daniel Dominic Sleator. Competitive algorithms for server problems. *J. Algorithms*, 11(2):208–230, 1990.

[20] Moni Naor, Benny Pinkas, and Reuban Sumner. Privacy preserving auctions and mechanism design. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, EC '99, pages 129–139, New York, NY, USA, 1999. ACM.

[21] David C. Parkes. Online mechanisms. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 16, pages 411–439. Cambridge University Press, New York, NY, USA, 2007.

[22] D.C. Shoup and American Planning Association. *The High Cost of Free Parking*. Planners Press, American Planning Association, 2005.

A Appendix

A.1 An Illustration of the Fractional Traversal Algorithm The following is an example illustrating the execution of Algorithm 1 given three tasks as input:

Example 1 The fractional traversal algorithm applied to a metrical task system with states $\{1, 2, 3\}$, distances $d_{1,2} = 2$, $d_{1,3} = 3$, $d_{2,3} = 4$, and traversal sequence $\tau = (1, 2, 1, 3)^*$. The example and the underlying metric are also illustrated.

- Task $w^1 = (3, 6, 3)$:
 - $\lambda_1^1 = 2/3$, i.e., process two thirds of task w^1 in state $\tau_1 = 1$, for a work expenditure of $3 \cdot 2/3 = 2$. Increase j by one.
 - $\lambda_2^1 = 1/3$, i.e., process the remaining 1/3 of task w^1 in state $\tau_2 = 2$, for a work expenditure of $6/3 = 2$. At this point the algorithm will increase j by one. Move to state $\tau_3 = 1$.
- Task $w^2 = (1, 3, 4)$: $\lambda_3^2 = 1$, for a work expenditure of 1 in state $\tau_3 = 1$, at this point of time $\rho_j = \rho_3 = 1$.
- Task $w^3 = (10, 10, 10)$:
 - $\lambda_3^3 = 2/10$, for a work expenditure of 2 in state $\tau_3 = 2$. Increase j by one. Move to state $\tau_4 = 3$.
 - $\lambda_4^3 = 3/10$, for a work expenditure of 3 in state $\tau_4 = 3$. Increase j by one. Move to state $\tau_5 = 1$.
 - $\lambda_5^3 = 2/10$, for a work expenditure of 2 in state $\tau_5 = 1$. Increase j by one. Move to state $\tau_6 = 2$.
 - $\lambda_6^3 = 2/10$, for a work expenditure of 2 in state $\tau_6 = 2$. Increase j by one. Move to state $\tau_7 = 1$.
 - $\lambda_7^3 = 1/10$, for a work expenditure of 1 in state $\tau_7 = 1$, at this point of time $\rho_7 = 1$.

