

Competitively Pricing Parking in a Tree

Max Bender
Computer Science Department
University of Pittsburgh
MCB121@pitt.edu

Jacob Gilbert
Computer Science Department
University of Pittsburgh
JMG264@pitt.edu

Aditya Krishnan*
Computer Science Department
Johns Hopkins University
aditya.krishnan94@gmail.com

Kirk Pruhs†
Computer Science Department
University of Pittsburgh
kirk@cs.pitt.edu

October 5, 2020

Abstract

Motivated by demand-responsive parking pricing systems we consider posted-price algorithms for the online metrical matching problem and the online metrical searching problem in a tree metric. Our main result is a poly-log competitive posted-price algorithm for online metrical searching.

1 Introduction

Since 2011 SFpark has been San Francisco’s system for managing the availability of on-street parking [3, 27, 2]. The goal of the system is to reduce the time and fuel wasted by drivers searching for an open space. The system monitors parking usages via sensors embedded in the pavement and distributes this information in real-time to drivers via SFpark.org and phone apps. SFpark periodically adjusts parking meter pricing to manage demand, to lower prices in underutilized areas, and to raise prices in overutilized areas. Prices can range from a minimum of 25 cents to a maximum of 7 dollar per hour during normal hours with a 18 dollars per hour cap for special events such as baseball games or street fairs. Several other cities in the world have similar demand-responsive parking pricing systems, for example Calgary has had the ParkPlus system since 2008 [1].

The problem of centrally assigning drivers to parking spots to minimize time and fuel usage is naturally modeled by the online metrical matching problem. The setting for online metrical matching consists of a collection of k servers (the parking spots) located at various locations within a metric space. The algorithm then sees an online sequence of requests over time that arrive at various locations in the metric space (the drivers arriving to look for a parking spot). In response to a request, the online algorithm must match the request (car) to some server (parking spot) that has not been previously matched; Conceptually we interpret this matching as the request (car) moving to the location of the matched server (parking spot). The objective goal is to minimize the aggregate distance traveled by the requests (cars).

*This work was done in part while this author was a student at Carnegie Mellon University

†Supported in part by NSF grants CCF-1421508 and CCF-1535755, and an IBM Faculty Award.

We also consider what we call the online metrical search problem, which is an important special case of the online metrical matching problem. This is a promise problem in that the adversary is constrained to guarantee that there is an optimal matching for which only one edge has positive cost. It is useful to conceptually think of online metrical search as the following parking problem: the setting consists of many parking spots at various locations in a metric space and a single car that is initially parked at some location in the metric space. Over time the parking spots are decommissioned one by one until only one parking spot is left in commission. If at any time the car is not parked at an in-commission parking spot, then the car must move to a parking spot that is still in commission. The objective is to minimize the aggregate distance traveled by the car. The optimal solution is to move the car directly to the last remaining parking spot.

The online metrical search problem is a special case of the online metrical matching problem because the parking spots can be viewed as servers and the decommissioning of a parking spot can be simulated by the arrival of a request at the location of that parking spot. So a lower bound on the competitive ratio for the online metrical search problem for a particular metric space also gives a lower bound for the online metrical matching problem on the metric space. Conversely it seems that in terms of the optimal competitive ratio, online metric search is no easier than metric matching. In particular, there is no known example of a metric space where the optimal competitive ratio for online metrical matching is known to be significantly greater than the optimal competitive ratio for online metrical search on that metric space. For example on a line metric, the online metrical search problem is better known as the “cow path problem”, and the optimal deterministic competitive ratio is known to be 9 [12], while the best known lower bound on the deterministic competitive ratio for online metrical matching on a line metric is 9.001 [17], worse only by a minuscule factor.

In order to be implementable within the context of SFpark, online algorithms must be posted-price algorithms. In this setting, posted-price means that before each request arrives, the online algorithm sets a price on each unused server (parking spot) without knowing the location where the next request will arrive. Furthermore, each request is assumed to be a selfish agent who moves to the available server (parking spot) that minimizes the sum of the price of and distance to that server. The objective remains to minimize the aggregate distance traveled by the requests. So conceptually the objective of the parking pricing agency is minimizing social cost, not maximizing revenue.

Research into posted-price algorithms for online metrical matching was initiated in [13] as part of a line of research to study the use of posted-price algorithms to minimize social cost in online optimization problems. As a posted-price algorithm is a valid online algorithm, one can not expect to obtain a better competitive ratio for posted-price algorithms than what is achievable by online algorithms. So this research line has primarily focused on problems where the optimal competitive ratio achievable by an online algorithm is (perhaps approximately) known and seeks to determine whether a similar competitive ratio can be (again perhaps approximately) achieved by a posted-price algorithm. The higher level goal is to determine the increase in social cost that is necessitated by the restriction that an algorithm has to use posted prices to incentivize selfish agents, instead of being able to mandate agent behavior.

An $O(\log \Delta)$ -competitive randomized posted-price algorithm for metric matching on a line metric is given in [13] where Δ is the ratio of the distance between the furthest two servers and the distance between the closest two servers. No $o(\log k)$ -competitive (not necessarily posted-price) algorithm is known for online metric matching on a line metric. So arguably, on a line metric there is a posted-price algorithm that is nearly as competitive as the best known centralized online algorithm.

Our original research goal was to determine whether posted-price algorithms can be similarly

competitive with a centralized online algorithm for tree metrics for online metrical matching. In order to be more specific about our goal, we need to review a bit. A tree metric is represented by a tree $T = (V, E)$ with positive real edge weights where the distance $d_T(u, v)$ between vertices $u, v \in V$ is the shortest path between vertices u and v in T . There is a deterministic online algorithm that is $(2k - 1)$ -competitive for online metric matching in any metric space, and no deterministic online algorithm can achieve a better competitive ratio for online metric searching in a tree metric [20, 21]. An $O(\log k)$ -competitive randomized algorithm for online metric matching in $O(\log k)$ -HST's (Hierarchically Separated Trees) is given in [24]. By combining this result with results about randomly embedding metric spaces into HST's [10, 11, 15], [24] obtained an $O(\log^3 k)$ -competitive randomized algorithm for online metric matching in a general metric space. Following this general approach, [9] later obtained an $O(\log^2 k)$ -competitive randomized algorithm for online metrical search in an arbitrary metric by giving an $O(\log k)$ -competitive randomized algorithm for 2-HST's. No better results are known for tree metrics, so all evidence points to tree metrics as being as hard as general metrics for online metrical matching. Thus, more specifically our original research goal was to determine whether there is poly-log competitive randomized posted-price algorithm for the online metrical matching problem on a tree metric. Before stating our progress toward this goal, it will be useful to review the literature a bit more.

1.1 Prior Related Work

The most obvious algorithmic design approach for posted-price problems is to directly design a pricing algorithm from scratch, as is done for metrical task systems in [13], but this is not the most common approach in the literature. Two less direct algorithmic design paradigms have emerged in the literature. The first algorithmic design paradigm is what we will call *mimicry*. A posted-price algorithm A *mimics* an online algorithm B if the probability that B will take a particular action is equal to the probability that a self-interested agent will choose this same action when the prices of actions are set using A . For example, [13] shows how to set prices to mimic the $O(\log \Delta)$ -competitive Harmonic algorithm for online metric matching on a line metric from [18]. As another example, [16] shows how to set prices to mimic the $O(1)$ -competitive algorithm Slow-Fit from [7, 8] for the problem of minimizing makespan on related machines. However, for some problems it is not possible to mimic known competitive algorithms using posted prices. For such problems, another algorithmic design paradigm is what we will call *monotonization*. In the *monotonization* algorithm design approach, one first seeks to characterize the online algorithms that can be mimicked, and then designs such an online algorithm. In the examples in the literature, this characterization involves some sort of monotonicity property. For example, *monotonization* is used in [13] to obtain an $O(k)$ -competitive posted-price algorithm for the k -server problem on a line metric, and in [14] to obtain an $O(k)$ -competitive posted-price algorithm for the k -server problem on a tree metric. Since no deterministic algorithm can be better than k -competitive for the k -server problem in any metric [23], this shows that in these settings, there is minimal increase in social cost necessitated by the use of posted-prices. As another example, *monotonization* is used in [19] to obtain an $O(1)$ -competitive posted-price algorithm for minimizing maximum flow time on related machines.

For online metric matching on a line metric, better competitive ratios are achievable. An $O(k^{.59})$ -competitive deterministic online algorithm was given in [4]. Subsequently several different $O(\log n)$ -competitive randomized online algorithms for a line are given in [18]; these algorithms leverage special properties of HST's constructed from a line metric. As already mentioned, [18] also showed that the natural Harmonic algorithm is $O(\log \Delta)$ -competitive. An $O(\log^2 k)$ -competitive deterministic online algorithm was given in [25], and this was later improved to $O(\log k)$ in [26]. Super-constant lower bounds for various types of algorithms are given in [5, 22]. More generally,

the algorithm for online metric matching given in [25] has the property that for every metric space, its competitive ratio is at most $O(\log^2 k)$ times the optimal competitive ratio achievable by any deterministic algorithm on that metric space.

1.2 Our Contribution

There is no hope to mimic any of the online algorithms for online metrical matching that are based on HST's as HST's by their very nature lose too much information about the structure of a tree metric. Therefore we adopt the monotonization approach. In Section 2 we identify a monotonicity property that characterizes mimicable algorithms for online metrical matching in tree metrics. Roughly speaking this property says that if a request were to have arrived on the route to its desired server, then the probability that the request would still have been matched to this server can not decrease. Thus we reduce finding a post-priced algorithm to finding a monotone algorithm.

In Section 3 we give an algorithm **TreeSearch** for the online metrical search problem on a tree metric. The algorithm is based on the classic multiplicative weights algorithm for online learning from experts [6]. Conceptually there is one expert E^ℓ for each leaf ℓ of the tree T . Expert E^ℓ always recommends that the car/request travels toward the leaf ℓ . Thus expert E^ℓ pays a cost of one whenever a parking spot on the path from the root to ℓ is decommissioned, a cost of zero when other parking spots are decommissioned, and an infinite cost if there are no remaining parking spots on the path from the root to ℓ . Let π_t^ℓ be the probability that the multiplicative weights algorithm has associated with expert E^ℓ right before request r_t arrives. Let v_t^ℓ be the location of the car just before request r_t arrives if the advice of expert E^ℓ had always been followed. The algorithm **TreeSearch** maintains the invariant that right before request r_t arrives, the probability that the car is at a vertex v is $\sum_{\ell: v_t^\ell = v} \pi_t^\ell$, the sum of the probabilities of the experts that recommend that the car should be parked at v . The most technically difficult part of the algorithm design process was maintaining this invariant. We then upper bound the expected number of jumps made by the **TreeSearch** algorithm, where a jump is a movement of the car by a positive amount. Finally, we show how to extend **TreeSearch** to be a monotone algorithm **TreeMatch** for online metrical matching on a tree metric.

In Section 4 we give an algorithm for online metric searching on a tree metric. Before any requests arrive, an algorithm **GroveBuild** embeds the tree metric into what we will call a grove, which is a refinement of an HST that retains more information about the topology of the original metric space. It is probably easiest to explain what a grove is by explaining the difference in how one is constructed in comparison to how an HST is constructed. The construction of each starts with a Low Diameter Decomposition (LDD) of the metric space. A LDD is a partition $\mathcal{P} = \{P_1, \dots, P_n\}$ of the vertices of the metric space where each part is connected and the diameter of each part is an α factor smaller than the diameter of the whole metric space. The top of the HST consists of a star where the center of the star is the root of the HST, and there is one child of the root for each part P_i . In contrast, the top of a grove consists of the tree that remains after collapsing each part to a single vertex. For both an HST and a grove, the construction then proceeds recursively on each part. So intuitively the key difference is that groves retain information about the distances between parts in the LDD that the HST instead discards. See Figure 1 for a comparison of an HST and a grove constructed from the same LDD.

We then give a monotone algorithm **GroveMatch** for online metrical matching on a tree metric that utilizes the algorithm **TreeMatch** on each tree in the grove constructed from the tree metric. We show that **GroveMatch** is poly-log competitive (more precisely $O(\log^6 \Delta \log^2 n)$ -competitive) on metric search instances by induction on the levels of the grove. This is an extension of a

similar induction argument in [24] that shows that a $O(\log n)$ -competitive algorithm for a star (or a complete unit metric) can be extended to an algorithm for a $O(\log n)$ -HST with the loss of a poly-log factor in the competitiveness. However, our situation is complicated by the fact the possible ways that a request can potentially move within a grove is more complicated than the possible ways a request can move within an HST, and thus the induction is more complicated as the induction depends on when the request is moving “up” and when the request is moving “down” in trees within the grove. The bound on the number of jumps made by **TreeSearch** translates to a bound on the number of recursive calls made by **GroveMatch**. There is not a lot of wiggle room in our analysis, and thus both the algorithm design and algorithm analysis process are necessarily quite delicate. For example, if **TreeSearch** made just 1% more jumps than the bound that we can show, then the resulting competitiveness of **GroveMatch** would not be poly-logarithmic. One consequence of this delicateness is that we can not use a black box LDD construction to build our grove, we need to construct our LDD in a way that tightly controls the variance of random properties of our grove.




Figure 1: An example of a LDD, the corresponding HST, and the corresponding grove.

2 Pricing Monotone Algorithms

In this section, we show that an algorithm for the online metrical matching can be implemented as a posted-price algorithm if and only if the algorithm satisfies the following monotonicity property. Intuitively, an algorithm is monotone if a request and server pair get matched with nondecreasing probability as the request is dragged towards the server. We note that monotonicity does not have a natural interpretation within the context of online metrical searching, which explains why we give a monotone algorithm for online metrical matching, even though we only analyze its competitiveness for online metrical search.

Definition 1. An algorithm A for online metric matching is monotone if for every instance, every request r_t in that instance, every possible sequence R of random events internal to A prior to r_t 's arrival, and all vertices u, v, s where v is on the path from u to s it is the case that: $\Pr[A_R(r_t) = s \mid E_R \text{ and } r_t = u] \leq \Pr[A_R(r_t) = s \mid E_R \text{ and } r_t = v]$ where $A_R(r_t) = s$ is the event that A matches r_t to s , and E_R is the event that the past random events internal to A are equal to R .

Theorem 2. Any algorithm A for the online metrical matching problem can be implemented as a posted-price algorithm if and only if A is monotone.

In section 2.1, we show this correspondence for the deterministic setting and then extend it to randomized algorithms in section 2.2. Lastly, in section 2.3 we show that all pricing schemes induce monotone matching algorithms, giving us the equivalency of pricing schemes and monotone matching algorithms. For notational simplicity we use $d(u, v)$ instead of $d_T(u, v)$ in this section.

2.1 Pricing Deterministic Monotone Algorithms

We first need to define monotone partitions.

Definition 3. A monotone partition of the tree metric space $T = (V, E, d_T)$ consists of two components. The first component is a partition $P = \{Q_1, \dots, Q_t\}$ of the vertices V , such that for each part $Q_i \in P$ it is the case that the induced subgraph on Q_i is connected. The second component consists of a designated leader for each nonempty $Q_i \in P$, where a leader is an available server located in Q_i .

A monotone algorithm for serving a request can be simply derived from a monotone partition by matching a request in each part to that part's designated leader. The converse is proved in the course of the proof of lemma 4.

Lemma 4. Every monotone deterministic algorithm \mathcal{A} for online metrical matching can be implemented by a pricing algorithm \mathcal{B} .

Proof. We first explain how to derive monotone partition P from \mathcal{A} at each time step t . For each $s_i \in S_t$, let part Q_i consist of the vertices v such that a request arriving on vertex v would be served by s_i . The leader of each nonempty Q_i is s_i . The fact that each part Q_i induces a connected subgraph follows directly from the monotonicity of the algorithm. Let $u_{i,j}$ be the vertex in Q_i closest to s_j .

We now define the pricing scheme $p : S_j \rightarrow \mathbb{R}$ for time step t for pricing algorithm \mathcal{B} .

```

Set  $p(s_i) = 0$  for an arbitrary nonempty partition  $Q_i \in P$ 
for every nonempty  $Q_j$  whose leader is not yet priced and that is adjacent to a part  $Q_i$  whose
  leader is already priced do
|   Set  $p(s_j) = p(s_i) + d(u_{i,j}, s_i) - d(u_{j,i}, s_j)$ 
  end
for any server  $s_i \in S_t$  that is not already priced do
|   Set  $p(s_i) = +\infty$ 
  end

```

To show that that this pricing scheme implements \mathcal{A} , we will show that if \mathcal{A} services a request at a vertex v by server s_i then $d(v, s_i) + p(s_i) < d(v, s_j) + p(s_j)$ for all $j \neq i$. Assume otherwise to reach a contradiction. Let s_j be the server that minimizes $d(v, s_j) + p(s_j)$ (with ties broken arbitrarily). We now break the proof into cases.

In the first case assume that Q_i and Q_j are adjacent parts, or equivalently that the edge $(u_{i,j}, u_{j,i})$ connects Q_i and Q_j . Then we can conclude that:

$$\begin{aligned}
p(s_i) &= p(s_j) + d(u_{j,i}, s_j) - d(u_{i,j}, s_i) && \text{by definition of } p(s_i) \\
&\leq p(s_i) + d(v, s_i) - d(v, s_j) + d(u_{j,i}, s_j) - d(u_{i,j}, s_i) && \text{by assumption} \\
&= p(s_i) + d(v, s_i) - d(v, u_{j,i}) - d(u_{i,j}, s_i) && \text{as } u_{j,i} \text{ is on the path from } v \text{ to } s_j \\
&\leq p(s_i) + d(v, u_{i,j}) - d(v, u_{j,i}) && \text{by triangle inequality} \\
&\leq p(s_i) - d(u_{i,j}, u_{j,i}) && \text{as } u_{i,j} \text{ is on the path from } v \text{ to } u_{j,i} \\
&< p(s_i) && \text{by definition of metric space}
\end{aligned}$$

Note that the first inequality holds independently of which of Q_i and Q_j was priced first, so our assumption causes a contradiction.

In the second case assume that Q_j and Q_i are not adjacent, i.e. that there is not an edge of the form $(u_{i,j}, u_{j,i})$ in T . However, let Q_k be the last nonempty part before Q_j on the unique path from s_i to s_j in T . Thus the edge $(u_{k,j}, u_{j,k})$ exists in T . Then we can conclude that:

$$\begin{aligned}
d(v, s_k) + p(s_k) &\leq d(v, u_{k,j}) + d(u_{k,j}, s_k) + p(s_k) && \text{by triangle inequality} \\
&= d(v, u_{k,j}) + d(u_{k,j}, s_k) \\
&\quad + d(u_{j,k}, s_j) - d(u_{k,j}, s_k) + p(s_j) && \text{by definition of } p(s_k) \\
&= d(v, u_{k,j}) + d(u_{j,k}, s_j) + p(s_j) \\
&= d(v, u_{k,j}) + d(u_{k,j}, s_j) - d(u_{j,k}, u_{k,j}) + p(s_j) && \text{as } u_{j,k} \text{ is on the path} \\
&\quad \text{from } u_{k,j} \text{ to } s_j \\
&< d(v, u_{k,j}) + d(u_{k,j}, s_j) + p(s_j) && \text{by definition of metric space} \\
&= d(v, s_j) + p(s_j) && \text{as } u_{k,j} \text{ is on the path} \\
&\quad \text{from } v \text{ to } s_j.
\end{aligned}$$

This again is a contradiction to our minimality assumption for s_j . \square

2.2 Randomized Algorithms

From a randomized monotone algorithm \mathcal{A} we derive a corresponding distribution \mathcal{P} over monotone partitions such that can implement \mathcal{A} by picking a random monotone partition P from \mathcal{P} , and setting prices as described in subsection 2.1. We will need the following definitions.

Definition 5.

- Let π_i^v be the probability that algorithm \mathcal{A} matches a request at v to s_i .
- Let $\Pr_{\mathcal{P}}(P)$ denote the probability of partition P under distribution \mathcal{P} .
- Let $v \rightarrow_P s$ denote that it is the case that in monotone partition P , server s is the leader of the part Q where $v \in Q \in P$.
- Let $v \rightarrow_{\mathcal{P}} s$ be the event that in a sampled $P \sim \mathcal{P}$ it holds that $v \rightarrow_P s$.
- Let $\mathcal{P}_i(w)$ be the monotone partitions in \mathcal{P} that have s_i as the leader of the part containing w .

Our goal for the rest of this subsection is to prove the following lemma, which asserts the existence of an appropriate \mathcal{P} .

Lemma 6. Consider a monotone matching algorithm \mathcal{A} for the online metrical matching on a tree T , a particular request r_t , and a particular collection S of remaining available servers. Then there exists a distribution \mathcal{P} over monotone partitions of T such that: $\Pr(w \rightarrow_{\mathcal{P}} s_i \mid r_t = w) = \pi_i^w$ for all servers $s_i \in S$ and all vertices w .

2.2.1 The Construction of \mathcal{P}

Without loss of generality it is sufficient to take $t = 1$. The proof will be via induction on the number of vertices in T . The base case is the case where T is a single vertex v containing all n available servers on it. Then \mathcal{P} consists of n monotone partitions, where monotone partition P_i consists of one part containing vertex v with server s_i as the leader, and with associated probability π_i^v .

For the inductive step, pick an arbitrary leaf u . Let v be the unique neighbor of u in T . By renumbering assume servers s_1, \dots, s_m are located at u . Let T' be the tree derived from T by deleting u , and moving servers s_1, \dots, s_m to v . By induction there exists a probability distribution \mathcal{P}' over monotone partitions of T' such that

$$\sum_{P \in \mathcal{P}'(w)} \Pr_{\mathcal{P}'}(P) = \pi_i^w \quad (1)$$

for all vertices w in T' and servers s_i . We now obtain \mathcal{P} from \mathcal{P}' by extending each monotone partition P in \mathcal{P}' to a collection of monotone partitions in \mathcal{P} . So consider an arbitrary $P \in \mathcal{P}'$. We consider two cases.

Case 1: P is of type 1 if $P \in \mathcal{P}'_i(v)$ for some $1 \leq i \leq m$; That is, one of s_1, \dots, s_m is the leader of the part Q satisfying $v \in Q \in P$. There will be one partition P_1 in \mathcal{P} derived from P . The partition P_1 is identical to P except that the vertex u is added to the part Q . P_1 inherits the probability of P , that is:

$$\Pr_{\mathcal{P}}(P_1) = \Pr_{\mathcal{P}'}(P). \quad (2)$$

Case 2: P is of type 2 if $P \in \mathcal{P}'_i(v)$ for some $i > m$; That is, one of s_{m+1}, \dots, s_n is the leader of the part Q satisfying $v \in Q \in P$. There will be $m+1$ partitions P_1, \dots, P_{m+1} in \mathcal{P} derived from P . For $j \in [1, m]$ the partition P_j is identical to P except that P_j contains a new part consisting of only the vertex u with the leader of this part being s_j . Partition P_{m+1} is identical to P except that in P_{m+1} the part Q satisfying $v \in Q \in P$ also contains the vertex u .

To set the probabilities on P_1, \dots, P_{m+1} in \mathcal{P} let $\delta_k = \pi_k^u - \pi_k^v$ for $k \in [1, m]$ and let $\Delta = \sum_{k=1}^m \delta_k$. For each $j \in [1, m]$, we set

$$\Pr_{\mathcal{P}}(P_j) = \Pr_{\mathcal{P}'}(P) \left(\frac{(\pi_j^u - \pi_j^v)(\pi_i^v - \pi_i^u)}{\Delta \cdot \pi_i^v} \right) \quad (3)$$

and we set

$$\Pr_{\mathcal{P}}(P_{m+1}) = \Pr_{\mathcal{P}'}(P) \frac{\pi_i^u}{\pi_i^v} \quad (4)$$

The following two observations ensure that the probability of P_1, \dots, P_{m+1} are well-defined and that \mathcal{P} is indeed a probability distribution.

Observation 7. For each $k \in \{1, \dots, m\}$, we have that $\pi_k^u \geq \pi_k^v$; for each $k \in \{m+1, \dots, n\}$, we have that $\pi_k^v \geq \pi_k^u$.




Figure 2: An example of case 2 P_j for $1 \leq j \leq m$ and P_{m+1}

Proof. This follows from the monotonicity of \mathcal{A} ; it is assumed that s_1, \dots, s_m are all on vertex u , so for $k \in \{1, \dots, m\}$ it must be the case that $\pi_k^u \geq \pi_k^v$. Since u is a leaf vertex and v is u 's only neighbor, v must be on the path from u to any other server. Thus $\pi_k^v \geq \pi_k^u$ for all $k \in \{m+1, \dots, n\}$. \square

Observation 8. $\sum_{j=1}^{m+1} \Pr_{\mathcal{P}}(P_j) = \Pr_{\mathcal{P}'}(P)$.

Proof.

$$\begin{aligned}
\sum_{j=1}^{m+1} \Pr_{\mathcal{P}}(P_j) &= \sum_{j=1}^m \Pr_{\mathcal{P}'}(P) \left(\frac{(\pi_j^u - \pi_j^v)(\pi_i^v - \pi_i^u)}{\Delta \cdot \pi_i^v} \right) + \Pr_{\mathcal{P}'}(P) \frac{\pi_i^u}{\pi_i^v} && \text{by Eq.s 3 and 4} \\
&= \Pr_{\mathcal{P}'}(P) \frac{\pi_i^v - \pi_i^u}{\Delta \cdot \pi_i^v} \sum_{j=1}^m (\pi_j^u - \pi_j^v) + \Pr_{\mathcal{P}'}(P) \frac{\pi_i^u}{\pi_i^v} \\
&= \Pr_{\mathcal{P}'}(P) \frac{\pi_i^v - \pi_i^u}{\Delta \cdot \pi_i^v} \Delta + \Pr_{\mathcal{P}'}(P) \frac{\pi_i^u}{\pi_i^v} && \text{by definition of } \Delta \\
&= \frac{\Pr_{\mathcal{P}'}(P)}{\pi_i^v} (\pi_i^v - \pi_i^u + \pi_i^u) \\
&= \frac{\Pr_{\mathcal{P}'}(P)}{\pi_i^v} (\pi_i^v) = \Pr_{\mathcal{P}'}(P).
\end{aligned}$$

\square

2.2.2 The analysis of \mathcal{P}

We now turn to proving that our constructed \mathcal{P} has the desired properties.

Definition 9. We partition the support of \mathcal{P} as follows

$$\text{supp}(\mathcal{P}) = \bigcup_{i=1}^m \Phi_i \cup \left(\bigcup_{i=m+1}^n \left(\Phi_i \cup \left(\bigcup_{j=1}^m \Pi_i^j \right) \right) \right)$$

where

$$\Phi_i = \{P_1 \mid P \in \mathcal{P}'_i(v)\} \tag{5}$$

for each $1 \leq i \leq m$,

$$\Pi_i^j = \{P_j \mid P \in \mathcal{P}'_i(v)\} \quad (6)$$

for $i \in \{m+1, \dots, n\}$ and $j \in \{1, \dots, m\}$, and

$$\Phi_i = \{P_{m+1} \mid P \in \mathcal{P}'_i(v)\} \quad (7)$$

for $i \in \{m+1, \dots, n\}$.

Lemma 10. For $i \in \{1, \dots, m\}$, we have that $\sum_{P \in \Phi_i} \Pr_{\mathcal{P}}(P) = \pi_i^v$.

Proof.

$$\begin{aligned} \sum_{P \in \Phi_i} \Pr_{\mathcal{P}}(P) &= \sum_{P \in \mathcal{P}'_i(v)} \Pr_{\mathcal{P}}(P_1) && \text{by Eq. 5} \\ &= \sum_{P \in \mathcal{P}'_i(v)} \Pr_{\mathcal{P}'}(P) && \text{by Eq. 2} \\ &= \pi_i^v && \text{by Eq. 1} \end{aligned}$$

□

Lemma 11. For $i \in \{m+1, \dots, n\}$, we have that $\sum_{P \in \Phi_i} \Pr_{\mathcal{P}}(P) = \pi_i^u$.

Proof.

$$\begin{aligned} \sum_{P \in \Phi_i} \Pr_{\mathcal{P}}(P) &= \sum_{P \in \mathcal{P}'_i(v)} \Pr_{\mathcal{P}}(P_{m+1}) && \text{by Eq. 7} \\ &= \sum_{P \in \mathcal{P}'_i(v)} \Pr_{\mathcal{P}'}(P) \frac{\pi_i^u}{\pi_i^v} && \text{by Eq. 4} \\ &= \frac{\pi_i^u}{\pi_i^v} \sum_{P \in \mathcal{P}'_i(v)} \Pr_{\mathcal{P}'}(P) && \text{by Eq. 4} \\ &= \frac{\pi_i^u}{\pi_i^v} \pi_i^v && \text{by Eq. 1} \\ &= \pi_i^u \end{aligned}$$

□

Lemma 12. For $i \in \{m+1, \dots, n\}$, we have that $\sum_{j=1}^m \sum_{P \in \Pi_i^j} \Pr_{\mathcal{P}}(P) = \pi_i^v - \pi_i^u$.

Proof.

$$\begin{aligned}
\sum_{j=1}^m \sum_{P \in \Pi_i^j} \Pr(P) &= \sum_{j=1}^m \sum_{P \in \mathcal{P}'_i(v)} \Pr(P_j) && \text{by Eq. 6} \\
&= \sum_{j=1}^m \sum_{P \in \mathcal{P}'_i(v)} \Pr(P) \left(\frac{(\pi_j^u - \pi_j^v)(\pi_i^v - \pi_i^u)}{\Delta \cdot \pi_i^v} \right) && \text{by Eq. 3} \\
&= \sum_{j=1}^m \left(\frac{(\pi_j^u - \pi_j^v)(\pi_i^v - \pi_i^u)}{\Delta \cdot \pi_i^v} \right) \sum_{P \in \mathcal{P}'_i(v)} \Pr(P) \\
&= \sum_{j=1}^m \left(\frac{(\pi_j^u - \pi_j^v)(\pi_i^v - \pi_i^u)}{\Delta \cdot \pi_i^v} \right) \pi_i^v && \text{by Eq. 1} \\
&= \frac{\pi_i^v - \pi_i^u}{\Delta} \sum_{j=1}^m (\pi_j^u - \pi_j^v) \\
&= \frac{\pi_i^v - \pi_i^u}{\Delta} \Delta && \text{by definition of } \Delta \\
&= \pi_i^v - \pi_i^u.
\end{aligned}$$

□

Proof of lemma 6. By Observation 8 we know that $\Pr(w \rightarrow_{\mathcal{P}} s_i \mid r_t = w) = \pi_i^w$ holds for all vertices w except for possibly for vertex u . Thus we just need to verify that this holds for $w = u$.

For $i \in \{1, \dots, m\}$, we have that

$$\begin{aligned}
\sum_{P \in \mathcal{P}_i(u)} \Pr(P) &= \sum_{P \in \Phi_i} \Pr(P) + \sum_{j=m+1}^n \sum_{P \in \Pi_i^j} \Pr(P) \\
&= \sum_{P \in \mathcal{P}'_i(v)} \Pr(P_1) + \sum_{j=m+1}^n \sum_{P \in \mathcal{P}'_j(v)} \Pr(P_i) && \text{by Eq.s 5 and 6} \\
&= \sum_{P \in \mathcal{P}'_i(v)} \Pr(P) + \sum_{j=m+1}^n \sum_{P \in \mathcal{P}'_j(v)} \Pr(P) \left(\frac{(\pi_i^u - \pi_i^v)(\pi_j^v - \pi_j^u)}{\Delta \cdot \pi_j^v} \right) && \text{by Eq.s 2 and 3} \\
&= \pi_i^v + \sum_{j=m+1}^n \sum_{P \in \mathcal{P}'_j(v)} \Pr(P) \left(\frac{(\pi_i^u - \pi_i^v)(\pi_j^v - \pi_j^u)}{\Delta \cdot \pi_j^v} \right) && \text{by Eq. 1} \\
&= \pi_i^v + \sum_{j=m+1}^n \left(\frac{(\pi_i^u - \pi_i^v)(\pi_j^v - \pi_j^u)}{\Delta \cdot \pi_j^v} \right) \sum_{P \in \mathcal{P}'_j(v)} \Pr(P) \\
&= \pi_i^v + \sum_{j=m+1}^n \left(\frac{(\pi_i^u - \pi_i^v)(\pi_j^v - \pi_j^u)}{\Delta \cdot \pi_j^v} \right) \pi_j^v && \text{by Eq. 1} \\
&= \pi_i^v + \frac{\pi_i^u - \pi_i^v}{\Delta} \sum_{j=m+1}^n (\pi_j^v - \pi_j^u) \\
&= \pi_i^v + \frac{\pi_i^u - \pi_i^v}{\Delta} \sum_{j=1}^m (\pi_j^u - \pi_j^v) && \text{as both } \pi^u \text{ and } \pi^v
\end{aligned}$$

are distributions
over $1, \dots, n$

$$\begin{aligned}
&= \pi_i^v + \frac{\pi_i^u - \pi_i^v}{\Delta} \Delta \\
&= \pi_i^v + \pi_i^u - \pi_i^v = \pi_i^u
\end{aligned}$$

For $i \in \{m+1, \dots, n\}$, we have that

$$\begin{aligned}
\sum_{P \in \mathcal{P}_i(u)} \Pr_{\mathcal{P}}(P) &= \sum_{P \in \Phi_i} \Pr_{\mathcal{P}}(P) \\
&= \sum_{P \in \mathcal{P}'_i(v)} \Pr_{\mathcal{P}}(P_{m+1}) && \text{by Eq. 7} \\
&= \sum_{P \in \mathcal{P}'_i(v)} \Pr_{\mathcal{P}'}(P) \frac{\pi_i^u}{\pi_i^v} && \text{by Eq. 4} \\
&= \frac{\pi_i^u}{\pi_i^v} \sum_{P \in \mathcal{P}'_i(v)} \Pr_{\mathcal{P}'}(P) \\
&= \frac{\pi_i^u}{\pi_i^v} \pi_i^v = \pi_i^u && \text{by Eq. 1}
\end{aligned}$$

This gives us that $\Pr(u \rightarrow_{\mathcal{P}} s_i) = \pi_i^u$ for all $i \in \{1, \dots, n\}$. \square

2.3 Pricing Schemes Induce Monotone Matching Algorithms

In this section, we show matching requests in an instance of online metrical matching according to a pricing scheme gives us a monotone algorithm. Specifically, given servers S and a pricing scheme \mathcal{P} , let p be the pricing function created by \mathcal{P} prior to the arrival of the first request, and let $f_1 : V \rightarrow S$ be the matching function such that $f(v) = \min_{s \in S} d(v, s) + p(s)$. Since any instance of online metrical matching after j servers have been used is equivalent to an instance of online metrical matching with $n - j$ initial available servers, it's sufficient to just show that matching according to the pricing scheme is monotone on the first request. Now, let $s \in S$ and $u, v \in V$ such that v is on the path from u to s . Then, since $d(v, s) \leq d(u, s)$, we have that $\Pr(f(u) = s) \leq \Pr(f(v) = s)$. Thus f is a monotone matching function. Since the associated matching function of any pricing scheme is monotone, this gives us the equivalence of pricing schemes and monotone matching algorithms for online metrical matching.

3 The Algorithm TreeMatch

In subsection 3.1 we define algorithm **TreeSearch** for the metric search problem on a tree $T = (V, E)$ rooted at vertex a . The distance metric on T will not be of interest to us in this section. We will use the interpretation of a car moving when its parking spot is decommissioned, as introduced earlier, as we think that this interpretation is more intuitive. The description of **TreeSearch** in subsection 3.1 uses a probability distribution $q_t^\sigma(\tau)$ that is complicated to define, so its exact definition is postponed until subsection 3.2, in which we also show that it achieves our goal of matching the experts distribution. In subsection 3.3 we analyze the number of jumps used by the **TreeSearch** algorithm. Finally in subsection 3.4, we show how to convert **TreeSearch** into a monotone algorithm **TreeMatch** for online metrical matching that is identical to **TreeSearch** on online metrical search instances.

3.1 Algorithm Description

We start with some needed definitions and notation.

Definition 13. A parking spot s_i in the collection S of parking spots is a leaf-spot if there are no other parking spots in the subtree rooted at s_i . Let $L(T) = \{\ell_1, \dots, \ell_d\}$ denote the collection of leaf-spots. Let H be the maximum initial number of parking spots in T on the path from the root ρ to a leaf-spot in $L(T)$. For $\sigma \in [d]$, define $T_\sigma \subseteq V$ as the set of parking spots on the path from the root ρ to ℓ_σ , inclusive. We define T_σ to be alive if there is still an in-commission parking spot in T_σ , and dead otherwise. A T_σ is killed by r_t if r_t is the last parking spot to be decommissioned in T_σ . Let $\mathcal{A}_t = \{\sigma \in [d] \mid T_\sigma \text{ is alive just before the arrival of } r_t\}$. For a vertex $v \in V$, let $L(v)$ denote the collection of leaf-spots that are descendants of v in T . Let c_t be the location of the car just before the arrival of request r_t .

Algorithm TreeSearch: The algorithm has two phases: the prologue phase and the core phase. The algorithm starts in the prologue phase and transitions to the core phase after the first time m when there is no available parking space on the path from the new parking spot c_{m+1} to the root ρ , inclusive. The algorithm then remains in the core phase until the end. In the prologue phase, whenever the car is not parked at a vertex with an in-commission parking spot, the following actions are taken:

1. If there is an in-commission parking spot at c_t then no action is taken.
2. Else if there is an in-commission parking spot on the path between c_t and the root ρ , inclusive, then the car moves to the first in-commission parking spot on this path nearest to c_t .
3. Else the car moves to the root ρ and enters the core phase to determine where to go from there. So for analysis purposes, the movement to ρ counts as being part of the prologue phase, and the rest of the movement counts as being in the core phase.

If the car is at the root ρ and the algorithm is just transitioning into the core phase, then a live T_τ is picked uniformly at random from \mathcal{A}_{t+1} , an internal variable γ is set to be τ , and the car moves to the first in-commission parking spot on the path from ρ to ℓ_τ . Subsequently in the core phase, when a parking spot r_t is decommissioned then:

1. If the car is not parked at r_t , that is if $c_t \neq r_t$, then no action is taken.
2. Else the car moves to the first in-commission parking spot in T_τ with probability $q_t^\gamma(\tau)$ and sets γ to be τ . ($q_t^\gamma(\tau)$ is defined in the next subsection.)

Intuitively γ stores the last random choice of the algorithm.

3.2 The Definition of $q_t^\sigma(\tau)$

In this section we only consider times in the core phase. We conceptually divide up the tree T into three regions. Given vertex v and time t , we let z_t^v be the number of in-commission parking spots on the path from v to ρ , inclusive, just before decommission r_t . We then define the regions as follows:

1. The *root region* is the set of all vertices v such that $z_t^v = 0$. Note that this region is connected, and no decommissioning can occur in this region since there are no parking spots left.

2. The *frontier region* is the set of all vertices v such that $z_t^v = 1$. A decommissioning r_t is called a frontier decommissioning if r_t is in the frontier region.
3. The *outer region* is the set of all vertices v such that $z_t^v > 1$. A decommissioning r_t is called a outer decommissioning if r_t is in the outer region.

Observe that these regions have no dependence on random events internal to the algorithm. Further observe that step 2 of the core phase in algorithm **TreeSearch** maintains the invariant that the car is always parked at a spot in the frontier region. This means that any outer decommissionings will not move the car from its current parking spot.

Definition 14. Let r_m be the last decommissioning handled in the prologue phase of **TreeSearch**. Define $\mathcal{X}_t = \mathcal{A}_t \cap L(r_t)$ to be the collection of σ 's such that T_σ is alive and contains r_t and define $\mathcal{Y}_t = \mathcal{A}_t \setminus \mathcal{X}_t = \mathcal{A}_t \cap \overline{L(r_t)}$ to be the collection of σ 's such that T_σ is alive and doesn't contain r_t . Define $\mathcal{F}_t = \mathcal{X}_t \cap \overline{\mathcal{A}_{t+1}}$ to be the collection of σ 's such that T_σ is killed by r_t . Let n_t^σ denote the number of frontier decommissionings strictly before time t from T_σ . Define $w_t^\sigma = (1 - \epsilon)^{n_t^\sigma}$ for each $\sigma \in [d]$. Define $W_t(\mathcal{J}) = \sum_{\sigma \in \mathcal{J}} w_t^\sigma$ for any $\mathcal{J} \subseteq \{1, \dots, d\}$. Define π_t^σ as the probability the experts algorithm would give to expert σ , that is $\pi_t^\sigma = \frac{w_t^\sigma}{\sum_{\tau \in [d]} w_t^\tau}$. Define $\tilde{\pi}_t^\sigma$ as π_t^σ normalized amongst all experts in \mathcal{A}_t , that is $\tilde{\pi}_t^\sigma = \frac{w_t^\sigma}{\sum_{\tau \in \mathcal{A}_t} w_t^\tau}$ if $\sigma \in \mathcal{A}_t$, and 0 otherwise. Define p_t^σ as the probability that $\gamma = \sigma$ right before time t .

We are now ready to define $q_t^\sigma(\tau)$. Note that by the definition of **TreeSearch**, $q_t^\sigma(\tau)$ is only used for $\sigma \in \mathcal{X}_t$ since the algorithm only reaches step 2 of the core phase when $r_t \in T_\gamma$. We show in Lemma 16 that this definition of $q_t^\sigma(\tau)$ indeed defines a probability distribution over $\tau \in [d]$. We then show in Lemma 18 that the definition of $q_t^\sigma(\tau)$ guarantees that our desired invariant $p_t^\sigma = \tilde{\pi}_t^\sigma$ holds.

Definition 15.

$$q_t^\sigma(\tau) = \begin{cases} \frac{\epsilon w_t^\tau}{(1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)} & \text{if } \tau \in \mathcal{Y}_t \text{ and } \sigma \in \mathcal{X}_t \setminus \mathcal{F}_t \\ \frac{w_t^\tau}{(1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)} & \text{if } \tau \in \mathcal{Y}_t \text{ and } \sigma \in \mathcal{F}_t \\ \frac{1 - \sum_{\varsigma \in \mathcal{Y}_t} q_t^\sigma(\varsigma)}{|\mathcal{X}_t \setminus \mathcal{F}_t|} & \text{if } \tau \in \mathcal{X}_t \setminus \mathcal{F}_t \\ 0 & \text{if } \tau \in \mathcal{F}_t \text{ or } \tau \in \overline{\mathcal{A}_t} \end{cases}$$

Lemma 16. For all times t in the core phase and for all $\sigma \in \mathcal{X}_t$, $q_t^\sigma(\tau)$ forms a distribution over $\tau \in [d]$.

Proof. First, note that the cases of Definition 15 partition $[d]$ since

$$\begin{aligned} [d] &= \mathcal{A}_t \cup \overline{\mathcal{A}_t} && \text{by definition of } \mathcal{A}_t \\ &= (\mathcal{X}_t \cup \mathcal{Y}_t) \cup \overline{\mathcal{A}_t} && \text{by definition of } \mathcal{X}_t, \mathcal{Y}_t \\ &= ((\mathcal{X}_t \setminus \mathcal{F}_t) \cup \mathcal{F}_t \cup \mathcal{Y}_t) \cup \overline{\mathcal{A}_t} && \text{by definition of } \mathcal{F}_t \end{aligned}$$

Next, we show that the distribution sums to 1.

$$\begin{aligned}
\sum_{\tau \in [d]} q_t^\sigma(\tau) &= \sum_{\tau \in \mathcal{X}_t \setminus \mathcal{F}_t} q_t^\sigma(\tau) + \sum_{\tau \in \mathcal{Y}_t} q_t^\sigma(\tau) && \text{since } q_t^\sigma(\tau) = 0 \text{ for } \tau \in \mathcal{F}_t \text{ or } \tau \in \overline{\mathcal{A}}_t \\
&= \sum_{\tau \in \mathcal{X}_t \setminus \mathcal{F}_t} \frac{1 - \sum_{\varsigma \in \mathcal{Y}_t} q_t^\sigma(\varsigma)}{|\mathcal{X}_t \setminus \mathcal{F}_t|} + \sum_{\tau \in \mathcal{Y}_t} q_t^\sigma(\tau) && \text{By Defn. 15} \\
&= |\mathcal{X}_t \setminus \mathcal{F}_t| \frac{1 - \sum_{\varsigma \in \mathcal{Y}_t} q_t^\sigma(\varsigma)}{|\mathcal{X}_t \setminus \mathcal{F}_t|} + \sum_{\tau \in \mathcal{Y}_t} q_t^\sigma(\tau) \\
&= 1
\end{aligned}$$

Finally, we show that every event has non-negative probability. Trivially, $q_t^\sigma(\tau) \geq 0$ for any $\tau \in [d] \setminus (\mathcal{X}_t \setminus \mathcal{F}_t)$ by Definition 15. For $\tau \in \mathcal{X}_t \setminus \mathcal{F}_t$, $q_t^\sigma(\tau) = \frac{1 - \sum_{\varsigma \in \mathcal{Y}_t} q_t^\sigma(\varsigma)}{|\mathcal{X}_t \setminus \mathcal{F}_t|}$, so it is sufficient to show that $\sum_{\tau \in \mathcal{Y}_t} q_t^\sigma(\tau) \leq 1$. If $\sigma \in \mathcal{X}_t \setminus \mathcal{F}_t$, then

$$\begin{aligned}
\sum_{\tau \in \mathcal{Y}_t} q_t^\sigma(\tau) &= \sum_{\tau \in \mathcal{Y}_t} \frac{\epsilon w_t^\tau}{(1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)} && \text{by Defn. 15} \\
&\leq \frac{\epsilon}{W_t(\mathcal{Y}_t)} \sum_{\tau \in \mathcal{Y}_t} w_t^\tau && \text{since } (1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) \geq 0 \\
&= \frac{\epsilon W_t(\mathcal{Y}_t)}{W_t(\mathcal{Y}_t)} && \text{by definition of } \mathcal{Y}_t \\
&= \epsilon \\
&\leq 1
\end{aligned}$$

If $\sigma \in \mathcal{F}_t$, the only difference from the previous case is that the constant ϵ is replaced with a 1 in Definition 15, so the analysis still holds. \square

Before proving that the invariant $p_t^\sigma = \tilde{\pi}_t^\sigma$ holds using distribution q , we prove some necessary properties about the set X_t .

Lemma 17. *Let $X \subseteq \mathcal{X}_t$ be an arbitrary subset of \mathcal{X}_t and let $\sigma \in \mathcal{X}_t$ be any parking spot, we then have that:*

1. $w_t^\sigma = \frac{W_t(X)}{|X|}$
2. $\sum_{\tau \in X} \tilde{\pi}_t^\tau = \frac{W_t(X)}{W_t(\mathcal{A}_t)}$

Proof. Given decommissioning r_t , let $\sigma_i, \sigma_j \in \mathcal{X}_t$ be arbitrary spots. Then, by definition of \mathcal{X}_t , T_{σ_i} and T_{σ_j} overlap at r_t as well as every parking spot from r_t to the root. So, any frontier decommissioning before r_t that decommissioned a spot from T_{σ_i} also decommissioned a spot from T_{σ_j} . In other words, $n_t^{\sigma_i} = n_t^{\sigma_j}$ and $w_t^{\sigma_i} = w_t^{\sigma_j}$. Since σ_i and σ_j were arbitrary from \mathcal{X}_t , for any subset $X \subseteq \mathcal{X}_t$ and $\sigma \in \mathcal{X}_t$,

$$W_t(X) = \sum_{\tau \in X} w_t^\tau = |X|w_t^\sigma$$

Property 2 comes just from the definition of $\tilde{\pi}_t^\sigma$:

$$\begin{aligned} \sum_{\tau \in X} \tilde{\pi}_t^\tau &= \frac{\sum_{\tau \in X} w_t^\tau}{\sum_{\tau \in \mathcal{A}_t} w_t^\tau} && \text{by definition of } \tilde{\pi}_t^\tau \\ &= \frac{W_t(X)}{W_t(\mathcal{A}_t)} && \text{by definition of } W_t(X), W_t(\mathcal{A}_t) \end{aligned}$$

□

Lemma 18. *For all times t during the core phase and for all $\sigma \in [d]$, $p_t^\sigma = \tilde{\pi}_t^\sigma$.*

Proof. The proof is by induction on the time t . The transition from the prologue to core phases of **TreeSearch** guarantees that this invariant is initially true. Now assuming $p_t^\sigma = \tilde{\pi}_t^\sigma$ for all σ , we want to show that $p_{t+1}^\sigma = \tilde{\pi}_{t+1}^\sigma$ for all σ . The proof is broken into cases.

In the first case, assume $\sigma \in \mathcal{X}_t \setminus \mathcal{F}_t$ and r_t is a frontier decommissioning. We will begin with two equations:

$$\begin{aligned} \sum_{\tau \in \mathcal{X}_t \setminus \mathcal{F}_t} p_t^\tau \frac{1 - \sum_{\varsigma \in \mathcal{Y}_t} q_t^\tau(\varsigma)}{|\mathcal{X}_t \setminus \mathcal{F}_t|} &= \sum_{\tau \in \mathcal{X}_t \setminus \mathcal{F}_t} p_t^\tau \frac{1 - \sum_{\varsigma \in \mathcal{Y}_t} \frac{\epsilon w_t^\varsigma}{(1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)}}{|\mathcal{X}_t \setminus \mathcal{F}_t|} && \text{by Defn. 15} \\ &= \sum_{\tau \in \mathcal{X}_t \setminus \mathcal{F}_t} p_t^\tau \frac{1 - \frac{\epsilon W_t(\mathcal{Y}_t)}{(1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)}}{|\mathcal{X}_t \setminus \mathcal{F}_t|} && \text{by defn. of } W_t(\mathcal{Y}_t) \\ &= \sum_{\tau \in \mathcal{X}_t \setminus \mathcal{F}_t} p_t^\tau \frac{(1-\epsilon)W_t(\mathcal{A}_t \setminus \mathcal{F}_t)}{|\mathcal{X}_t \setminus \mathcal{F}_t|((1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t))} && \text{by defn. of } \mathcal{A}_t, \mathcal{F}_t \\ &= \frac{(1-\epsilon)W_t(\mathcal{A}_t \setminus \mathcal{F}_t)}{|\mathcal{X}_t \setminus \mathcal{F}_t|((1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t))} \sum_{\tau \in \mathcal{X}_t \setminus \mathcal{F}_t} \tilde{\pi}_t^\tau && p_t^\tau = \tilde{\pi}_t^\tau \\ &= \frac{(1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t)W_t(\mathcal{A}_t \setminus \mathcal{F}_t)}{|\mathcal{X}_t \setminus \mathcal{F}_t|W_t(\mathcal{A}_t)((1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t))} && \text{by Lemma 17 part 2} \end{aligned} \tag{8}$$

and similarly,

$$\begin{aligned} \sum_{\tau \in \mathcal{F}_t} p_t^\tau \frac{1 - \sum_{\varsigma \in \mathcal{Y}_t} q_t^\tau(\varsigma)}{|\mathcal{X}_t \setminus \mathcal{F}_t|} &= \sum_{\tau \in \mathcal{F}_t} p_t^\tau \frac{1 - \sum_{\varsigma \in \mathcal{Y}_t} \frac{w_t^\varsigma}{(1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)}}{|\mathcal{X}_t \setminus \mathcal{F}_t|} && \text{by Defn. 15} \\ &= \sum_{\tau \in \mathcal{F}_t} p_t^\tau \frac{1 - \frac{W_t(\mathcal{Y}_t)}{(1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)}}{|\mathcal{X}_t \setminus \mathcal{F}_t|} && \text{by defn. of } W_t(\mathcal{Y}_t) \\ &= \sum_{\tau \in \mathcal{F}_t} p_t^\tau \frac{(1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t)}{|\mathcal{X}_t \setminus \mathcal{F}_t|((1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t))} \\ &= \frac{(1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t)}{|\mathcal{X}_t \setminus \mathcal{F}_t|((1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t))} \sum_{\tau \in \mathcal{F}_t} \tilde{\pi}_t^\tau && p_t^\tau = \tilde{\pi}_t^\tau \\ &= \frac{(1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t)W_t(\mathcal{F}_t)}{|\mathcal{X}_t \setminus \mathcal{F}_t|W_t(\mathcal{A}_t)((1-\epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t))} && \text{by Lemma 17 part 2} \end{aligned} \tag{9}$$

Now, notice that by the definition of $q_t^\sigma(\tau)$, if $\gamma = \sigma \in \mathcal{X}_t$ after r_t , then it must have been the case that $\gamma \in \mathcal{X}_t$ before r_t , and thus:

$$\begin{aligned}
p_{t+1}^\sigma &= \sum_{\tau \in \mathcal{X}_t} p_t^\tau q_t^\tau(\sigma) \\
&= \sum_{\tau \in \mathcal{X}_t} p_t^\tau \frac{1 - \sum_{\varsigma \in \mathcal{Y}_t} q_t^\tau(\varsigma)}{|\mathcal{X}_t \setminus \mathcal{F}_t|} && \text{by Defn. 15} \\
&= \sum_{\tau \in \mathcal{X}_t \setminus \mathcal{F}_t} p_t^\tau \frac{1 - \sum_{\varsigma \in \mathcal{Y}_t} q_t^\tau(\varsigma)}{|\mathcal{X}_t \setminus \mathcal{F}_t|} + \sum_{\tau \in \mathcal{F}_t} p_t^\tau \frac{1 - \sum_{\varsigma \in \mathcal{Y}_t} q_t^\tau(\varsigma)}{|\mathcal{X}_t \setminus \mathcal{F}_t|} \\
&= \frac{(1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t)(W_t(\mathcal{A}_t \setminus \mathcal{F}_t) + W_t(\mathcal{F}_t))}{|\mathcal{X}_t \setminus \mathcal{F}_t|W_t(\mathcal{A}_t)((1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t))} && \text{by Eq. 8 and 9} \\
&= \frac{(1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t)}{|\mathcal{X}_t \setminus \mathcal{F}_t|((1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t))} && \text{by definition of } W(\mathcal{A}_t) \\
&= \frac{(1 - \epsilon)w_t^\sigma}{(1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)} && \text{by Lemma 17 part 1} \\
&= \tilde{\pi}_{t+1}^\sigma && \text{by definition of } \tilde{\pi}_{t+1}^\sigma
\end{aligned}$$

In the second case, assume that $\sigma \in \mathcal{Y}_t$ and r_t is a frontier decommissioning. We will begin with three equations:

$$\begin{aligned}
\tilde{\pi}_{t+1}^\sigma - \tilde{\pi}_t^\sigma &= \frac{w_t^\sigma}{(1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)} - \frac{w_t^\sigma}{W_t(\mathcal{X}_t) + W_t(\mathcal{Y}_t)} && \text{by defn. of } \pi_t^\sigma \\
&= \frac{(W_t(\mathcal{F}_t) + \epsilon W_t(\mathcal{X}_t \setminus \mathcal{F}_t))w_t^\sigma}{(W_t(\mathcal{X}_t) + W_t(\mathcal{Y}_t))((1 - \epsilon)W_t(\mathcal{X}_t) + W_t(\mathcal{Y}_t))} && (10) \\
&= \tilde{\pi}_t^\sigma \frac{W_t(\mathcal{F}_t) + \epsilon W_t(\mathcal{X}_t \setminus \mathcal{F}_t)}{(1 - \epsilon)W_t(\mathcal{X}_t) + W_t(\mathcal{Y}_t)} && \text{by defn. of } \tilde{\pi}_t^\sigma \\
&= p_t^\sigma \frac{W_t(\mathcal{F}_t) + \epsilon W_t(\mathcal{X}_t \setminus \mathcal{F}_t)}{(1 - \epsilon)W_t(\mathcal{X}_t) + W_t(\mathcal{Y}_t)} && \text{as } p_t^\sigma = \tilde{\pi}_t^\sigma.
\end{aligned}$$

$$\begin{aligned}
\sum_{\tau \in \mathcal{X}_t \setminus \mathcal{F}_t} \tilde{\pi}_t^\tau q_t^\tau(\sigma) &= \sum_{\tau \in \mathcal{X}_t \setminus \mathcal{F}_t} \tilde{\pi}_t^\tau \frac{\epsilon w_t^\sigma}{(1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)} && \text{by Defn. 15} \\
&= \frac{\epsilon w_t^\sigma}{(1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)} \sum_{\tau \in \mathcal{X}_t \setminus \mathcal{F}_t} \tilde{\pi}_t^\tau \\
&= \frac{\epsilon w_t^\sigma}{(1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)} \left(\frac{W_t(\mathcal{X}_t \setminus \mathcal{F}_t)}{W_t(\mathcal{A}_t)} \right) && \text{by Lemma 17 part 2} \\
&= \tilde{\pi}_t^\sigma \frac{\epsilon W_t(\mathcal{X}_t \setminus \mathcal{F}_t)}{(1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)} && \text{by defn. of } \tilde{\pi}_t^\sigma \\
&= p_t^\sigma \frac{\epsilon W_t(\mathcal{X}_t \setminus \mathcal{F}_t)}{(1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)} && \text{as } p_t^\sigma = \tilde{\pi}_t^\sigma
\end{aligned} \tag{11}$$

and similarly,

$$\begin{aligned}
\sum_{\tau \in \mathcal{F}_t} \tilde{\pi}_t^\tau q_t^\tau(\sigma) &= \sum_{\tau \in \mathcal{F}_t} \tilde{\pi}_t^\tau \frac{w_t^\sigma}{(1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)} && \text{by Defn. 15} \\
&= p_t^\sigma \frac{W_t(\mathcal{F}_t)}{(1 - \epsilon)W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)} && (12)
\end{aligned}$$

If right after r_t it was the case that $\gamma = \sigma \in \mathcal{Y}_t$, then it must be the case that right before r_t either $\gamma = \sigma$ and the car did not move at time t , or the car moved and γ was updated. This gives us

$$p_{t+1}^\sigma = p_t^\sigma + \sum_{\tau \in \mathcal{X}_t} p_t^\tau q_t^\tau(\sigma)$$

or equivalently:

$$\begin{aligned} p_{t+1}^\sigma - p_t^\sigma &= \sum_{\tau \in \mathcal{X}_t} p_t^\tau q_t^\tau(\sigma) \\ &= \sum_{\tau \in \mathcal{X}_t} \tilde{\pi}_t^\tau q_t^\tau(\sigma) & p_t^\tau &= \tilde{\pi}_t^\tau \\ &= \sum_{\tau \in \mathcal{X}_t \setminus \mathcal{F}_t} \tilde{\pi}_t^\tau q_t^\tau(\sigma) + \sum_{\tau \in \mathcal{F}_t} \tilde{\pi}_t^\tau q_t^\tau(\sigma) \\ &= p_t^\sigma \frac{(\epsilon W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{F}_t))}{(1 - \epsilon) W_t(\mathcal{X}_t \setminus \mathcal{F}_t) + W_t(\mathcal{Y}_t)} & \text{by Eq. 11 and 12} \\ &= \tilde{\pi}_{t+1}^\sigma - \tilde{\pi}_t^\sigma & \text{by Eq. 10} \end{aligned}$$

Hence we can conclude that $p_{t+1}^\sigma = \tilde{\pi}_{t+1}^\sigma$.

In the third case, assume $\sigma \in \mathcal{F}_t$ and r_t is a frontier decommissioning. By definition of \mathcal{F}_t , T_σ has been killed by r_t and so $p_{t+1}^\sigma = 0$. Similarly, since $\sigma \notin \mathcal{A}_{t+1}$, $\tilde{\pi}_{t+1}^\sigma = 0 = p_{t+1}^\sigma$.

In the last case assume that r_t is an outer decommissioning. Then r_t does not increase n_t^σ for any $\sigma \in [d]$, so $\tilde{\pi}_t^\sigma = \tilde{\pi}_{t+1}^\sigma$. Moreover, by the invariant that we also keep the car at a parking spot in the frontier, r_t cannot decommission the parking spot at c_t , and thus the car does not move. Thus $p_{t+1}^\sigma = p_t^\sigma$. \square

3.3 Cost Analysis

Definition 15 with Lemmas 16 and 18 give us the following bound on the cost:

Theorem 19. *During the prologue phase, $\sum_{t=1}^m \mathbf{1}^{\text{TM}}(t) \leq H$ and during the core phase,*

$$\mathbf{E} \left[\sum_{t=m+1}^{k-1} \mathbf{1}^{\text{TM}}(t) \right] \leq (1 + \epsilon)H + \frac{\ln d}{\epsilon}$$

where $\mathbf{1}^{\text{TM}}(t)$ is an indicator random variable that is 1 if **TreeSearch** moves the car to a new parking spot on the decommissioning r_t and 0 otherwise.

Proof. The first inequality for the prologue phase follows from the fact that the car always moves towards the root and the definition of H . So now consider a time t in the core phase, and $\sigma \in [d]$. Let $\mathcal{D}_t = [d] \setminus \mathcal{A}_t$ be the set of dead paths just before decommissioning r_t and define

$$c_t^\sigma = \begin{cases} 1 & \sigma \in \mathcal{X}_t \text{ or } \sigma \in \mathcal{D}_t \\ 0 & \text{otherwise.} \end{cases}$$

Let δ_t^σ be defined for all $\sigma \in \mathcal{A}_t$ such that $\tilde{\pi}_t^\sigma = \pi_t^\sigma + \delta_t^\sigma$. Then by summing over all $\sigma \in \mathcal{A}_t$ we have

$$\sum_{\sigma \in \mathcal{A}_t} \tilde{\pi}_t^\sigma = \sum_{\sigma \in \mathcal{A}_t} \pi_t^\sigma + \sum_{\sigma \in \mathcal{A}_t} \delta_t^\sigma. \quad (13)$$

Note that $\sum_{\sigma \in \mathcal{A}_t} \tilde{\pi}_t^\sigma = 1$ by definition of $\tilde{\pi}_t$ and $\sum_{\sigma \in \mathcal{A}_t} \pi_t^\sigma = 1 - \sum_{\tau \in \mathcal{D}_t} \pi_t^\tau$ since \mathcal{A}_t and \mathcal{D}_t partition $[d]$. Thus Equation 13 can be reformulated as

$$\sum_{\sigma \in \mathcal{A}_t} \delta_t^\sigma = \sum_{\tau \in \mathcal{D}_t} \pi_t^\tau. \quad (14)$$

Thus

$$\begin{aligned} \mathbf{E} \left[\mathbf{1}^{\text{TM}}(i) \right] &= \sum_{\sigma \in \mathcal{X}_t} p_t^\sigma && \text{By Defn. of } p_t^\sigma \\ &= \sum_{\sigma \in \mathcal{X}_t} \tilde{\pi}_t^\sigma && \text{by Lemma 18} \\ &= \sum_{\sigma \in \mathcal{X}_t} \delta_t^\sigma + \sum_{\sigma \in \mathcal{X}_t} \pi_t^\sigma && \text{By Defn. of } \delta_t^\sigma \\ &\leq \sum_{\sigma \in \mathcal{A}_t} \delta_t^\sigma + \sum_{\sigma \in \mathcal{X}_t} \pi_t^\sigma && \text{Since } \mathcal{X}_t \subseteq \mathcal{A}_t \\ &= \sum_{\sigma \in \mathcal{D}_t} \pi_t^\sigma + \sum_{\sigma \in \mathcal{X}_t} \pi_t^\sigma && \text{by Eq. 14} \\ &= \sum_{\sigma} c_t^\sigma \pi_t^\sigma && \text{By Defn. of } c_t^\sigma \\ &\leq (1 + \epsilon) \min_{\sigma \in [d]} \left(\sum_{i=m+1}^{k-1} c_i^\sigma \right) + \frac{\ln d}{\epsilon} && \text{By Multiplicative Weights Analysis [6]} \\ &\leq (1 + \epsilon)H + \frac{\ln(d)}{\epsilon} && \text{By considering the last alive path } T_\sigma \end{aligned}$$

For the last inequality, note that if T_σ is the last alive path then $\sum_{i=m+1}^{k-1} c_i^\sigma \leq H$ by the definition of H . \square

3.4 Monotonicity

We show that any neighbor algorithm for online metrical search can be extended to a monotone algorithm for online metrical matching, where a neighbor algorithm has the property that if it moves the car to a parking spot s_i with positive probability then it must be the case that there is no in-commission parking spot on the route to s_i . As **TreeSearch** is obviously a neighbor algorithm, it then follows that it can be extended to a monotone algorithm for online metrical matching, which we will call **TreeMatch**.

Lemma 20. *Let A be a neighbor algorithm for online metrical search. Then there exists a monotone algorithm B for online metrical matching on a tree metric that is identical to A for online metrical search instances.*

Proof. We construct B using A as a subroutine. As long as B continues to receive requests for which there is a co-located available server, it will send that request to a co-located server. By renumber the requests, let r_1 for the first request for which there is not a co-located server. We now know that at this point, there optimal matching has positive cost. Now B starts running A , with the current server locations and the car parked at r_1 . As long as there is an optimal matching with only one positive edge, B will continue to run A . Whenever A doesn't move the car, the request

arrived at a co-located server, which is where B will send the request. If a request arrives at the location of the car in A , then B moves the request to where A moves the car.

If a request r_t arrives at a vertex without a co-located server and that is not the location of the car, then let us call it a deviating request. B handles a deviating request in the following way. Let c be the current parking spot of the car. Let Q be the collection of in-commission parking spots that can be reached from c without passing over another in-commission parking spot. For a parking spot $s_i \in Q$ let p_i be the probability that A moves the car to s_i if c is decommissioned. Let R be the vertices that can reach c without passing through a vertex in Q . Note R includes no vertices in Q . Let $X = V - R - Q$, the vertices separated from c by Q . If a request arrives at vertex $v \in X$ then the B moves the request to the first in-commission parking spot on the path from v to c . If a request arrives at a vertex $v \in Q$ then B moves the request to a co-located server. If a request arrives at a vertex $v \in R$ then moves the request to each server $s_i \in Q$ with probability p_i . It is clear that B is monotone for a deviating request r_t . After receiving a deviating request, then B knows that the instance is not an online metrical searching instance, and can then mimic any monotone algorithm, for exactly the greedy algorithm that moves a request to the nearest available server. \square

4 The GroveMatch Algorithm

In subsection 4.1 we describe an algorithm **GroveBuild** that builds a grove G from a tree metric T with distance metric d_T before any request arrives. We assume without loss of generality that the minimum distance in T is 1. In subsection 4.2 we then give an algorithm **GroveMatch** for online metrical matching on a tree metric that utilizes the algorithm **TreeMatch** on each tree in the grove constructed by **GroveBuild**, and we prove some basic properties of the grove G . In subsection 4.3 we show that **GroveMatch** is a monotone online metrical matching algorithm on a tree metric, and is $O(\log^6 \Delta \log^2 n)$ -competitive for online metrical search instances.

4.1 The GroveBuild Algorithm

Definition 21. *A grove G is either: a rooted tree X consisting of a single vertex, or an unweighted rooted tree X with a grove $X(v)$ associated with each vertex $v \in X$. The tree X is the canopy of the grove G . Each $X(v)$ is a subgrove of X . The canopy of a subtree $X(v)$ is a child of X . Trees in G are descendants of X .*

GroveBuild Description: **GroveBuild** is a recursive algorithm that takes as input a tree metric T , a designated root ρ of T , positive real R , a positive real α and a positive integer d . In the initial call to **GroveBuild**, T is the original tree metric, ρ is an arbitrary vertex in T , R is the maximum distance Δ between ρ and any other vertex in T , d is 1, and α is a parameter to be determined later in the analysis.

If T consists of a single vertex v , then the recursion ends and the algorithm outputs a rooted tree consisting of only the vertex v . We call this tree a leaf of the grove. Otherwise the algorithm's first goal is to partition the vertices of T into parts P_1, \dots, P_k , and designate one vertex ℓ_i of each partition P_i as being the leader of P_i . To accomplish this, the algorithm sets partition P_1 to consist of the vertices in T that are within a distance z of ρ , where z is selected uniformly at random from the range $[0, \frac{R}{\alpha}]$. The leader ℓ_1 is set to be ρ . To compute P_i and ℓ_i after the first $i - 1$ parts and leaders are computed the algorithm takes the following steps. Let ℓ_i be a vertex such that $\ell_i \notin \cup_{j=1}^{i-1} P_j$ and for each vertex v on the path (ℓ_i, ρ) it is the case that $v \in \cup_{j=1}^{i-1} P_j$. So ℓ_i is not in but adjacent to the previous partitions. Then P_i consists of all vertices $v \in T - \cup_{j=1}^{i-1} P_j$ that are

within distance $\frac{R}{\alpha}$ from ℓ_i in T . So P_i intuitively is composed of vertices that are not in previous partitions and that are close to ℓ_i .

The tree X at this point in the recursion has a vertex for each part in the partition of T . There is an edge between vertices/parts P_i and P_j in X if and only if there is an edge (v, w) in T such that $v \in P_i$ and $w \in P_j$. We identify this edge in X with the edge $(v, w) \in T$. The root of X is the vertex/part P_1 . The tree X is at depth d in the grove. The grove $X(P_i)$ associated with vertex P_i in X is the result of calling **GroveBuild** on the subtree of T induced by the vertices in P_i , with ℓ_i designated as the root, parameter R decreased by an α factor, parameter α unchanged, and parameter d incremented by 1.

So from here on, let G denote the grove built by **GroveBuild** on the original tree metric T .

Definition 22.

- For an edge $(u, v) \in T$, let $\delta(u, v)$ be the depth in the grove G of the tree X that contains (u, v) . Note that each edge in T occurs in exactly one tree in G .
- For an edge $(u, v) \in T$, define $d_G(u, v)$ to be $\frac{\Delta}{\alpha^{\delta(u, v)-1}}$.
- For vertices $u_0, u_h \in T$, connected by the simple path (u_0, u_1, \dots, u_h) in T , define $d_G(u_0, u_h)$ to be $\sum_{i=0}^{h-1} d_G(u_i, u_{i+1})$. Obviously d_G forms a metric on the vertices of T .

Lemma 23. Recall that $d_T(u, v)$ is the shortest path distance between two vertices u, v of tree T . For all vertices $u, v \in T$, we have that $d_G(u, v) \geq d_T(u, v)$ and $\mathbf{E}[d_G(u, v)] \leq \alpha(1 + \log \Delta) \cdot d_T(u, v)$.

Proof. Let (u_0, u_1, \dots, u_h) be the path from u_0 to u_h in T . Since $d_G(u_0, u_h) = \sum_{i=0}^{h-1} d_G(u_i, u_{i+1})$, and $d_T(u_0, u_h) = \sum_{i=0}^{h-1} d_T(u_i, u_{i+1})$, it is sufficient to prove this for each $(u, v) \in T$. For notational simplicity let $\delta := \delta(u, v)$.

Notice that from the construction of G that $d_G(u, v) = \frac{\Delta}{\alpha^{\delta-1}}$. Showing that $d_T(u, v) \leq \frac{\Delta}{\alpha^{\delta-1}}$ proves the first inequality: if $\delta = 1$ then $d_T(u, v) \leq \Delta$, else, $d_T(u, v) \leq \frac{\Delta}{\alpha^{\delta-1}}$ since (u, v) did not get cut at depth $\delta - 1$.

To prove the second inequality, let A_i be the event that $\delta = i$, and $A_{<i}$ be the event that $\delta < i$. Notice that if $d_T(u, v) \geq \frac{\Delta}{\alpha^i}$ for some i , then, $\delta \leq i$ since cuts at depth i of the recursion are made in increments of $\frac{\Delta}{\alpha^i}$ distance. Hence, the value of δ is at most $1 + \log_\alpha(\frac{\Delta}{d_T(u, v)})$. By the linearity of expectation we have:

$$\begin{aligned} \mathbf{E}[d_G(u, v)] &= \sum_{i=1}^{1+\log_\alpha(\Delta/d_T(u, v))} \Pr[A_i] \cdot \frac{\Delta}{\alpha^{i-1}} = \sum_{i=1}^{1+\log_\alpha(\Delta/d_T(u, v))} \Pr[A_i | \overline{A_{<i}}] \Pr[\overline{A_{<i}}] \cdot \frac{\Delta}{\alpha^{i-1}} \\ &\leq \sum_{i=1}^{1+\log_\alpha(\Delta/d_T(u, v))} \Pr[A_i | \overline{A_{<i}}] \cdot \frac{\Delta}{\alpha^{i-1}} = \sum_{i=1}^{1+\log_\alpha(\Delta/d_T(u, v))} \alpha d_T(u, v). \end{aligned}$$

The last equality follows from the fact that $\Pr[A_i | \overline{A_{<i}}] = \frac{d_T(u, v)}{\Delta/\alpha^i}$ since cuts at depth i of the recursion are made in increments of $\frac{\Delta}{\alpha^i}$ distance with an offset randomly chosen from $[0, \frac{\Delta}{\alpha^i}]$. \square

Corollary 24. An algorithm \mathcal{B} that is c -competitive for online metric matching on T with distance metric d_G is $O(c \cdot \alpha \log \Delta)$ -competitive for online metric matching on T with distance metric d_T .

Proof. This is an immediate consequence of Lemma 23. \square

4.2 GroveMatch Description

We now describe an algorithm **GroveMatch** for online metrical matching for tree metrics.

GroveMatch Description: Conceptually within **GroveMatch**, a separate copy **TreeMatch**(X) of the online metric matching algorithm **TreeMatch** will be run on each tree X in the grove G constructed by the algorithm **GroveBuild**. In order to accomplish this, we need to initially place servers at the vertices in X . We set the number of servers initially located at each vertex $x \in X$ to the number of servers in T that are located at vertices $v \in T$ such that $v \in x$ (recall that each vertex in a tree in the grove G corresponds to a collection of vertices in T).

When a request r_t arrives at a vertex v in T , the algorithm **GroveMatch** calls the algorithm **TreeMatch** on a sequence $(X_1, x_1), (X_2, x_2), \dots$ where each X_i is a tree of depth i in G and x_i is a vertex in X_i . Initially X_1 is the depth 1 tree in G , and x_1 is the vertex in X_1 that contains v . Assume that **TreeMatch** has already been called on $(X_1, x_1), (X_2, x_2), \dots, (X_{i-1}, x_{i-1})$, then the algorithm **GroveMatch** processes (X_i, x_i) in the following manner. First, **TreeMatch**(X_i) is called to respond to a request at x_i . Let y_i be the vertex in X_i that **TreeMatch**(X_i) moved this request to. If X_i is a leaf in G , then **TreeMatch**(X_i) sets $y_i = x_i$, and **GroveMatch** moves request r_t to the unique vertex in T corresponding to x_i . If X_i is not a leaf in G , then X_{i+1} is set to be the canopy of the grove $X_i(y_i)$, and $x_{i+1} = \arg \min_{w \in T: w \in X_{i+1}} d_T(v, w)$ or equivalently x_{i+1} is the first vertex in X_{i+1} that one encounters if one walks in T from v to the vertices of X_{i+1} .

Lemma 25. *Consider a tree X at depth δ with root ρ in grove G . For any vertex v in X , the number of hops in X between ρ and v is at most $\alpha + 1$. Furthermore, by the time that **TreeMatch**(X) enters its core phase, it must be the case that for every descendent tree Y of X in G there will be no future movement of the car on edges in Y while **TreeMatch**(Y) is in its prologue phase.*

Proof. This follows immediately from the fact that the parameter R decreases by an α factor on each recursion. \square

4.3 GroveMatch Analysis

We now analyze **GroveBuild** and **GroveMatch** under the assumption that $\alpha = (\ln n)(\log_\alpha^2 \Delta)$ and $\epsilon = \frac{1}{\log_\alpha \Delta}$.

Lemma 26. *The algorithm **GroveMatch** is $O(\log n \log^3 \Delta)$ -competitive for online metrical search instances with the metric d_G .*

Proof. If **GroveMatch** directs a request to traverse an edge $(u, v) \in T$, we will say that the cost of this traversal is charged to the unique tree in G that contains (u, v) . Define $P(\delta)$ to be the charge incurred by a tree X of depth δ in G and all subgroves $X(v)$ of X during the prologue phase of **TreeMatch**(X). Define $C(\delta)$ to be the charge incurred by a tree X of depth δ in G and all subgroves $X(v)$ of X during the core phase of **TreeMatch**(X).

Recall that the distance under the d_G metric of every edge in X is $\frac{\Delta}{\alpha^{\delta-1}}$ and by Lemma 25 there are at most $\alpha + 1$ vertices on the path from any leaf to the root of X . This gives us that the distance in X under d_G from the root to any leaf is at most $\alpha \frac{\Delta}{\alpha^{\delta-1}} = \frac{\Delta}{\alpha^{\delta-2}}$ and that the diameter of X is at most $2 \frac{\Delta}{\alpha^{\delta-2}}$. The only subgroves $X(v)$ of X that incur costs during the prologue phase of **TreeMatch**(X) are those subgroves for which v is traversed by the car on its path to the root of X . Thus we obtain the following recurrence:

$$P(\delta) \leq (\alpha + 1)(P(\delta + 1) + C(\delta + 1)) + \frac{\Delta}{\alpha^{\delta-2}}. \quad (15)$$

Note that once the core phase begins in $\text{TreeMatch}(X)$, by Lemma 25 all instances of $\text{TreeMatch}(Y)$ on any tree Y that is a descendent of X in G can incur no most costs in their prologue phase. By Theorem 19 the core phase cost on X is at most $(1+\epsilon)(\alpha+1) + \frac{\ln n}{\epsilon}$ times the diameter of X , which is at most $2\frac{\Delta}{\alpha^{\delta-2}}$. Thus we obtain the following recurrence:

$$C(\delta) \leq \left(C(\delta+1) + 2\frac{\Delta}{\alpha^{\delta-2}} \right) \left((1+\epsilon)(\alpha+1) + \frac{\ln n}{\epsilon} \right) \quad (16)$$

We expand the recurrence relation for $C(\delta)$ first. Treating $((1+\epsilon)(\alpha+1) + \frac{\ln(n)}{\epsilon})$ as a constant Z , and expanding $C(\delta)$ we obtain:

$$\begin{aligned} C(\delta) &\leq \left(C(\delta+1) + 2\frac{\Delta}{\alpha^{\delta-2}} \right) Z \\ &= C(\delta+1)Z + 2\frac{\Delta}{\alpha^{\delta-2}}Z \\ &\leq 2\frac{\Delta}{\alpha^{\delta-1}} \sum_{i=1}^{\log_{\alpha} \Delta} \left(\frac{Z}{\alpha} \right)^i \\ &\leq \frac{2\Delta \log_{\alpha} \Delta}{\alpha^{\delta-1}} \left(\frac{Z}{\alpha} \right)^{\log_{\alpha}(\Delta)} \\ &\leq \frac{2\Delta \log_{\alpha} \Delta}{\alpha^{\delta-1}} \left((1+\epsilon)\frac{\alpha+1}{\alpha} + \frac{\ln(n)}{\epsilon\alpha} \right)^{\log_{\alpha}(\Delta)} \\ &\leq \frac{2\Delta \log_{\alpha} \Delta}{\alpha^{\delta-1}} \left(1+\epsilon + \frac{1}{\alpha} + \frac{\epsilon}{\alpha} + \frac{\ln n}{\epsilon\alpha} \right)^{\log_{\alpha}(\Delta)} \\ &\leq \frac{2\Delta \log_{\alpha} \Delta}{\alpha^{\delta-1}} \left(1 + \frac{1}{\log_{\alpha} \Delta} + \frac{1}{(\ln n)(\log_{\alpha}^2 \Delta)} + \frac{1}{(\ln n)(\log_{\alpha}^3 \Delta)} + \frac{1}{\log_{\alpha} \Delta} \right)^{\log_{\alpha}(\Delta)} \\ &\leq \frac{2\Delta \log_{\alpha} \Delta}{\alpha^{\delta-1}} \left(1 + \frac{4}{\log_{\alpha} \Delta} \right)^{\log_{\alpha} \Delta} \\ &\leq \frac{2e^4 \Delta \log_{\alpha} \Delta}{\alpha^{\delta-1}} \end{aligned}$$

Now expanding the recurrence relation for $P(\delta)$ we obtain:

$$\begin{aligned} P(\delta) &\leq (\alpha+1) (P(\delta+1) + C(\delta+1)) + \frac{\Delta}{\alpha^{\delta-2}} \\ &\leq (\alpha+1) \left(P(\delta+1) + \frac{2e^4 \Delta}{\alpha^{\delta-1}} \log_{\alpha} \Delta \right) + \frac{\Delta}{\alpha^{\delta-2}} \\ &= (\alpha+1)P(\delta+1) + \frac{2e^4 \Delta(\alpha+1)}{\alpha^{\delta-1}} \log_{\alpha} \Delta + \frac{\Delta}{\alpha^{\delta-2}} \\ &\leq (\alpha+1)P(\delta+1) + \frac{2e^4 \Delta(\alpha+1)}{\alpha^{\delta-1}} \log_{\alpha} \Delta + \frac{\Delta(\alpha+1)}{\alpha^{\delta-1}} \\ &\leq (\alpha+1)P(\delta+1) + \frac{\Delta(\alpha+1)}{\alpha^{\delta-1}} (2e^4 \log_{\alpha} \Delta + 1) \\ &\leq (\alpha+1)P(\delta+1) + \frac{3e^4(\alpha+1)\Delta \log_{\alpha} \Delta}{\alpha^{\delta-1}} \\ &\leq \frac{3e^4 \Delta \log_{\alpha} \Delta}{\alpha^{\delta-2}} \sum_{i=1}^{\log_{\alpha} \Delta} \left(\frac{\alpha+1}{\alpha} \right)^i \end{aligned}$$

$$\begin{aligned}
&\leq \frac{3e^4 \Delta \log_\alpha^2 \Delta}{\alpha^{\delta-2}} \left(\frac{\alpha+1}{\alpha} \right)^{\log_\alpha \Delta} \\
&= \frac{3e^4 \Delta \log_\alpha^2 \Delta}{\alpha^{\delta-2}} \left(1 + \frac{1}{\alpha} \right)^{\log_\alpha \Delta} \\
&= \frac{3e^4 \Delta \log_\alpha^2 \Delta}{\alpha^{\delta-2}} \left(1 + \frac{1}{(\ln n)(\log_\alpha^2 \Delta)} \right)^{\log_\alpha \Delta} \\
&\leq \frac{3e^5 \Delta \log_\alpha^2 \Delta}{\alpha^{\delta-2}}
\end{aligned}$$

Hence the cost of the algorithm **GroveMatch** is $O\left(\frac{\Delta}{\alpha^{\delta-2}} \log^2 \Delta\right)$. However, note that **TreeMatch** only pays positive cost on X if for any optimal solution there is at least one request that such a solution must pay positive cost for in X . The reason for this is that if **TreeMatch**(X) moves the car out of a vertex v in X , then there are no in-commission parking spots left in v , and therefore every algorithm would have to move the car out of v . Since every edge in X has distance $\frac{\Delta}{\alpha^{\delta-1}}$, this gives us that **GroveMatch** must be $O(\alpha \log^2 \Delta) = O(\log n \log^3 \Delta)$ competitive on the metric d_G . \square

Together with Corollary 24, Lemma 26 gives us the following theorem:

Theorem 27. *GroveMatch is $O(\log^6 \Delta \log^2 n)$ -competitive for online metrical search instances.*

Lemma 28. *GroveMatch is a monotone algorithm for online metrical matching.*

Proof. Consider a tree X_i and vertex $x_i \in X_i$ considered in the **GroveMatch** algorithm. Then by the monotonicity of the algorithm **TreeMatch** the probability that a request would arrive at a particular y_i if the request arrived in a vertex x_j on the path from x_i to y_i has to be at least the probability that a request arriving at x_i moves to y_i . And the routing in the X_k 's, $k > i$, is independent of whether the request arrived at x_i or x_j . \square

Acknowledgements

We thank Anupam Gupta for his guidance, throughout the research process, that was absolutely critical to obtaining these results. We thank Amos Fiat for introducing us to this posted-price research area, and for several helpful discussions.

References

- [1] *Calgar ParkPlus Homepage*. URL: <https://www.calgaryparking.com/parkplus>.
- [2] *SFpark Homepage*. URL: <http://sfpark.org/>.
- [3] *SFpark Wikipedia page*. URL: <https://en.wikipedia.org/wiki/SFpark>.
- [4] Antonios Antoniadis, Neal Barcelo, Michael Nugent, Kirk Pruhs, and Michele Scquizzato. A $o(n)$ -competitive deterministic algorithm for online matching on a line. In *Workshop on Approximation and Online Algorithms*, pages 11–22, 2014.
- [5] Antonios Antoniadis, Carsten Fischer, and Andreas Tönnis. A collection of lower bounds for online matching on the line. In *Latin American Symposium on Theoretical Informatics*, volume 10807 of *LNCS*, pages 52–65. Springer, 2018.

- [6] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: A meta-algorithm and applications. *Theory of Computing*, 8:121–164, 2012.
- [7] James Aspnes, Yossi Azar, Amos Fiat, Serge Plotkin, and Orli Waarts. On-line routing of virtual circuits with applications to load balancing and machine scheduling. *Journal of the ACM*, 44(3), May 1997.
- [8] Yossi Azar, Bala Kalyanasundaram, Serge A. Plotkin, Kirk Pruhs, and Orli Waarts. On-line load balancing of temporary tasks. *Journal of Algorithms*, 22(1):93–110, 1997.
- [9] Nikhil Bansal, Niv Buchbinder, Anupam Gupta, and Joseph Naor. A randomized $O(\log^2 k)$ -competitive algorithm for metric bipartite matching. *Algorithmica*, 68(2):390–403, 2014.
- [10] Yair Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. In *Symposium on Foundations of Computer Science*, pages 184–193, 1996.
- [11] Yair Bartal. On approximating arbitrary metrics by tree metrics. In *ACM Symposium on Theory of Computing*, pages 161–168, 1998.
- [12] Allan Borodin and Ran El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge University Press, USA, 1998.
- [13] Ilan Reuven Cohen, Alon Eden, Amos Fiat, and Lukasz Jez. Pricing online decisions: Beyond auctions. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 73–91, 2015.
- [14] Ilan Reuven Cohen, Alon Eden, Amos Fiat, and Lukasz Jez. Dynamic pricing of servers on trees. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 145 of *LIPIcs*, pages 10:1–10:22, 2019.
- [15] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *Journal of Computer and System Sciences*, 69(3):485–497, 2004.
- [16] Michal Feldman, Amos Fiat, and Alan Roytman. Makespan minimization via posted prices. In *ACM Conference on Economics and Computation*, pages 405–422, 2017.
- [17] Bernhard Fuchs, Winfried Hochstättler, and Walter Kern. Online matching on a line. *Theoretical Computer Science*, 332(1-3):251–264, 2005.
- [18] Anupam Gupta and Kevin Lewi. The online metric matching problem for doubling metrics. In *International Colloquium on Automata, Languages, and Programming*, pages 424–435, 2012.
- [19] Sungjin Im, Benjamin Moseley, Kirk Pruhs, and Clifford Stein. Minimizing maximum flow time on related machines via dynamic posted pricing. In *European Symposium on Algorithms*, pages 51:1–51:10, 2017.
- [20] Bala Kalyanasundaram and Kirk Pruhs. Online weighted matching. *Journal of Algorithms*, 14(3):478–488, 1993.
- [21] Samir Khuller, Stephen G. Mitchell, and Vijay V. Vazirani. On-line algorithms for weighted bipartite matching and stable marriages. *Theoretical Computer Science*, 127(2):255–267, 1994.

- [22] Elias Koutsoupias and Akash Nanavati. The online matching problem on a line. In *Workshop on Approximation and Online Algorithms*, volume 2909 of *Lecture Notes in Computer Science*, pages 179–191. Springer, 2003.
- [23] Mark Manasse, Lyle McGeoch, and Daniel Sleator. Competitive algorithms for on-line problems. In *ACM Symposium on Theory of Computing*, page 322–333, 1988.
- [24] Adam Meyerson, Akash Nanavati, and Laura J. Poplawski. Randomized online algorithms for minimum metric bipartite matching. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 954–959, 2006.
- [25] Krati Nayyar and Sharath Raghvendra. An input sensitive online algorithm for the metric bipartite matching problem. In *Symposium on Foundations of Computer Science*, pages 505–515, 2017.
- [26] Sharath Raghvendra. Optimal analysis of an online algorithm for the bipartite matching problem on a line. In *Symposium on Computational Geometry*, volume 99 of *LIPICs*, pages 67:1–67:14, 2018.
- [27] Donald Shoup and Greg Pierce. *SFpark: Pricing Parking by Demand*, 2013. URL: <https://www.accessmagazine.org/fall-2013/sfpark-pricing-parking-demand/>.