

Edge Weighted Online Windowed Matching

ITAI ASHLAGI, Stanford University

MAXIMILIEN BURQ, Massachusetts Institute of Technology

CHINMOY DUTTA, Lyft

PATRICK JAILLET*, Massachusetts Institute of Technology

AMIN SABERI, Stanford University

CHRIS SHOLLEY, Lyft

Motivated by applications from ride-sharing and kidney exchange, we study the problem of matching agents who arrive at a marketplace over time and leave after d time periods. Agents can only be matched while they are present in the marketplace. Each pair of agents can yield a different match value, and the planner's goal is to maximize the total value over a finite time horizon.

First we study the case in which vertices arrive in an adversarial order. We provide a randomized $1/4$ -competitive algorithm building on a result by Feldman et al. [14] and Lehmann et al. [23]. We extend the model to the case in which departure times are drawn independently from a distribution with non-decreasing hazard rate, for which we establish a $1/8$ -competitive algorithm.

When the arrival order is chosen uniformly at random, we show that a batching algorithm, which computes a maximum-weighted matching every $(d + 1)$ periods, is 0.279 -competitive.

CCS Concepts: • **Theory of computation** → **Online algorithms**; *Theory and algorithms for application domains*; Graph algorithms analysis.

Additional Key Words and Phrases: Online Matching, Online Windowed Matching, Carpooling, Ride sharing, Kidney Exchange, Batching.

ACM Reference Format:

Itai Ashlagi, Maximilien Burq, Chinmoy Dutta, Patrick Jaillet, Amin Saberi, and Chris Sholley. 2019. Edge Weighted Online Windowed Matching. In *ACM EC '19: ACM Conference on Economics and Computation (EC '19)*, June 24–28, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3328526.3329573>

1 INTRODUCTION

Traffic congestion is a severe problem in metropolitan areas around the world. A resident in Los Angeles is estimated to lose around \$6,000 per year due to spending extra hours in traffic (Economist 2014). A couple of ways to relieve congestion are pricing [34] and carpooling, and online platforms and other technological advances are now available to assist with these tasks [28].¹

Online ride-sharing platforms now offer the option to share rides which may contribute to reducing congestion. Passengers who share rides also share the cost thereby paying a lower price for their rides and can further use high occupancy lanes. However, a passenger may also experience

*Jaillet acknowledges the research support of the Office of Naval Research grants N00014-15-1-2083 and N00014-18 -1-2122.

¹See [28] for a discussion about the complementarities between autonomous vehicles, carpooling and pricing.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '19, June 24–28, 2019, Phoenix, AZ, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6792-9/19/06...\$15.00

<https://doi.org/10.1145/3328526.3329573>

additional waiting, detours, and less privacy. Facing these trade-offs, ride-sharing platforms seek to increase the volume of shared rides, which will in turn help in reducing congestion.

This paper is concerned with the problem of matching two passengers together for carpooling in the context of an on-demand ride-sharing platform. We present and study a stylized graph-theoretic matching problem that captures the following three key features faced by ride-sharing platforms. First is spatial; the farther away two passengers' pickup locations or the more mismatch between their routes, the higher the disutility from being matched. Second is temporal; ride-sharing platforms offer passengers the option of waiting for a few minutes, during which time they are eligible to be matched to one another. Third, the platform faces uncertainty about future demand, and needs to commit to matching decisions within the waiting period that the passengers have agreed to with imperfect knowledge of future matching opportunities.

The spatial aspect is captured by modeling passengers as vertices of a graph. Time is discrete and one vertex arrives at each time period. An edge of the graph has a non-negative weight which is the utility from matching the two passenger-vertices. A vertex cannot match more than d periods after its arrival; after d units of time the vertex becomes *critical* and departs. It is useful to interpret d as a service quality set by the platform: it is the maximum waiting period for a passenger to get matched and if the platform is unable to match a given passenger within d units of time since arrival, then the passenger must be assigned to a single ride.

The goal is to find a weighted matching with a large total weight in an online manner. This means that the decision for every vertex has to be made no later than d periods after its arrival (this differs from the classic online bipartite matching literature, in which $d = 0$). There is no a priori information about edge weights and the underlying graph may be arbitrary and hence *non-bipartite*.

Contributions

Our first results are given for the setting in which the vertices arrive in an **adversarial order**. We introduce for this setting a $1/4$ -competitive algorithm, termed *Postponed Greedy* (PG). We further show that no algorithm achieves a competitive ratio that is higher than $1/2$.

The key idea behind PG is to look at a virtual bipartite graph, in which each vertex is duplicated into a "buyer" and a "seller" copy. We enforce that the seller copy does not match before the vertex becomes critical. This enables us to postpone the matching decision until we have more information about the graph structure and the likely matchings. We then proceed in a manner similar to [14]: tentatively match each new buyer copy to the seller that maximizes its margin, i.e., the difference between the weight of the edge with the seller and the value of the seller's current match.

We extend the model to the case where the departure of vertices are determined stochastically. We show that when the departure distribution is memoryless and realized departure times are revealed to the algorithm just when becoming critical, one can adapt the PG algorithm to achieve a competitive ratio of $1/8$. It is worth noting that when departure times are chosen in an adversarial manner no algorithm can achieve a constant competitive ratio.

Next we study the setting, in which vertices arrive in a **random order**. We analyze a *batching* algorithm which, every $d + 1$ time steps, computes a maximum weighted matching among the last $d + 1$ arrivals (batching-like algorithms are commonly used in ride-sharing platforms).² Vertices that are left unmatched are discarded forever (again, one can interpret discards as single rides). We show that when the number of vertices is sufficiently large, batching is 0.279-competitive.

The analysis proceeds in three steps. First, we show that the competitive ratio is bounded by the solution to a graph covering problem. Second, we show how a solution for small graphs can be extended to covers for larger graphs. Finally, we establish a reduction that allows us to consider

²Batching is also used in many kidney exchange platforms (from personal communication).

only a finite set of values for d . We conclude with a computer-aided argument for graphs in the finite family.

Related literature

There is a growing literature related to ride-sharing. Santi et al. [31] finds that about 80% of rides in Manhattan could be shared by two passengers. Many studies focus on rebalancing or dispatching problems without pooling, e.g., Banerjee et al. [9], Pavone et al. [30], Santi et al. [31], Spieser et al. [32], Zhang and Pavone [35]. Alonso-Mora et al. [3] studies real-time high-capacity ride-sharing. They do not consider, however, a graph-theoretic online formulation for matching rides.

This paper is closely related to the online matching literature. In the classic problem, introduced in Karp et al. [21], the graph is bipartite with vertices on one side waiting, while others are arriving sequentially and have to be matched *immediately* upon arrival. This work has numerous extensions, for example to stochastic arrivals and in the adwords context [15, 16, 20, 25, 27]. See [26] for a detailed survey. Our work contributes to this literature in three ways. First and foremost, our graph can be non-bipartite, which is the case in applications such as ride-sharing and kidney exchange. Second, all vertices arrive over time and remain for some given time until they are matched or hit their deadline and depart. Third, we provide algorithms that perform well on edge-weighted graphs. Closely related is Huang et al. [18, 19], which studies a similar model to ours in the non-weighted case, but allow departure times to be adversarial.

Several papers consider the problem of dynamic matching in the edge-weighted case. Feldman et al. [14] find that in the classic online bipartite setting, no algorithm achieves a constant approximation. They introduce a *free disposal* assumption, which allows to discard a matched vertex in favor of a new arriving vertex. They find, based on an algorithm by Lehmann et al. [23], that a greedy algorithm that matches a vertex to the highest marginal vertex, is 0.5-competitive. We build on this result for a special class of bipartite graphs. In the adversarial setting, Ashlagi et al. [5], Emek et al. [13] study the problem of minimizing the sum of distances between matched vertices and the sum of their waiting times. In their model no vertex leaves unmatched. Few papers consider the stochastic setting [8, 17, 29]. These papers find that some waiting before matching is beneficial for improving efficiency.

Related to our work are some works on job or packet scheduling. Jobs arrive online to a buffer, and reveal upon arrival the deadline by which they need to be scheduled. The algorithm can schedule at most one job per time and the value of scheduling a job is independent of the time slot. Constant approximation algorithms are given by Chin et al. [11] and Li et al. [24].

Finally, there is a growing literature that focuses on dynamic matching motivated by kidney exchange [4, 6, 12, 33]. These papers focus mostly on random graphs with no weights. Closer to our paper is [2], which finds that in a sparse random graph, knowledge about the departure time of a vertex is beneficial and matching a vertex only when it becomes critical performs well. Our work differs from these papers in two ways: we consider the edge-weighted case, and, we make no assumption on the graph structure.

2 MODEL

Consider a weighted graph G with n vertices indexed by $i = 1, \dots, n$. Vertices arrive sequentially over n periods and let $\sigma(i)$ denote the arrival time of vertex i . Let $v_{ij} \geq 0$ denote the weight on the undirected edge (i, j) between vertices i and j .

For vertices i and j with $\sigma(i) < \sigma(j)$, the weight v_{ij} on the edge between i and j is observed only after vertex j has arrived.

For $d \geq 1$, the **online graph with deadline** d , denoted by $G_{d,\sigma}$, has the same vertices as G , and the edge between i and j in G exists if and only if $|\sigma(i) - \sigma(j)| \leq d$. We say that i becomes **critical** at

time $\sigma(i) + d$, at which time the online algorithm needs to either match it and collect the associated edge weight, or let it remain unmatched.

We will consider two settings regarding how arrivals are generated. In the Adversarial Order (AO) setting, we assume that $\sigma(i) = i$. In the Random Order (RO) setting, we assume that σ is sampled uniformly at random among all possible permutations S_n of $[1, n]$.

The goal is to find an online algorithm that generates a matching with high total weight. More precisely, we seek to design a randomized online algorithm that obtains in expectation a high fraction of the expected maximum-weight of a matching over $G_{d,\sigma}$.

To illustrate a natural tradeoff, consider the example in Figure 1 for $d = 1$. At time 2, the planner can either match vertices 1 and 2 or let vertex 1 remain unmatched. This simple example shows that no deterministic algorithm can obtain a constant competitive ratio. Furthermore, no algorithm can achieve a competitive ratio higher than $1/2$.




Fig. 1. Let $d = 1$. Therefore, there is no edge between vertices 1 and 3. The algorithm needs to decide whether to match 1 with 2 and collect $v_{1,2}$ without knowing y .

3 ADVERSARIAL ORDER (AO) OF ARRIVALS

The example in Figure 1 illustrates a necessary condition for the algorithm to achieve a constant competitive ratio: with some probability, vertex 2 needs to forgo the match with vertex 1. We ensure this property by assigning every vertex to be either a *seller* or a *buyer*. We then prevent sellers from matching before they become critical, while we allow buyers to be matched at any time.

It will be useful to first study a special case, in which the underlying graph G is bipartite, with sellers on one side and buyers on the other, and in the online graph a buyer and a seller cannot match if the buyer arrives before the seller. For such online graphs we show that a greedy algorithm given by Feldman et al. [14] is 0.5-competitive. We then build on this algorithm to design a randomized $1/4$ -competitive algorithm for arbitrary graphs.

3.1 Bipartite constrained online graphs

Let G be a bipartite graph and σ be the order of arrivals. The online graph $G_{d,\sigma}$ is called **constrained bipartite** if for every seller s and every buyer b , there is no edge between s and b if $\sigma(b) < \sigma(s)$, i.e. b and s cannot match if b arrives before s .

Consider the following greedy algorithm, which attempts to match buyers in their arriving order. An arriving buyer b is matched to the seller with the highest marginal value if the marginal value is positive. If the seller is already matched to another buyer b' , b' becomes unmatched and never gets matched again. Formally:

PROPOSITION 3.1 (FELDMAN ET AL. [14]). *The greedy algorithm is 0.5-competitive for online bipartite constrained graphs.*

Feldman et al. [14] prove that this algorithm is 0.5-competitive for an online matching problem with *free disposal*. In their setting all seller exists and buyer arrive one at a time. The algorithm provides the same guarantees for constrained bipartite graph since, by construction, there is no harm in assuming that all sellers exist rather than arriving over time. The key behind the proof is

Algorithm 1 Greedy algorithm [14]

- Input: constrained bipartite graph, $G_{d,\sigma}$.
 - For each arrival $i = 1, \dots, n$:
 - If i is a seller, initialize $p(i) = 0$, and $m(s) = \text{null}$.
 - If i is a buyer:
 - * Set $s = \operatorname{argmax}_{s' \in S} \{v_{is'} - p(s')\}$, where S denote the set of sellers already arrived.
 - * If $v_{is} - p(s) > 0$, set $m(s) = i$ and set $p(s) = v_{is}$.
 - When a seller s becomes critical: match it to $b = m(s)$ if $m(s) \neq \text{null}$.
-

that the value $p(s)$ function for each seller s is submodular. In fact the result is a special case of a result by Lehmann et al. [23], who study combinatorial auctions with submodular valuations.

3.2 Arbitrary graphs

In this section we extend the greedy algorithm for constrained bipartite graphs to arbitrary graphs. A naive way to generate an online constrained bipartite graph from an arbitrary one is to randomly assign each vertex to be either a seller or a buyer, independently and with probability $1/2$. Then only keep the edges between each buyer and all the sellers who arrived before her. Formally:

Algorithm 2 Naive Greedy

- Input: an online graph with deadline d , $G_{d,\sigma}$.
 - For each vertex $t = 1, \dots, n$:
 - Toss a fair coin to decide whether i is a *seller* or a *buyer*. Construct the online constrained bipartite graph $\tilde{G}(d, \sigma)$ by keeping only the edges between each buyer and the sellers who arrived before her.
 - Run the Greedy algorithm on $\tilde{G}(d, \sigma)$.
-

COROLLARY 3.2. *The naive greedy algorithm is $1/8$ -competitive for arbitrary online graphs.*

Observe that for vertices i, j with $\sigma(i) < \sigma(j)$, edge (i, j) in the original graph remains in the generated constrained bipartite graph with probability $1/4$ (if i is a seller and j is a buyer). We then use proposition 3.1 to prove that naive greedy is $1/8$ -competitive.

One source of inefficiency in the naive greedy algorithm is that the decision whether a vertex becomes a seller or a buyer is done independently at random and without taking the graph structure into consideration. We next introduce the *Postponed Greedy* algorithm that defers these decisions as long as possible in order to construct the constrained bipartite graph more carefully.

When a vertex k arrives, we add two copies of k to a virtual graph: a seller s_k and a buyer b_k . Let S_t and B_t be the set of sellers and buyers at arrival time t . On arrival, seller s_k does not have any edges, and buyer b_k has edges towards any vertex $s_l \in S_k$ with value $v_{l,k}$. Then we run the greedy algorithm with the virtual graph as input. When a vertex k becomes critical, s_k becomes critical in the virtual graph, and we compute its matches generated by greedy.

Both s_k and b_k can be matched in this process. If we were to honor both matches, the outcome would correspond to a 2-matching, in which each vertex has degree at most 2. Now observe that because of the structure of the constrained bipartite graph, this 2-matching does not have any cycles; it is just a collection of disjoint paths. We decompose each path into two disjoint matchings and choose each matching with probability $1/2$.

In order to do that, the algorithm must determine, for each original vertex k , whether the virtual buyer b_k or virtual seller s_k will be used in the final matching. We say that k is a *buyer* or *seller*

depending on which copy is used. We say that vertex k is *undetermined* when the algorithm has not yet determined which virtual vertex will be used. When an undetermined vertex becomes critical, the algorithm flips a fair coin to decide whether to match according to the buyer or seller copy. This decision is then propagated to the next vertex in the 2-matching: if k is a *seller* then the next vertex will be a *buyer* and vice-versa. That ensures that assignments are correlated and saves a factor 2 compared to uncorrelated assignments in the naive greedy algorithm.

Algorithm 3 Postponed Greedy (PG)

- Input: an online graph with deadline d , $G_{d,\sigma}$.
 - Process events at time t in the following way:
 - (1) *Arrival of a vertex k* :
 - (a) Set k 's status to be *undetermined*.
 - (b) *Add a virtual seller*: $S_t \leftarrow S_{t-1} \cup \{s_k\}$, $p(s_k) \leftarrow 0$ and $m(s_k) = \emptyset$.
 - (c) *Add a virtual buyer*: $B_t \leftarrow B_{t-1} \cup \{b_k\}$.
 - (d) *Find a virtual seller for the virtual buyer*: $s = \operatorname{argmax}_{s' \in S_t} v_{s', b_k} - p(s')$.
 - (e) *Match if marginal utility is positive*: If $v_{s, b_k} - p(s) > 0$, then tentatively match b_k to s by setting $m(s) \leftarrow b_k$ and $p(s) \leftarrow v_{s, b_k}$.
 - (2) *Vertex k becomes critical*:
 - (a) *Proceed if no match found*: If $m(s_k) = \emptyset$, do nothing.
 - (b) *Match in the virtual graph*: If $m(s_k) = b_l$. Set $S_t \leftarrow S_t \setminus \{s_k\}$, and $B_t \leftarrow B_t \setminus \{b_l\}$.
 - (c) If k 's status is *undetermined*, w.p $1/2$ set it to be either *seller* or *buyer*.
 - (i) *If k is a seller*: finalize the matching of k to l and collect the reward $v_{k,l}$. Set l 's status to be a buyer.
 - (ii) *If k is a buyer*: Set l 's status to be a seller.
-

THEOREM 3.3. *The postponed greedy (PG) algorithm is $1/4$ -competitive for arbitrary online graphs.*

All missing proofs appear in the full version.

3.3 Lower bounds

CLAIM 3.4. *When the input is a constrained bipartite graph:*

- No deterministic algorithm can obtain a competitive ratio above $\frac{\sqrt{5}-1}{2} \approx 0.618$.
- No randomized algorithm can obtain a competitive ratio above $\frac{4}{5}$.

Next we show that our analysis for PG is tight.

CLAIM 3.5. *There exists a constrained bipartite graph for which PG is $1/(4-2\epsilon)$ -competitive.*

4 RANDOM PERMUTATION (RP) OF ARRIVALS

In some cases, the vertices can be assumed to come from a distribution that is unknown to the online algorithm. One way to model this is to assume that the adversary chooses the underlying graph, but that the vertices arrive in random order.

4.1 The batching algorithm

The *batching* algorithm computes a maximum-weight matching every $d + 1$ time steps. Every vertex in the matching is then matched, and all other vertices in the batch are discarded.

THEOREM 4.1. *Batching is $(0.279 + O(1/n))$ -competitive.*

The proof of Theorem 4.1 works in three steps. In a first step, we reduce the analysis of the competitive ratio of *Batching* to a graph covering problem. More precisely, we show that it is enough to cover C_n^d , the cycle with n vertices to the power d , with ensembles of cliques. Second, we show how a cover for small n can be extended to any n at the cost of a small rounding error. Finally, we establish a reduction that allows us to consider only a finite set of values for d . We conclude with a computer-aided argument for graphs in the finite family.

We provide here the main steps of the proof.

4.1.1 Reducing to a graph theoretic problem. There is no harm in assuming that the underlying graph G is a complete. Recall that S_n is the set of all permutations over integers $1, \dots, n$. For any deadline d and any arrival sequence $\sigma \in S_n$, we define the *path* graph $P_n^d(\sigma)$ with edge-weight $v_{ij} = 1$ if $|\sigma(i) - \sigma(j)| \leq d$, and $v_{ij} = 0$ otherwise.³

Note that every batch in the algorithm has $d + 1$ vertices except the last batch which may have fewer vertices. Let $b_i(\sigma, d)$ be the batch of vertex i under permutation σ and batch size $d + 1$: $b_i(\sigma, d)$ is the unique integer such that $(d + 1)(b_i - 1) < \sigma(i) \leq (d + 1)b_i$. We define the *batched* graph $B_n^d(\sigma)$ with edge-weight $v_{ij} = 1$ if i and j are in the same batch (i.e. $b_i(\sigma, d + 1) = b_j(\sigma, d + 1)$), and $v_{ij} = 0$ otherwise.⁴

For any $n \geq d \geq 1$, denote C_n^d to be the n -cycle to the power d .

Definition 4.2 (Graph operations). For any two graphs H and H' with vertices $1, \dots, n$ and respective edge weights v_{ij}, v'_{ij} , we define the following:

- (i) The linear combination $aH + bH'$ denotes the graph with edge weights $av_{ij} + bv'_{ij}$,
- (ii) The product $H * H'$ denotes the graph with edge weights $v_{ij} * v'_{ij}$, and
- (iii) We say that H is a *cover* of H' if for all i, j , $v_{i,j} \geq v'_{ij}$.

For any graph H , let $m(H)$ denote the value of a maximum-weight matching over H . Observe that when the arrival sequence is σ , the graph $P_n^d(\sigma) * G = G(d, \sigma)$ and therefore the offline algorithm collects $m(P_n^d(\sigma) * G)$. Note that the online algorithm collects $m(B_n^d(\sigma) * G)$.

REMARK 4.3. Observe that for any graphs H, H', G and any $a, b \in \mathbb{R}$, we have:

- $m(aH + bH') \leq am(H) + bm(H')$.
- If H is a cover of H' , then, $m(H * G) \geq m(H' * G)$.

Definition 4.4 (Periodic permutation). For $p < n$ such that p divides n , we say that a permutation $\sigma \in S_n$ is p -periodic if for all $i \in [1, n - p]$, $\sigma(i + p) \equiv \sigma(i) + p \pmod n$.

We say that a permutation σ is periodic if there exists p such that σ is p -periodic.

Definition 4.5 ((α, d)-cover). Let F be an unweighted graph with n vertices. We say that a set of permutations $\{\sigma_1, \dots, \sigma_K\} \in S_n$ forms an (α, d) -cover of F if there exist values $\lambda_1, \dots, \lambda_K \in [0, 1]$ such that:

- (i) $\sum_{k \leq K} \lambda_k B_n^d(\sigma_k)$ is a cover of F .
- (ii) $\sum_{k \leq K} \lambda_k = \alpha$.

We say that an (α, d) -cover is p -periodic if for all k , σ_k is p -periodic.

The next proposition will allow us to abstract away from the weights that are chosen by the adversary. For any graph H , we denote by H_{ij} the weight v_{ij} in H .

PROPOSITION 4.6. *If there exists an (α, d) -cover of C_n^d , then batching is $1/\alpha$ -competitive.*

³Note that $P_n^d(\sigma)$ corresponds to the path $(\sigma(1), \sigma(2)), (\sigma(2), \sigma(3)), \dots, (\sigma(n - 1), \sigma(n))$ taken to the power d .

⁴Note that $B_n^d(\sigma)$ is a collection of disjoint $(d + 1)$ -cliques.

We have reduced the analysis of Batching to a graph-theoretic problem without edge weights. In what follows, we will show that we can reduce the problem further to find covers of C_n^d for only small values of n and d .

4.1.2 *Reducing n : periodic covers.* We now wish to find (α, d) -covers for C_n^d for every n and d . In Proposition 4.7, we show that it is sufficient to find periodic covers for small values of n .

PROPOSITION 4.7. *Let p be a multiple of $d + 1$, and n_1 a multiple of p . Any p -periodic (α, d) -cover of $C_{n_1}^d$ can be extended into an $(\alpha + O(p/n), d)$ -cover of C_n^d for any $n \geq n_1$.*

The proof of Proposition 4.7 first considers the case in which n is a multiple of p and then analyzes the general case.

4.1.3 *Reducing d : cycle contraction.* In Proposition 4.7, we show that it is enough to find periodic (α, d) -covers of C_n^d for small values of n . Next, we provide a reduction that enables us to consider only a finite set of values for d .

The key idea of the reduction is that we can contract vertices of C_n^d into n/u groups of u vertices. The resulting graph also happens to be a cycle $C_{n/u}^{(d+1)/u}$. In Proposition 4.9, we provide a way to expand an $(\alpha, u - 1)$ -cover on the contracted graph into an (α, d) cover on the original graph.

Definition 4.8 (Cycle contraction). For any n, d and an integer u which divides n , we define the u -contraction $f_u(C_n^d)$ to be the graph with vertices $a_k = \{uk + 1, \dots, u(k + 1)\}$ for $k \in [0, n/u - 1]$, and edges (a_k, a_l) if and only if there exist $i \in a_k$ and $j \in a_l$ with an edge (i, j) in C_n^d .




Fig. 2. Left: C_{12}^3 , with contraction for $u = 2$. Right: Contracted graph $f(C_{12}^3) = C_6^2$ with vertices $a = \{1, 2\}$, $b = \{3, 4\}$, ... $f = \{11, 12\}$.

PROPOSITION 4.9. *Fix $d \geq 1$. For $d + 1 > k \geq 1$, suppose that there is a periodic $(\alpha, k - 1)$ -cover of C_{rk}^k .*

- (i) *For any integer r , if k divides $d + 1$ then there exists a periodic (α, d) -cover of $C_{r(d+1)}^d$.*
- (ii) *In general, if v is the remainder of the euclidian division of $d + 1$ by k , then there exists a periodic $(\alpha(1 + v/d_{d+1-v})^2, d)$ -cover of $C_{r(d+1)}^d$.*

4.1.4 *Final step: Computer-aided proof of factor 2.79.* We apply here apply Proposition 4.7 with $p = 2(d + 1)$ and $n_1 = 4(d + 1)$. Let Ω_d be the set of $2(d + 1)$ -periodic permutations of $1, \dots, 4(d + 1)$. We can find covers for $C_{4(d+1)}^d$ using the following linear program:

$$\begin{aligned} \min \quad & \sum_{\sigma \in \Omega_d} \lambda_\sigma \\ \text{s.t.} \quad & \sum_{\sigma \in \Omega_d} \lambda_\sigma \mathbb{I}[b_i(\sigma, d) = b_j(\sigma, d)] \geq 1, \quad \forall (i, j) \in C_{4(d+1)}^d \\ & \lambda_\sigma \in \mathbb{R}^+, \quad \sigma \in \Omega_d \end{aligned} \tag{LP}_d$$

PROPOSITION 4.10. *Let α_d be the solution to LP_d . Let $\alpha = \sup_{d \geq 1} \alpha_d$. Batching is $(1/\alpha + O(1/n))$ -competitive.*

The Linear program (LP_d) has $O(d!)$ variables, and solving it may not be computationally possible when d is large. Using Proposition 4.9, we provide a way to find upper bounds on α_d by solving a different LP on a smaller graph.

4.2 Lower bound in random order.

PROPOSITION 4.11. *No algorithm is more than $\frac{1}{2}$ -competitive even under the random arrival order.*

5 EXTENSIONS

5.1 Stochastic departures in the AO setting

We relax the assumption that all vertices depart after exactly d time steps.

We therefore focus on the stochastic case, in which the departure time d_i of vertex i is sampled independently from a distribution \mathcal{D} . We assume that the realizations d_i are only known at the time i becomes critical.

PROPOSITION 5.1. *Suppose that there exists $\alpha \in (0, 1)$ such that \mathcal{D} satisfies the property that for all $i < j$,*

$$\mathbb{P}[i + d_i \leq j + d_j | i + d_i \geq j] \geq \alpha.$$

Then PG is $\alpha/4$ -competitive.

For a memory-less process, conditional on i still being present when j arrives, the probability that i departs first is exactly $1/2$. Therefore the above condition is satisfied with $\alpha = 1/2$.

COROLLARY 5.2. *PG is $1/8$ -competitive when \mathcal{D} is memory-less.*

5.2 Look-ahead under RP setting

We assume now that the online algorithm knows vertices that will arrive in l time steps (and their adjacent edges). We can update the Batching Algorithm in the following way: every $d + l + 1$ time steps, compute a maximum-weight matching on both the current vertices and the next l arrivals. Match vertices as they become critical according to the matching, and discard unmatched vertices. Note that this is the same as running Batching when the deadline is $d + l$.

PROPOSITION 5.3. *There exists an $(\frac{d+l+1}{l+1}, d + l)$ -cover of C_n^d .*

COROLLARY 5.4. *Batching with l -lookahead is $\frac{l+1}{d+l+1}$ -competitive when n is large.*

5.3 AO setting: alternative to Greedy

Observe that the greedy algorithm discards a buyer that becomes unmatched and does not attempt to subsequently rematch it. In this section, we introduce the Dynamic Deferred Acceptance (DDA) algorithm, which takes as input a constrained bipartite graph and returns a matching (formally presented below). Although the DDA provides the same theoretical guarantees as greedy, we will see in Section 6 that it rationalizes *Re-Optimization* algorithms, which perform very well on practical instances of the carpooling problem.

The main idea is to maintain a tentative maximum-weight matching m at all times during the run of the algorithm. This tentative matching is updated according to an auction mechanism: every seller s is associated with a *price* p_s , which is initiated at zero upon arrival. Every buyer b that already arrived and yet to become critical is associated with a *profit margin* q_b which corresponds to the value of matching to their most preferred seller minus the price associated with that seller. Every time a new buyer arrives, she bids on her most preferred seller at the current set of prices. This triggers a bidding process that terminates when no unmatched buyer can profitably bid on a seller.

When a seller becomes critical, she is irrevocably matched to her tentative match. A buyer is discarded only if she is unmatched and becomes critical.

At any point t throughout the algorithm, we maintain a set of sellers S_t , a set of buyers B_t , as well as a matching m , a price p_s for every seller $s \in S_t$, and a marginal profit q_b for every buyer $b \in B_t$.

Algorithm 4 Dynamic Deferred Acceptance

- Input: an online graph with deadline d , $G_{d, \sigma}$.
 - Process each event in the following way:
 - (1) *Arrival of a seller s* : Initialize $p_s \leftarrow 0$ and $m(s) \leftarrow \emptyset$.
 - (2) *Arrival of a buyer b* : Start the following *ascending auction*.

Repeat

 - (a) Let $q_b \leftarrow \max_{s' \in S_t} v_{s', b} - p_{s'}$ and $s \leftarrow \operatorname{argmax}_{s' \in S_t} v_{s', b} - p_{s'}$.
 - (b) If $q_b > 0$ then
 - (i) $p_s \leftarrow p_s + \epsilon$.
 - (ii) $m(s) \leftarrow b$ (tentatively match s to b)
 - (iii) Set b to \emptyset if s was not matched before. Otherwise, let b be the previous match of s .

Until $q_b \leq 0$ or $b = \emptyset$.
 - (3) *Departure of a seller s* : If seller s becomes critical and $m(s) \neq \emptyset$, finalize the matching of s and $m(s)$ and collect the reward of $v_{s, m(s)}$.
-

Our algorithm bears similarities to the auction algorithm by [10]. Prices in this auction increase by ϵ to ensure termination, and optimality is proven through ϵ -complementary slackness conditions. For the analysis, we consider the limit $\epsilon \rightarrow 0$ and assume the auction phase terminates with the maximum weight matching.⁵

The auction phase is always initiated at the existing prices and profit margins. This, together with the fact that the graph is bipartite, ensures that prices never decrease and marginal profits never increase throughout the algorithm. Furthermore, the prices and marginal profits of the sellers and buyers that are present in the “market” form an optimum dual for the matching linear program.

⁵One way to formalize this argument is through the Hungarian algorithm [22], where prices are increased simultaneously along an alternating path that only uses edges for which the dual constraint is tight.

LEMMA 5.5. Consider the DDA algorithm on a constrained bipartite graph.

- (1) Sellers' prices never decrease, and buyers' profit margins never increase.
- (2) At the end of every ascending auction, prices of the sellers and the marginal profits of the buyers form an optimal solution to the dual of the matching linear program associated with buyers and sellers present at that particular time.

Maintaining a maximum-weight matching along with optimum dual variables does not guarantee an efficient matching for the whole graph. The dual values are not always feasible for the offline problem. Indeed, the profit margin of some buyer b may decrease after some seller departs the market. This is because b may face increasing competition from new buyers, while the bidding process excludes sellers that have already departed the market (whether matched or not). The main result in this section is established using the primal-dual framework.

PROPOSITION 5.6. DDA is $1/2$ -competitive for constrained bipartite graphs.

6 NUMERICAL RESULTS

The results from section 5.3 motivate a modified algorithm, termed *Re-Optimization*, in which vertices are no longer separated into buyers and sellers. At each time step, *Re-Opt* re-computes a tentative (non-bipartite) maximum weight matching over all the vertices that are present in the graph. Only vertices that are critical are matched to their tentative match at that time. This modified algorithm does not have theoretical guarantees, but we will show that it performs well on data-driven compatibility graphs.

We compare the *Re-Opt* against three benchmarks that have been previously considered, or that are commonly used in practice:

- The *Greedy* algorithm. The algorithm matches vertices as soon as possible to their available neighbor with the highest value (ties are broken in favor of the earliest arrival).
- The *Batching*(b) algorithm. The algorithm waits k times-steps and then finds a maximum-weight matching. Unmatched vertices are kept in the next batch.⁶ We report the best simulation results across parameters a number of batch sizes.
- The *Patient* algorithm. This algorithm waits until a vertex becomes critical, and matches it to the neighbor with the highest match value (ties are broken in favor of the earliest arrival). This allows to separate the value from knowing the time in which vertices become critical and the value of optimization.

Data

For this numerical experiment, we use a data set of all taxi rides in New York City over an hour time period⁷. For any pair k, l of trips, we can compute the Euclidean distance that would have been traveled had the two passengers taken the same taxi (with multiple stops). The value $v_{k,l}$ represents the "distance saved" by combining the trips.

This enables us to generate a dynamic graph in the following way. For $t \in [1, T]$:

- (1) Sample with replacement an arrival t from the data set.
- (2) For any vertex l that is present in the pool, compute the value $v_{t,l}$ of matching t to l .
- (3) Sample a departure time $d_t \sim \mathcal{D}$.

We consider three settings, termed *deterministic* and *exponential* and *uniform* respectively, in which \mathcal{D} is either constant with value d , exponentially distributed with mean d , or uniform $[0, 2d]$. We will report simulation results for d between 50 and 150.

⁶See [1, 7] in the case of ride-sharing and kidney exchange respectively.

⁷<http://www.andresmh.com/nyctaxitrips/>

Results

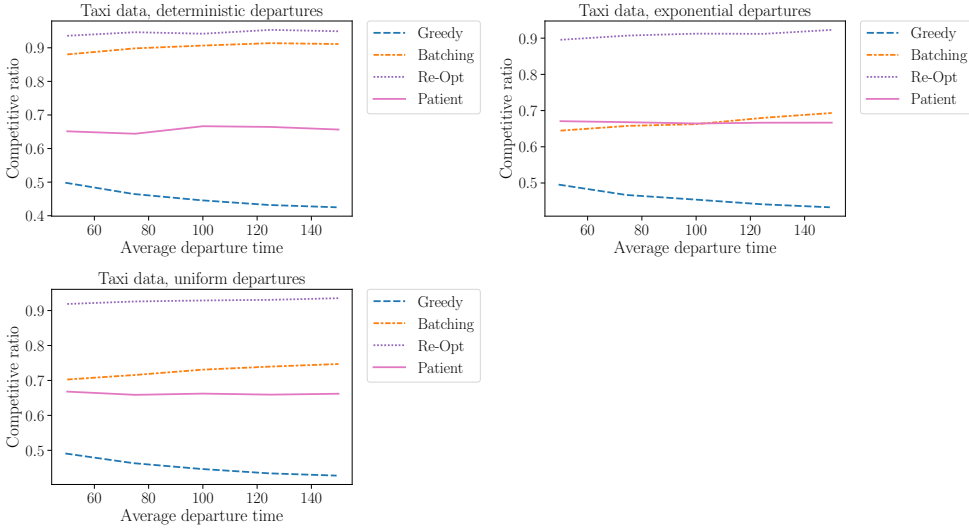


Fig. 3. Performance of our 4 algorithms on taxi data (weighted compatibility graph).

In Figure 3, we observe that both the *Patient* and *Batching* algorithms outperform *Greedy* across all three departure settings. Intuitively, having vertices wait until they become critical helps to thicken the market and gives vertices higher valued match options. We notice that when departures are deterministic, *Batching* with the optimal batch size will be almost as efficient as *Re-Opt*. However when the departures are stochastic, there is value in matching vertices as they become critical.

We observe that *Re-Opt* outperforms all other algorithm, although in the cases where the departures are deterministic, *Batching* performs close to *Re-Opt* when the batch size k is carefully chosen. This shows the value of both having information on agents' departure times and also subsequent optimization.

It is important to note that the experiments we ran do not take into account the cost of waiting. We think that a richer model that accounts for this would be an interesting future direction. Two interesting areas for future work include the setting when the information about agent's departure times is uncertain, as well as models that are less restrictive than the adversarial setting (see, e.g., Ozkan and Ward [29]).

7 CONCLUSION

This paper introduces a model for dynamic matching, in which all agents arrive and depart after some deadline. Match values are heterogeneous and the underlying graph is non-bipartite. We study online algorithms for two settings, where vertices arrive in an adversarial or random order.

In the adversarial arrival case, we introduce two new $1/4$ -competitive algorithms when departures are deterministic and known in advance. We also provide a $1/8$ -competitive algorithm when departures are stochastic, i.i.d, memoryless, and known at the time a vertex becomes critical. Finally we show that no online algorithm is more than $1/2$ -competitive.

In the random arrival case, we show that a batching algorithm is 0.279 -competitive. We also show that with knowledge of future arrivals, its performance guarantee increases towards 1.

Importantly, our model imposes restrictions on the departure process and requires the algorithm to know when vertices become critical. Other than closing the gaps between the upper bound $1/2$ and the achievable competitive ratios, we point out a just a few interesting directions for future research. Our model imposes that matches retain the same value regardless of when they are conducted. An interesting direction is to account for agents' waiting times. A different interesting objective is to achieve both a high total value and a large fraction of matched agents. Finally, it is interesting to consider the stochastic setting with prior information over weights and future arrivals.

REFERENCES

- [1] Niels AH Agatz, Alan L Erera, Martin WP Savelsbergh, and Xing Wang. 2011. Dynamic ride-sharing: A simulation study in metro Atlanta. *Transportation Research Part B: Methodological* 45, 9 (2011), 1450–1464.
- [2] Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. 2017. Thickness and information in dynamic matching markets. *Available at SSRN 2394319* (2017).
- [3] Javier Alonso-Mora, Samitha Samaranayake, Alex Wallar, Emilio Frazzoli, and Daniela Rus. 2017. On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proc Natl Acad Sci USA* (2017).
- [4] Ross Anderson, Itai Ashlagi, David Gamarnik, and Yash Kanoria. 2015. A dynamic model of barter exchange. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 1925–1933.
- [5] Itai Ashlagi, Yossi Azar, Moses Charikar, Ashish Chiplunkar, Ofir Geri, Haim Kaplan, Rahul Makhijani, Yuyi Wang, and Roger Wattenhofer. 2017. Min-cost bipartite perfect matching with delays. In *LIPICs-Leibniz International Proceedings in Informatics*, Vol. 81. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [6] Itai Ashlagi, Maximilien Burq, Patrick Jaillet, and Vahideh Manshadi. 2017. On matching and thickness in heterogeneous dynamic markets. (2017).
- [7] I. Ashlagi, P. Jaillet, and V.H. Manshadi. 2013. Kidney Exchange in Dynamic Sparse Heterogenous Pools. *arXiv preprint arXiv:1301.3509* (2013).
- [8] M. Baccara, S. Lee, and L. Yariv. 2015. Optimal Dynamic Matching. (2015). Working paper.
- [9] Siddhartha Banerjee, Yash Kanoria, and Pengyu Qian. 2018. State Dependent Control of Closed Queueing Networks. In *Abstracts of the 2018 ACM International Conference on Measurement and Modeling of Computer Systems*. ACM, 2–4.
- [10] Dimitri P Bertsekas. 1988. The auction algorithm: A distributed relaxation method for the assignment problem. *Annals of operations research* 14, 1 (1988), 105–123.
- [11] Francis YL Chin, Marek Chrobak, Stanley PY Fung, Wojciech Jawor, Jiří Sgall, and Tomáš Tichý. 2006. Online competitive algorithms for maximizing weighted throughput of unit jobs. *Journal of Discrete Algorithms* 4, 2 (2006), 255–276.
- [12] John P Dickerson, Ariel D Procaccia, and Tuomas Sandholm. 2013. Failure-aware kidney exchange. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 323–340.
- [13] Yuval Emek, Shay Kutten, and Roger Wattenhofer. 2016. Online matching: haste makes waste!. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. ACM, 333–344.
- [14] Jon Feldman, Nitish Korula, Vahab Mirrokni, S Muthukrishnan, and Martin Pál. 2009. Online ad assignment with free disposal. In *International Workshop on Internet and Network Economics*. Springer, 374–385.
- [15] J. Feldman, A. Mehta, V. S. Mirrokni, and S. Muthukrishnan. 2009. Online Stochastic Matching: Beating 1-1/e. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 117–126.
- [16] G. Goel and A. Mehta. 2008. Online budgeted matching in random input models with applications to Adwords. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms (SODA)*. 982–991.
- [17] Ming Hu and Yun Zhou. 2016. Dynamic type matching. (2016).
- [18] Zhiyi Huang, Ning Kang, Zhihao Gavin Tang, Xiaowei Wu, Yuhao Zhang, and Xue Zhu. 2018. How to match when all vertices arrive online. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 17–29.
- [19] Zhiyi Huang, Binghui Peng, Zhihao Gavin Tang, Runzhou Tao, Xiaowei Wu, and Yuhao Zhang. 2019. Tight Competitive Ratios of Classic Matching Algorithms in the Fully Online Model. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*. 2875–2886.
- [20] P. Jaillet and X. Lu. 2013. Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations Research* 39, 3 (2013), 624–646.
- [21] R. M Karp, U. V. Vazirani, and V. V. Vazirani. 1990. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing (STOC)*. 352–358.

- [22] Harold W Kuhn. 1955. The Hungarian method for the assignment problem. *Naval Research Logistics (NRL)* 2, 1-2 (1955), 83–97.
- [23] Benny Lehmann, Daniel Lehmann, and Noam Nisan. 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55, 2 (2006), 270–296.
- [24] Fei Li, Jay Sethuraman, and Clifford Stein. 2005. An optimal online algorithm for packet scheduling with agreeable deadlines. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 801–802.
- [25] V. H. Manshadi, S. Oveis-Gharan, and A. Saberi. 2011. Online stochastic matching: online actions based on offline statistics. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1285–1294.
- [26] Aranyak Mehta. 2013. Online matching and ad allocation. *Foundations and Trends® in Theoretical Computer Science* 8, 4 (2013), 265–368.
- [27] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. 2007. Adwords and generalized online matching. *Journal of the ACM (JACM)* 54, 5 (2007), 22.
- [28] Michael Ostrovsky and Michael Schwarz. 2018. *Carpooling and the Economics of Self-Driving Cars*. Technical Report. National Bureau of Economic Research.
- [29] Erhun Ozkan and Amy R Ward. 2016. Dynamic matching for real-time ridesharing. (2016).
- [30] Marco Pavone, Stephen L Smith, Emilio Frazzoli, and Daniela Rus. 2012. Robotic load balancing for mobility-on-demand systems. *Int J Rob Res* (2012).
- [31] Paolo Santi, Giovanni Resta, Michael Szell, Stanislav Sobolevsky, Steven H. Strogatz, and Carlo Ratti. 2014. Quantifying the benefits of vehicle pooling with shareability networks. In *Proc Natl Acad Sci USA*.
- [32] Kevin Spieser, Samitha Samaranyake, Wolfgang Gruel, and Emilio Frazzoli. 2016. Shared-vehicle mobility-on-demand systems: A fleet operator’s guide to rebalancing empty vehicles. In *Transportation Research Board 95th Annual Meeting*.
- [33] M. U. Ünver. 2010. Dynamic Kidney Exchange. *Review of Economic Studies* 77(1) (2010), 372–414.
- [34] William Vickrey. 1965. Pricing as a tool in coordination of local transportation. In *Transportation economics*. NBER, 275–296.
- [35] Rick Zhang and Marco Pavone. 2014. Control of robotic mobility-on-demand systems: a queueing-theoretical perspective. In *Proceedings of Robotics: Science and Systems Conference*.