

How to Match when All Vertices Arrive Online

Zhiyi Huang* Ning Kang* Zhihao Gavin Tang* Xiaowei Wu[†] Yuhao Zhang*
Xue Zhu*

Abstract

We introduce a fully online model of maximum cardinality matching in which all vertices arrive online. On the arrival of a vertex, its incident edges to previously-arrived vertices are revealed. Each vertex has a deadline that is after all its neighbors' arrivals. If a vertex remains unmatched until its deadline, the algorithm must then irrevocably either match it to an unmatched neighbor, or leave it unmatched. The model generalizes the existing one-sided online model and is motivated by applications including ride-sharing platforms, real-estate agency, etc.

We show that the Ranking algorithm by Karp et al. (STOC 1990) is 0.5211-competitive in our fully online model for general graphs. Our analysis brings a novel charging mechanic into the randomized primal dual technique by Devanur et al. (SODA 2013), allowing a vertex other than the two endpoints of a matched edge to share the gain. To our knowledge, this is the first analysis of Ranking that beats 0.5 on general graphs in an online matching problem, a first step towards solving the open problem by Karp et al. (STOC 1990) about the optimality of Ranking on general graphs. If the graph is bipartite, we show that the competitive ratio of Ranking is between 0.5541 and 0.5671. Finally, we prove that the fully online model is strictly harder than the previous model as no online algorithm can be $0.6317 < 1 - \frac{1}{e}$ -competitive in our model even for bipartite graphs.

*Department of Computer Science, University of Hong Kong. {zhiyi, nkang, zhtang, yhzhang2, xzhu2}@cs.hku.hk

[†]Department of Computing, Hong Kong Polytechnic University. csxwu@comp.polyu.edu.hk. The work was done when the author was a postdoc at the University of Hong Kong.

1 Introduction

Online Bipartite Matching is a central problem in the area of online algorithms with a wide range of applications. Consider a bipartite graph where the left-hand-side is known in advance, while vertices on the right-hand-side arrive online in an arbitrary order. On the arrival of a vertex, its incident edges are revealed and the algorithm must irrevocably either match it to one of its unmatched neighbors or leave it unmatched. Karp et al. [KVV90] introduced the Ranking algorithm, which picks at the beginning a random permutation over offline vertices, and matches each online vertex to the first unmatched neighbor according to the permutation. Further, they proved that Ranking is $(1 - \frac{1}{e})$ -competitive and the best possible among online algorithms. The analysis of Ranking has been subsequently simplified in a series of papers [GM08, BM08, DJK13]. Further, it has been generalized to several extended settings, including the vertex-weighted case [AGKM11], the random arrival model [KMT11, MY11], and the Adwords problem [MSVV07, BJK07, DJ12].

However, all the above successful applications of Ranking crucially rely on the assumption that one side of the bipartite graph is known upfront. This assumption prevents us from applying the known positive results to some applications. Here is an example:

Example (Real Estate Agency). During a typical day of a real estate agent in Hong Kong, both tenants and landlords drop by in an online fashion. Tenants specify what kinds of apartments they are looking for as well as the deadlines before which they need to move in¹. Similarly, landlords list certain rules for tenant screening together with their deadlines. Tenants and landlords can be modeled as the vertices in a bipartite graph. There is an edge between a tenant-landlord pair if (1) they mutually satisfy each other’s conditions, and (2) their time windows (between their respective arrivals and deadlines) overlap. Real estate agents charge for each successful deal and, thus, seek to maximize the size of the bipartite matching.

This is clearly a bipartite matching problem with an online nature. However, it does not fit into the existing model in two fundamental ways. First, vertices from both sides of the bipartite graphs arrive online. Second, matching decision of each vertex is made at its deadline rather than its arrival. There are many other applications with similar flavors such as job market intermediary and organ transplantation.

A Fully Online Model. Motivated by these applications, we formulate the following alternative online model of bipartite matching. We call it **Fully Online Bipartite Matching** since vertices from both sides arrive online. Let there be an underlying bipartite graph that is completely unknown to the algorithm at the beginning. Each time step falls into one of the following two kinds:

- Arrival of v : A vertex v arrives; edges between v and previously-arrived vertices are revealed.
- Deadline of v : This is the last time a vertex v can be matched (if it is not matched yet).

The model further guarantees that all edges incident to a vertex are revealed before its deadline. Indeed, a tenant-landlord pair must have overlapping time windows in order to have an edge between them in the above example. We further assume without loss of generality that algorithms are lazy in the sense that they only make decisions on the deadlines of the vertices. On the deadline of a vertex v , it might be the case that v has already been matched to another vertex u on u ’s deadline. Otherwise, the algorithm must irrevocably either match v to one of its unmatched neighbors, or leave it unmatched.

¹Tenants may also specify their earliest move-in dates. This is omitted in our model because, from the algorithmic point of view, it is equivalent to having each tenant arrive on the earliest move-in date. The same applies to landlords.

The previous one-sided online model is a special case in which all offline vertices arrive at the beginning and have deadlines at the end, and each online vertex has its deadline right after arrival.

Further, there are many applications for which the underlying graph is not necessarily bipartite. Consider the following example:

Example (Ride-sharing Platform). DiDi is a major ride-sharing platform in China, handling tens of millions of rides on a daily basis. Requests are submitted to the platform in an online fashion. Each request is active in the system for a few minutes. The platform may match a pair of requests and serve them with the same taxi (or self-employed driver), provided that the pick-up locations and destinations are compatible, and their active time windows overlap. Requests can be modeled as vertices in a general graph and the compatibilities of pairs of requests can be modeled as edges.

Our model generalizes straightforwardly to general graphs by removing the bipartite assumption on the underlying graph. We refer to the generalization as **Fully Online Matching**.

It is easy to check that the naïve greedy algorithm that simply matches a vertex to an arbitrary unmatched neighbor remains to be 0.5-competitive. Can we do better?

1.1 Our Results and Techniques

We consider a natural generalization of the **Ranking** algorithm that picks a random permutation over all vertices at the beginning, and matches each vertex (if unmatched at its deadline) to the first unmatched neighbor according to the permutation. This algorithm can be implemented in our fully online model following the interpretation of **Ranking** by Devanur et al. [DJK13]: On the arrival of $v \in V$, the rank of vertex v , denoted by y_v , is chosen uniformly at random from $[0, 1]$; each vertex (if unmatched at its deadline) is matched to its unmatched neighbor with the highest, i.e., smallest, rank. We show that the **Ranking** algorithm is strictly better than 0.5-competitive:

Theorem 1.1 *Ranking is 0.5211-competitive for Fully Online Matching.*

Theorem 1.2 *Ranking is 0.5541-competitive for Fully Online Bipartite Matching.*

To our knowledge, our result for the **Fully Online Matching** problem is the first generalization of **Ranking** that achieves a competitive ratio strictly better than 0.5 in an online matching model that allows general graphs, making the first step towards providing a positive answer to the open question of whether **Ranking** is optimal for general graphs by [KVV90].

Our Techniques (Bipartite Case). We build on the randomized primal dual technique introduced by Devanur et al. [DJK13]. It can be viewed as a charging argument for sharing the gain of each matched edge between its two endpoints. Whenever an edge (u, v) is added to the matching, where v is an offline vertex and u is an online vertex, imagine a total gain of 1 being shared between u and v based on the rank of the offline vertex v . The higher the rank of v , the smaller share it gets. For **Online Bipartite Matching**, Devanur et al. [DJK13] introduced a gain sharing method such that, for any edge (u, v) and for any fixed ranks of offline vertices other than v , the expected gains of u and v (from all of their incident edges) combined is at least $1 - \frac{1}{e}$ over the randomness of v 's rank. This implies the $1 - \frac{1}{e}$ competitive ratio.

Next, consider an edge (u, v) in our model. Suppose u is the one with an earlier deadline. Since algorithms are lazy, the edge can only be added into the matching as a result of u 's decision at its deadline. In this sense, u plays a similar role as the online vertex and v as the offline vertex in the

analysis of Devanur et al. [DJK13]. Hence, a natural attempt is to consider the expected gains of u and v combined in the charging argument over the randomness of v 's rank alone.

To explain why the above approach fails, we first introduce the notions of active and passive vertices. We say a vertex u matches actively (or is active) if it is added to the matching by the algorithm at u 's deadline; other matched vertices are passive. All previous analyses [KVV90, AGKM11, DJK13, CCWZ14, ACC⁺16] crucially rely on a structural property that whenever vertex v is unmatched, its neighbor u must be matched to some other vertex with rank higher than v . In our model, however, this holds only if v 's neighbor u is active.

One may try to resolve this issue with a global amortized argument. If we go over all the edges in the graph, it cannot be the case that the endpoint with an earlier deadline of every edge always matches passively. After all, the numbers of active and passive vertices are equal. Interestingly, we instantiate this intuition with a local amortized argument by taking expectation over the randomness of u 's rank as well. Recall that u is the vertex with earlier deadline and, thus, plays a similar role as the online vertex in the argument of Devanur et al. [DJK13]. Taking expectation over u 's rank can be viewed as amortizing the case when u 's rank is low (active) and the case when u 's rank is high (passive).

Our Techniques (General Case). Moving from bipartite graphs to general graphs takes away another crucial structural property that the previous arguments rely on. In a bipartite graph, if a vertex u is matched by the Ranking algorithm for a realization of ranks while one of its neighbors v is not, then u remains matched no matter how the rank of v changes. In a general graph, however, it is possible that u becomes unmatched when v gets a higher rank.

We introduce a novel charging mechanic on top of the gain sharing rule used in the bipartite case. After a matching has been chosen by Ranking, for each active vertex w , consider an alternative run of Ranking with the same ranks but with w removed from the graph. The difference between the two matchings will be an alternating path and, thus, at most one vertex v would change from unmatched to matched in the absence of w . We shall refer to such a vertex v as the *victim* of w . Note that each active vertex has at most one victim, but an unmatched vertex could be the victim of many vertices. If the victim of w turns out to be its neighbor², our new charging mechanic will have w send to v a portion of w 's share from its incident edge in the matching, which we shall refer to as the *compensation* from w to v . Further, we show that whenever the aforementioned structural property fails, namely, some vertex u becomes unmatched when its unmatched neighbor v gets a higher rank, we can always identify a unique neighbor w of v that sends a compensation to v to remedy the loss in the charging argument.

Putting together the gain sharing mechanic from the bipartite case and the new mechanic of compensations, we can prove that for any edge (u, v) , the expected net gains of u and v combined is strictly greater than 0.5 over the randomness of the ranks of both u and v .

To our knowledge, this is the first charging mechanic that allows a vertex other than the two endpoints of a matched edge to get a share. We believe this novel charging mechanic will find further applications in other matching problems that consider general graphs.

Hardness Results. We complement our competitive analysis with two hardness results. The first hardness applies to arbitrary online algorithms, showing a separation between the best competitive ratio in our fully online model and the optimal ratio of $1 - \frac{1}{e} \approx 0.6321$ in the existing one-sided online model. The second hardness focuses on the Ranking algorithm, certifying that our analysis for the bipartite case is close to the best possible.

²This would induce an odd cycle and, thus, can only happen in non-bipartite graphs.

Theorem 1.3 *No randomized algorithm can achieve a competitive ratio better than 0.6317 for Fully Online Bipartite Matching.*

Theorem 1.4 *Ranking is at most 0.5671-competitive for Fully Online Bipartite Matching.*

1.2 Other Related Works

An alternative generalization of **Ranking** to general graphs has been considered for the problem of oblivious matching [CCWZ14, ACC⁺16]: Pick a permutation of vertices uniformly at random; then, go over the vertices one by one according to the permutation; for each unmatched vertex, match it to the first unmatched neighbor according to the same permutation. Chan et al. [CCWZ14] showed that it is a 0.523-approximation algorithm, improving the previous $(\frac{1}{2} + \frac{1}{400000})$ -approximation by a greedy algorithm [ADFS95]. Abolhassani et al. [ACC⁺16] improved the ratio to 0.526 and analyzed the weighted case. We stress that the alternative generalization is an offline algorithm because it needs to consider the vertices in random order, while our generalization is online. Nevertheless, our result for general graphs can be viewed as a 0.5211-approximation in the oblivious matching problem. We believe the new algorithm and analysis in this paper, in particular, the new charging mechanic of compensations, will find further applications in oblivious matching and other matching problems that consider general graphs.

Another online matching model in the literature considers online edge arrivals, upon which the algorithm must immediately decide whether to add the edge to the matching. McGregor [McG05] gave a deterministic $\frac{1}{3+2\sqrt{2}} \approx 0.1715$ -competitive algorithm in the edge-weighted preemptive setting. This ratio is later shown to be tight for deterministic algorithms [Var11]. Epstein et al. [ELSW13] designed a $\frac{1}{5.356} \approx 0.1867$ -competitive randomized algorithm and proved a hardness of $\frac{1}{1+\ln 2} \approx 0.59$. Chiplunkar et al. [CTV15] considered a restricted setting where the input graph is an unweighted growing tree and gave a $\frac{15}{28}$ -competitive algorithm. Finally, Buchbinder et al. [BST17] introduced an optimal $\frac{5}{9}$ -competitive algorithm for unweighted forests.

Wang and Wong [WW15] considered a more restrictive model of online bipartite matching with both sides of vertices arriving online: A vertex can only actively match other vertices at its arrival; if it fails to match at its arrival, it may still get matched passively by other vertices later. They showed a 0.526-competitive algorithm for a fractional version of the problem. We argue that the model in this paper better captures our aforementioned motivating applications.

2 Preliminaries

We consider the standard competitive analysis against an oblivious adversary. The competitive ratio of an algorithm is the ratio between the expected size of the matching by the algorithm over its own random bits to the size of the maximum matching of the underlying graph in hindsight. The adversary must in advance choose an instance, i.e., the underlying graph as well as arrivals and deadlines of vertices, without observing the random bits used by the algorithm. Otherwise, no algorithm can get any competitive ratio better than 0.5.

2.1 Ranking Algorithm and Some Basic Properties

See Algorithm 1 for a formal definition of **Ranking** in our model. Let $M(\vec{y})$ denote the matching produced when **Ranking** is run with \vec{y} as the ranks.

Algorithm 1 The Ranking Algorithm

- (1) a vertex v arrives:
pick $y_v \in [0, 1)$ uniformly at random.
 - (2) a vertex v 's deadline reaches:
if v is unmatched,
let $N(v)$ be the set of unmatched neighbors of v .
if $N(v) = \emptyset$, then v remains unmatched;
else match v to $\arg \min_{u \in N(v)} y_u$.
-

Recall the following definition of active/passive vertices. In the one-sided online model, only online vertices can be active and only offline vertices can be passive. In our fully online model, however, vertices can in general be of either types depending on the random ranks of the vertices.

Definition 2.1 (Active, Passive) *For any edge (u, v) added to the matching by Ranking at u 's deadline, we say that u is active and v is passive.*

The proofs of the following lemmas are deferred to Appendix A. The first lemma is a variant of the monotonicity property in previous works, incorporating the notions of active and passive vertices in our fully online model.

Lemma 2.1 (Monotonicity) *For any rank vector \vec{y} and any vertex u , we have*

1. if u is active/unmatched, then $M(\vec{y})$ remains the same when y_u increases;
2. if u is passive, then u remains passive when y_u decreases.

Let $\vec{y}_{-u} \in [0, 1)^{V \setminus \{u\}}$ be the ranks of all vertices but u , i.e., \vec{y}_{-u} is obtained by removing the u -th entry in \vec{y} . Let $M(\vec{y}_{-u})$ denote the matching produced by Ranking on $G - \{u\}$, i.e., the subgraph with vertex u removed, with \vec{y}_{-u} as the ranks.

For ease of notation, for any $y \in [0, 1]$, we use y^- to denote a value that is arbitrarily close to, but smaller than y . For example, our arguments consider functions discontinuous at 1 and use $f(1^-)$ to denote the limit of $f(x)$ as x goes to 1 from below. We also consider the matching w.r.t. ranks $(y_u = \theta, \vec{y}_{-u})$ to avoid confusions in marginal cases where ranks $(y_u = \theta, \vec{y}_{-u})$ need tie-breaking.

By Lemma 2.1, we can uniquely define the following marginal rank for every vertex.

Definition 2.2 (Marginal Rank) *For any u and any ranks \vec{y}_{-u} of other vertices, the marginal rank θ of u with respect to \vec{y}_{-u} is the largest value such that u is passive in $M(y_u = \theta^-, \vec{y}_{-u})$.*

Note that a vertex may still match another vertex (actively) when its rank is below the marginal rank in our fully online model. Nevertheless, it is consistent with the previous definition in the one-sided online model that concerns offline vertices, which cannot match actively.

Lemma 2.2 (Unmatched Neighbor) *Suppose v has marginal rank $\theta < 1$ with respect to \vec{y}_{-v} . Then, for any neighbor u of v that has an earlier deadline than v , and for any rank vector $(y_v = y, \vec{y}_{-v})$ with $y \in [\theta, 1)$, u either is passive, or actively matches a vertex with rank at most θ .*

It is well known that removing a matched vertex from the graph results in an alternating path in the matching produced by Ranking. The next lemma provides a more fine-grained characterization.

Lemma 2.3 (Alternating Path) *If u is matched in \vec{y} , then the symmetric difference between the matchings $M(\vec{y})$ and $M(\vec{y}_{-u})$ is an alternating path (u_0, u_1, \dots, u_l) with $u_0 = u$ such that*

1. for all even $i < l$, we have $(u_i, u_{i+1}) \in M(\vec{y})$; for all odd $i < l$, we have $(u_i, u_{i+1}) \in M(\vec{y}_{-u})$;
2. from $M(\vec{y})$ to $M(\vec{y}_{-u})$, vertices $\{u_1, u_3, \dots\}$ get worse, vertices $\{u_2, u_4, \dots\}$ get better.

Here, passive is better than active, which is in turns better than unmatched. Conditioned on being passive, matching to a vertex with earlier deadline is better. Conditioned on being active, matching to a vertex with higher rank is better.

2.2 Randomized Dual Fitting

Consider the following linear program relaxation of the matching problem and its dual.

$$\begin{array}{ll}
\max : & \sum_{(u,v) \in E} x_{uv} \\
\text{s.t.} & \sum_{v:(u,v) \in E} x_{uv} \leq 1 \quad \forall u \in V \\
& x_{uv} \geq 0 \quad \forall (u,v) \in E
\end{array}
\qquad
\begin{array}{ll}
\min : & \sum_{u \in V} \alpha_u \\
\text{s.t.} & \alpha_u + \alpha_v \geq 1 \quad \forall (u,v) \in E \\
& \alpha_u \geq 0 \quad \forall u \in V
\end{array}$$

It is known that the above linear program relaxation is integral for bipartite graphs, but it has a large integrality gap for general graphs (e.g., a complete graph of 3 vertices). Interestingly, this relaxation is sufficient for proving our positive results, even for general graphs.

Our approach builds on the randomized primal dual technique by Devanur et al. [DJK13]. We believe it is more appropriate to call our analysis (for general graphs) randomized dual fitting, however, because it relies on an extra phase of adjustments to the dual variables at the end that requires full knowledge of the instance.

Randomized Dual Fitting. We set the primal variables according to the matching by Ranking, which ensures primal feasibility, and set the dual variables such that the dual objective equals the primal objective. The dual assignment can be viewed as splitting the gain of 1 of every matched edge among the vertices; the dual variable α_v for every vertex v is equal to the total share it gets from all matched edges. Given primal feasibility and equal objectives, the usual primal dual and dual fitting techniques would further seek to show approximate dual feasibility, namely, $\alpha_u + \alpha_v \geq F$ for every edge (u, v) where F is the target competitive ratio. This is where the usual techniques fail and the smart insight by Devanur et al. [DJK13] comes to help. Due to the intrinsic randomness of Ranking, the above primal and dual assignments are themselves random variables. Devanur et al. [DJK13] observe that it suffices to have approximate dual feasibility in expectation. For completeness, we formulate this insight as the following lemma and include a proof in Appendix A.

Lemma 2.4 *Ranking is F -competitive if we can set (non-negative) dual variables such that*

1. $\sum_{(u,v) \in E} x_{uv} = \sum_{u \in V} \alpha_u$; and
2. $\mathbb{E}_{\vec{y}}[\alpha_u + \alpha_v] \geq F$ for all $(u, v) \in E$.

3 Bipartite Graphs: A Warm-up

Dual Assignment. We adopt the dual assignment by Devanur et al. [DJK13] and share the gain of each matched edge between its two endpoints as follows:

- *Gain Sharing:* Whenever an edge (u, v) is added to the matching with u active and v passive, let $\alpha_u = 1 - g(y_v)$ and $\alpha_v = g(y_v)$. Here, $g : [0, 1] \rightarrow [0, 1]$ is non-decreasing with $g(1) = 1$.

Randomized Primal Dual Analysis. The previous analysis of Ranking for Online Bipartite Matching relies on a structural property that for any edge (u, v) and any ranks \vec{y}_{-v} , u matches a vertex with rank no larger than v 's marginal rank regardless of v 's rank (e.g. Lemma 2.3 of [DJK13]). However, in our fully online setting, the same property holds only when u is active. By introducing the notions of passive and active vertices, we show the following weaker version of the property. It complements the basic property when y_u is larger than the marginal rank (Lemma 2.2).

Lemma 3.1 *Suppose v has marginal rank $\theta < 1$ with respect to \vec{y}_{-v} . Then, for any neighbor u of v that has an earlier deadline than v , and for any rank vector $(y_v = y, \vec{y}_{-v})$ with $y \in [0, \theta)$, u either is passive, or matches actively to a vertex with rank at most θ .*

Proof: We consider the matchings in 3 sets of ranks $\vec{y} = (y_v = y, \vec{y}_{-v})$, \vec{y}_{-v} and $\vec{y}_\theta = (y_v = \theta, \vec{y}_{-v})$. First, we show that u matches the same neighbor in $M(\vec{y}_\theta)$ and $M(\vec{y}_{-v})$. Since v is unmatched or active in $M(\vec{y}_\theta)$, removing v cannot affect vertices with earlier deadlines. In particular, u would match the same neighbor.

Consider the alternating path from $M(\vec{y})$ to $M(\vec{y}_{-v})$. If u is not in the alternating path, then u matches the same neighbor in all $M(\vec{y})$, $M(\vec{y}_{-v})$ and $M(\vec{y}_\theta)$. Otherwise, u appears in the alternating path with an odd distance from v since the graph is bipartite. Hence, by Lemma 2.3, u is better in $M(\vec{y})$ than in $M(\vec{y}_{-v})$ and, thus, is better than in $M(\vec{y}_\theta)$. In both cases, u is passive or actively matches a vertex with rank $\leq \theta$ in $M(\vec{y})$, since this holds for u in $M(\vec{y}_\theta)$ (by Lemma 2.2). ■

Recall that for any edge (u, v) we will consider the expected gain of α_u and α_v combined over the randomness of the ranks of both u and v . First, let us fix the rank of u , the vertex with an earlier deadline, and consider the expected gain over the randomness of v 's rank alone.

Lemma 3.2 *For any neighbor u of v that has an earlier deadline than v , and for any \vec{y}_{-v} , we have*

$$\mathbf{E}_{y_v} [\alpha_u + \alpha_v] \geq f(y_u) \stackrel{\text{def}}{=} \min_{\theta \in [0, 1]} \left\{ \int_0^\theta g(y_v) dy_v + \min\{1 - g(\theta), g(y_u)\} \right\}.$$

Proof: Let θ be the marginal rank of v with respect to ranks \vec{y}_{-v} . By definition, v is passive and gets $g(y_v)$ when $y_v < \theta$, i.e. $\mathbf{E}_{y_v} [\alpha_v \cdot \mathbb{1}(y_v < \theta)] = \int_0^\theta g(y_v) dy_v$. By Lemma 2.2 and 3.1, $\mathbf{E}_{y_v} [\alpha_u] \geq \min\{1 - g(\theta), g(y_u)\}$. Adding them together and taking the minimum over all possible θ 's concludes the statement. ■

It is worthwhile to make a comparison with a similar claim in the previous analysis by Devanur et al. [DJK13] for Online Bipartite Matching:

$$\mathbf{E}_{y_v} [\alpha_u + \alpha_v] \geq \min_{\theta \in [0, 1]} \left\{ \int_0^\theta g(y_v) dy_v + 1 - g(\theta) \right\},$$

where u is an online vertex and v is an offline vertex. As we have discussed in the introduction, for every edge in our model, the endpoint with an earlier deadline plays a similar role as the online vertex in the previous one-sided online model since the edge can only be added to the matching as a result of this endpoint's matching decision. In this sense, the bounds are indeed very similar, except for the last term, where the previous bound simply has $1 - g(\theta)$ while our bound has the smaller of $1 - g(\theta)$ and $g(y_u)$.

We interpret Lemma 3.2 as follows. It recovers the previous bound when the rank of u is large ($1 - g(\theta) \leq g(y_u)$), which roughly corresponds to the case when u is active (or unmatched) and the previous structural property holds. When the rank of u is small ($1 - g(\theta) > g(y_u)$), which roughly corresponds to the case when u is passive, it still provides some weaker lower bound on the expected gains of the two endpoints. The weaker bound, however, is at most 0.5 in the worst

case: the RHS becomes $\min\{1 - g(0), g(0)\} \leq 0.5$ for $\theta = y_u = 0$. Hence, it is crucial that we take expectation over the randomness of u 's rank as well, effectively amortizing the cases when u is active and when it is passive. This idea carries over to general graphs.

Proof of Theorem 1.2: Let $g(x) = e^{x-1}$, where $x \in [0, 1]$. For all $(u, v) \in E$, by Lemma 3.2,

$$\mathbf{E}_{\vec{y}}[\alpha_u + \alpha_v] = \mathbf{E}_{\vec{y}_{-v}}[\mathbf{E}_{y_v}[\alpha_u + \alpha_v]] \geq \mathbf{E}_{\vec{y}_{-v}}[f(y_u)] = \int_0^1 f(y_u) dy_u.$$

Observe that

$$\int_0^\theta g(z) dz + 1 - g(\theta) = e^{\theta-1} - \frac{1}{e} + 1 - e^{\theta-1} = 1 - \frac{1}{e}$$

and $\int_0^\theta g(z) dz + g(x) \geq g(x)$ for all $\theta \in [0, 1]$. We have $f(x) \geq \min\{g(x), 1 - \frac{1}{e}\}$, which implies (let $\theta = \ln(e-1)$ s.t. $g(\theta) = 1 - \frac{1}{e}$)

$$\int_0^1 f(x) dx \geq \int_0^\theta e^{x-1} dx + (1 - \theta)(1 - \frac{1}{e}) = \frac{e-2}{e} + (1 - \ln(e-1))(1 - \frac{1}{e}) \approx 0.55418.$$

By Lemma 2.4, we conclude that Ranking is at least 0.5541-competitive. \blacksquare

We are aware of a different function $g(y) = \min\{1, e^{y-1} + 0.0128\}$ that gives a (very slightly) better competitive ratio 0.5547. For convenience of presentation we only fix a simple form here.

4 General Graphs: An Overview

Dual Assignment. Moving from bipartite graphs to general graphs, even the weaker version of the structural property, i.e., Lemma 3.1, ceases to hold. Consider an edge (u, v) with u 's deadline being earlier. It is possible that decreasing y_v leads to a change of u 's status from matched to unmatched in a non-bipartite graph. As a result, the simple gain sharing rule in the previous analysis on the bipartite case no longer gives any bound strictly better than 0.5.

To handle general graphs, we design a novel charging mechanic on top of the gain sharing rule between the endpoints of matched edges. First, we introduce the following notion of *victim*.

Definition 4.1 (Victim) For any ranks \vec{y} and any active vertex w , v is w 's victim if

- v is an unmatched neighbor of w ;
- v is matched in $M(\vec{y}_{-w})$.

Observe that removing w results in an alternating path (Lemma 2.3) and, thus, at most one vertex changes from unmatched to matched. Hence, each active vertex has at most one victim.

Consider the following two-step approach for computing a dual assignment:

- *Gain Sharing:* Whenever an edge (u, v) is added to the matching with u active and v passive, let $\alpha_u = 1 - g(y_v)$ and $\alpha_v = g(y_v)$. Here, $g : [0, 1] \rightarrow [0, 1]$ is non-decreasing with $g(1) = 1$.
- *Compensation:* For every active vertex u that has a victim z , suppose u is matched to v . Decrease α_u and increase α_z by the same amount $h(y_v)$, where $h : [0, 1] \rightarrow [0, 1]$ is non-decreasing in $[0, 1)$, $h(y)/y$ is non-increasing, $h(1) = 0$ and $1 - g(y) - h(y) \geq 0$ for all y .

Note that the second step, in particular, identifying the victims of active vertices, can only be done after the entire instance has been revealed.

Each matched vertex will gain only from its incident matched edge. If it is further active and has a victim, it needs to send a compensation to the victim. Further, the active vertex can always afford the compensation from its gain since $1 - g(y) - h(y) \geq 0$ for all $y \in [0, 1)$. The monotonicity of $h(y)/y$ is for technical reasons in the analysis. Finally, note that an unmatched vertex may receive compensations from any number of active vertices.

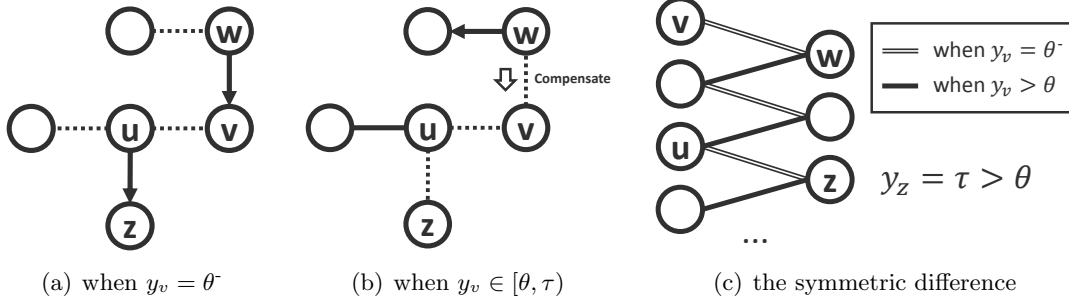


Figure 4.1: We use solid line to represent an edge in the matching, where the direction (if any) is from active vertex to passive vertex. (a) when v is higher than its marginal rank, u matches a vertex with rank $y_z = \tau > \theta$; (b) v is unmatched and compensated by w when it is lower than its marginal rank; u either is passive, or matches actively some vertex with rank higher than τ since v is unmatched; (c) the symmetric difference between the two matchings: an alternating path triggered when y_v increases to be larger than its marginal rank.

Randomized Dual Fitting Analysis. The main technical lemma is to establish a lower bound for $\mathbf{E}_{y_v} [\alpha_u + \alpha_v]$, as we have done in Lemma 3.2 for bipartite graphs. Due to space constraint, we present the analysis for a special case with following assumptions (θ is the marginal rank of v):

- v is unmatched in $M(y_v = y, \vec{y}_{-v})$ for all $y \geq \theta$;
- u actively matches the same vertex z with rank $y_z = \tau > \theta$ in $M(y_v = y, \vec{y}_{-v})$ for all $y < \theta$.

In other words, any rank of v higher than its marginal rank leads to the same (worse) situation for u , i.e. matching a vertex with rank $\tau \in (\theta, 1)$. See Figure 4.1 for an illustrative example.

In general, we need to also consider the case that v is active when its rank is lower than the marginal θ , and the possibility that u 's matching status may change multiple times as the rank of v gets higher. See Appendix B.1 for the analysis without the simplifying assumptions.

Subject to the above simplifying assumptions, we show the following:

Lemma 4.1 *For any neighbor u of v that has an earlier deadline than v , and for any \vec{y}_{-v} , we have*

$$\begin{aligned} \mathbf{E}_{y_v} [\alpha_v + \alpha_u] &\geq \int_0^\theta g(y_v) dy_v + (\tau - \theta) \cdot h(\theta) + \theta \cdot (1 - g(\tau) - h(\tau)) \\ &\quad + (1 - \theta) \cdot \min \{g(y_u), 1 - g(\theta) - h(\theta)\} . \end{aligned}$$

Suppose w matches actively to v in $M(y_v = \theta, \vec{y}_{-v})$ (refer to Figure 4.1(a)), that is, it is the first vertex after v in the alternating path when v 's rank moves below its marginal rank (refer to Figure 4.1(c)). We show in the following lemma that v receives a compensation from w whenever its rank y_v is between θ and τ (refer to Figure 4.1(b)).

Lemma 4.2 *For any $y \in [\theta, \tau)$, v is the victim of w in $M(y_v = y, \vec{y}_{-v})$.*

Proof: Let $\vec{y}_1 = (y_v = y, \vec{y}_{-v})$, where $y \in [\theta, \tau)$. By our assumption, v is unmatched and, thus, is an unmatched neighbor of w in $M(\vec{y}_1)$. To prove that v is the victim of w , we need to show that (1) w is active in $M(\vec{y}_1)$ and (2) v becomes matched when we remove w from the graph.

Consider $\vec{y}_2 = (y_v = \theta, \vec{y}_{-v})$. By our assumptions, v is passively matched to w and u actively matches z with $y_z = \tau$ in $M(\vec{y}_2)$. For this to happen, w must have an earlier deadline than v and none of matching decisions before w 's deadline pick w or v . Then, lowering v 's rank would not affect these decisions before w 's deadline and, thus, w must also be active in $M(\vec{y}_1)$.

Finally, consider what happens when w is removed from the graph. It triggers a portion of the alternating path (i.e., Figure 4.1(c)), the symmetric difference between $M(\vec{y}_1)$ and $M(\vec{y}_2)$. The portion starts from w (exclusive) and ends the first time when v becomes relevant, i.e., a vertex in the alternating path decides to pick v instead of the next vertex in the path. Further, we know for sure that v will be relevant at some point because otherwise u is in the path and the next vertex z has rank $\tau > y$. Therefore, v must be matched when w is removed from the graph. ■

The next two lemmas give lower bounds on the expected gain of α_v and α_u , respectively, over the randomness of v 's rank alone.

Lemma 4.3 $\mathbf{E}_{y_v} [\alpha_v] \geq \int_0^\theta g(y_v) dy_v + (\tau - \theta) \cdot h(\theta)$.

Proof: By definition, when $y_v < \theta$, v is passive and hence $\alpha_v = g(y_v)$.

Since w matches v actively in $M(y_v = \theta, \vec{y}_v)$ but not in $M(y_v = \theta, \vec{y}_{-v})$, we know that w must match a vertex with rank θ in $M(y_v \geq \theta, \vec{y}_{-v})$. For $y_v \in [\theta, \tau]$, Lemma 4.2 implies that v is the victim of w and, thus, v receives a compensation $h(\theta)$ from w . To sum up, we have

$$\mathbf{E}_{y_v} [\alpha_v] \geq \mathbf{E}_{y_v} [\alpha_v \cdot \mathbb{1}(y_v < \theta)] + \mathbf{E}_{y_v} [\alpha_v \cdot \mathbb{1}(y_v \in [\theta, \tau])] \geq \int_0^\theta g(y_v) dy_v + (\tau - \theta)h(\theta),$$

as claimed. ■

Lemma 4.4 $\mathbf{E}_{y_v} [\alpha_u] \geq \theta \cdot (1 - g(\tau) - h(\tau)) + (1 - \theta) \cdot \min\{g(y_u), 1 - g(\theta) - h(\theta)\}$.

Proof: By assumption, u actively matches vertex z with rank $y_z = \tau$ when $y_v < \theta$. Thus, u gains $1 - g(\tau)$ during the gain sharing phase and gives away $h(\tau)$ to its victim (if any). Integrating y_v from 0 to θ gives the first term on the RHS.

For $y_v \geq \theta$, u either is passive, or actively matches a vertex with rank at most θ . In the first case, we have $\alpha_u = g(y_u)$. In the second case, we have $\alpha_u \geq 1 - g(\theta) - h(\theta)$, by the monotonicity of g, h . Integrating y_v from θ to 1 gives the second term on the RHS. ■

Summing up the inequalities in Lemma 4.3 and 4.4 proves Lemma 4.1.

Comparing the bounds of Lemma 4.3 and Lemma 4.4, the parameter τ presents a trade-off between the expected gains α_u and α_v of the two vertices. The larger τ is, the less u gets when v is above its marginal rank, e.g., $1 - g(\tau) - h(\tau)$, and the more v gets as compensations when it is below its marginal rank, e.g., $(\tau - \theta) \cdot h(\tau)$; and vice versa.

Charging Functions. If Lemma 4.1 holds unconditionally, it remains to show that there exists functions g, h (with desired properties) and constant $F > 0.5$ such that

$$\int_0^1 \min_{0 \leq \theta < \tau < 1} \left\{ \int_0^\theta g(y_v) dy_v + (\tau - \theta)h(\theta) + \theta \cdot (1 - g(\tau) - h(\tau)) \right. \\ \left. + (1 - \theta) \cdot \min\{g(y_u), 1 - g(\theta) - h(\theta)\} \right\} dy_u \geq F,$$

and to apply Lemma 2.4 to conclude that Ranking is F -competitive. The unconditional version of Lemma 4.1 turns out to give a more complicated bound due to the considerations of other cases. Nevertheless, we can use a linear program to optimize the ratio over fine-grained discretized versions of g and h . To give a rigorous proof, which is deferred to Appendix B.2, we approximate the solutions of the linear program with piecewise-linear g and h (with two segments). We conclude with our choice of g and h that Ranking is 0.5211-competitive.

References

- [ACC⁺16] Melika Abolhassani, T.-H. Hubert Chan, Fei Chen, Hossein Esfandiari, Mohammad-Taghi Hajiaghayi, Hamid Mahini, and Xiaowei Wu. Beating ratio 0.5 for weighted oblivious matching problems. In *ESA*, volume 57 of *LIPICs*, pages 3:1–3:18. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
- [ADFS95] Jonathan Aronson, Martin Dyer, Alan Frieze, and Stephen Suen. Randomized greedy matching. ii. *Random Struct. Algorithms*, 6(1):55–73, January 1995.
- [AGKM11] Gagan Aggarwal, Gagan Goel, Chinmay Karande, and Aranyak Mehta. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *SODA*, pages 1253–1264, 2011.
- [BJN07] Niv Buchbinder, Kamal Jain, and Joseph Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. In *ESA*, volume 4698 of *Lecture Notes in Computer Science*, pages 253–264. Springer, 2007.
- [BM08] Benjamin Birnbaum and Claire Mathieu. On-line bipartite matching made simple. *ACM SIGACT News*, 39(1):80–87, 2008.
- [BST17] Niv Buchbinder, Danny Segev, and Yevgeny Tkach. Online algorithms for maximum cardinality matching with edge arrivals. In *ESA*, volume 87 of *LIPICs*, pages 22:1–22:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.
- [CCWZ14] T.-H. Hubert Chan, Fei Chen, Xiaowei Wu, and Zhichao Zhao. Ranking on arbitrary graphs: Rematch via continuous lp with monotone and boundary condition constraints. In *SODA*, pages 1112–1122, 2014.
- [CTV15] Ashish Chiplunkar, Sumedh Tirodkar, and Sundar Vishwanathan. On randomized algorithms for matching in the online preemptive model. In *ESA*, volume 9294 of *Lecture Notes in Computer Science*, pages 325–336. Springer, 2015.
- [DJ12] Nikhil R. Devanur and Kamal Jain. Online matching with concave returns. In *STOC*, pages 137–144. ACM, 2012.
- [DJK13] Nikhil R. Devanur, Kamal Jain, and Robert D. Kleinberg. Randomized primal-dual analysis of RANKING for online bipartite matching. In *SODA*, pages 101–107. SIAM, 2013.
- [ELSW13] Leah Epstein, Asaf Levin, Danny Segev, and Oren Weimann. Improved bounds for online preemptive matching. In *STACS*, volume 20 of *LIPICs*, pages 389–399. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013.
- [GM08] Gagan Goel and Aranyak Mehta. Online budgeted matching in random input models with applications to adwords. In *SODA*, pages 982–991, 2008.
- [KMT11] Chinmay Karande, Aranyak Mehta, and Pushkar Tripathi. Online bipartite matching with unknown distributions. In *STOC*, pages 587–596, 2011.
- [KVV90] Richard M. Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An optimal algorithm for on-line bipartite matching. In *STOC*, pages 352–358, 1990.

- [McG05] Andrew McGregor. Finding graph matchings in data streams. In *APPROX-RANDOM*, volume 3624 of *Lecture Notes in Computer Science*, pages 170–181. Springer, 2005.
- [MSVV07] Aranyak Mehta, Amin Saberi, Umesh V. Vazirani, and Vijay V. Vazirani. Adwords and generalized online matching. *J. ACM*, 54(5):22, 2007.
- [MY11] Mohammad Mahdian and Qiqi Yan. Online bipartite matching with random arrivals: an approach based on strongly factor-revealing LPs. In *STOC*, pages 597–606, 2011.
- [Var11] Ashwinkumar Badanidiyuru Varadaraja. Buyback problem - approximate matroid intersection with cancellation costs. In *ICALP (1)*, volume 6755 of *Lecture Notes in Computer Science*, pages 379–390. Springer, 2011.
- [WW15] Yajun Wang and Sam Chiu-wai Wong. Two-sided online bipartite matching and vertex cover: Beating the greedy algorithm. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, pages 1070–1081, 2015.

A Missing Proofs in Section 2

Proof of Lemma 2.1: For the first statement, since u is active or unmatched, we know that for each neighbor v of u with an earlier deadline than u , v does not match u in $M(\vec{y})$ at their deadlines. Hence when y_u increases, they would make the same decision. In other words, when u 's deadline reaches, the partial matching produced is the same as before. As a consequence, the eventual matching produced would be identical, as u will actively match the same vertex as in $M(\vec{y})$.

The second statement is implied by the first statement. Suppose otherwise, e.g., u is active or unmatched when y_u is decreased from y to some $y' < y$. Then we know that by increasing y_u from y' to y , u becomes passive, which violates the first statement. ■

Proof of Lemma 2.2: Consider the matching $M(y_v = \theta, \vec{y}_{-v})$. By definition, v is either active or unmatched. Hence, at u 's deadline, which is earlier than v 's deadline, v is unmatched. Consequently, u either is passive or matches actively to some vertex w with $y_w \leq y_v = \theta$. By Lemma 2.1, there is no change in the matching when we increases y_u , which concludes the proof. ■

Proof of Lemma 2.3: We prove the lemma by mathematical induction on n , the total number of vertices. For the base case when $n = 2$, the symmetric difference is a single edge (u, u_1) and the second statement holds since u_1 is matched in $M(\vec{y})$ and unmatched in $M(\vec{y}_{-u})$.

Suppose the lemma holds for $1, 2, \dots, n-1$ and we consider the case when $|V| = n$.

Let u_1 be matched to u in $M(\vec{y})$. Observe that if we remove both u and u_1 from G (let $\vec{y}' \in [0, 1]^{V \setminus \{u, u_1\}}$ be the resulting vector), then we have $M(\vec{y}) = M(\vec{y}') \cup \{(u, u_1)\}$.

If u_1 is unmatched in $M(\vec{y}_{-u})$, then we have $M(\vec{y}_{-u}) = M(\vec{y}')$ and the lemma holds by induction hypothesis. Now suppose u_1 is matched in $M(\vec{y}_{-u})$.

By definition \vec{y}' is obtained by removing u_1 (which is matched in \vec{y}_{-u}) from \vec{y}_{-u} . By induction hypothesis, the symmetric difference between $M(\vec{y}_{-u})$ and $M(\vec{y}')$ is an alternating path (u_1, \dots, u_l) such that (a) for all odd $i < l$, we have $(u_i, u_{i+1}) \in M(\vec{y}_{-u})$; for all even $i < l$, we have $(u_i, u_{i+1}) \in M(\vec{y}')$; (b) from $M(\vec{y}_{-u})$ to $M(\vec{y}')$, vertices $\{u_2, u_4, \dots\}$ get worse, vertices $\{u_3, u_5, \dots\}$ get better.

Hence the symmetric difference between $M(\vec{y})$ and $M(\vec{y}_{-u})$ is the alternating path (u, u_1, \dots, u_l) (recall that $M(\vec{y}) = M(\vec{y}') \cup \{(u, u_1)\}$). It is easy to see that statement (a) holds, and statement (b) holds for vertices $\{u_2, \dots, u_l\}$.

Now consider vertex u_1 , which is matched to u in $M(\vec{y})$, and matched to u_2 in $M(\vec{y}_{-u})$.

If u_1 is passively matched (by u) in $M(\vec{y})$, then we know that u has an earlier deadline than u_1 . Hence in $M(\vec{y}_{-u})$, either u_1 is active, or passively matched by some u_2 with a deadline later than u . In other words, u_1 gets worse from $M(\vec{y})$ to $M(\vec{y}_{-u})$.

If u_1 matches u actively in $M(\vec{y})$, then we know that u_1 has an earlier deadline than u . Hence when u_1 is considered in \vec{y}_{-u} , the set of unmatched vertices (except for u) is identical as in $M(\vec{y})$. Consequently, u_1 actively matches some vertex u_2 with $y_{u_2} > y_u$ (otherwise u_1 will not match u in $M(\vec{y})$). In other words, u_1 gets worse from $M(\vec{y})$ to $M(\vec{y}_{-u})$. ■

Proof of Lemma 2.4: Let $\tilde{\alpha}_u := \mathbf{E}_{\vec{y}}[\alpha_u] / F$ for all $u \in V$. By the first assumption,

$$\sum_{u \in V} \tilde{\alpha}_u = \sum_{u \in V} \frac{\mathbf{E}_{\vec{y}}[\alpha_u]}{F} = \frac{1}{F} \mathbf{E}_{\vec{y}} \left[\sum_{u \in V} \alpha_u \right] = \frac{1}{F} \mathbf{E}_{\vec{y}} \left[\sum_{(u,v) \in E} x_{uv} \right].$$

Moreover, $\tilde{\alpha}$ is a feasible dual solution: by the second assumption, $\tilde{\alpha}_u + \tilde{\alpha}_v = \mathbf{E}_{\vec{y}}[\alpha_u + \alpha_v] / F \geq 1$

for all $(u, v) \in E$. By duality, we conclude that

$$\frac{1}{F} \mathbf{E}_{\vec{y}} \left[\sum_{(u,v) \in E} x_{uv} \right] = \sum_{u \in V} \tilde{\alpha}_u \geq \text{OPT},$$

where OPT is the optimal primal solution, which is at least the size of a maximum matching. \blacksquare

B Missing Proofs in Section 4

B.1 General Version of Lemma 4.1

In this section we prove the following lemma, which is a general version of Lemma 4.1 (without the simplifying assumptions on u, v we made in Section 4).

Lemma B.1 *For any neighbor u of v that has an earlier deadline than v , and for any \vec{y}_{-v} , we have*

$$\begin{aligned} \mathbf{E}_{y_v} [\alpha_u + \alpha_v] \geq f(y_u) \stackrel{\text{def}}{=} \min_{\theta} \left\{ \min_{\tau \in [\theta, 1]} \left\{ \int_0^{\theta} g(y_v) dy_v + \min \{ (1 - \theta)(1 - g(1^-) - h(1^-)), (\tau - \theta)h(\theta) \} \right. \right. \\ \left. \left. + (1 - \theta) \min \{ g(y_u), 1 - g(\theta) - h(\theta) \} + \theta \min \{ g(y_u), 1 - g(\tau) - h(\tau) \} \right\}, \right. \\ \left. \int_0^{\theta} g(y_v) dy_v + (1 - \theta) \min \{ 1 - g(1^-) - h(1^-), h(\theta) \} + (1 - \theta) \min \{ g(y_u), 1 - g(\theta) \} \right\}. \end{aligned}$$

Fix any neighbor u of v with an earlier deadline than v , and any \vec{y}_{-v} . Let θ be the marginal rank of v , i.e. v is passive only when $y_v < \theta$. By Lemma 2.2, we know that when $y_v \geq \theta$, u either is passive or actively matches some vertex with rank at most θ .

We define in the following two lists of thresholds $\{\theta_i\}_{i=0}^{m+1}$ and $\{\tau_i\}_{i=0}^{m+1}$ that captures the matching statuses of u when y_v is smaller than the marginal rank θ .

Imagine that we decrease y_v continuously starting from $y_v = \theta$. Let $\theta_0 = \theta$ and $\tau_0 = \theta$. We define $y_v = \theta_{i+1}$ to be the first moment after θ_i when u actively matches some vertex z_{i+1} with $y_{z_{i+1}} > \tau_i$. For convenience of description, we say that u actively matches a vertex with rank 1 if u is unmatched (by definition, the gain of α_u is 0 in both descriptions since $1 - g(1) = 0$). Define $\tau_{i+1} := y_{z_{i+1}}$. Let θ_m be the last non-zero threshold. For convenience, we define $\theta_{m+1} = 0$ and $\tau_{m+1} = 1$. By definition we have the following fact.

Fact B.1 *There exists a sequence of non-increasing thresholds $\{\theta_i\}_{i=0}^{m+1}$ and a sequence of non-decreasing thresholds $\{\tau_i\}_{i=0}^{m+1}$ such that*

1. for all $0 \leq i \leq m$ and $y \in [\theta_{i+1}, \theta_i)$, u is passive or actively matches some vertex with rank at most τ_i in $M(y_v = y, \vec{y}_{-v})$;
2. for all $1 \leq i \leq m$, u actively matches a vertex z_i with rank τ_i in $M(y_v = \theta_i, \vec{y}_{-v})$.

For all $i \in [m]$, let w_i be vertex that actively matches v in $M(y_v = \theta_i, \vec{y}_{-v})$.

Observe that all w_i 's must be different, e.g. the deadline of w_{i+1} must be earlier than w_i , in order for w_{i+1} to match v in $M(y_v = \theta_{i+1}, \vec{y}_{-v})$. Moreover, we know that in $M(y_v = \theta, \vec{y}_{-v})$, each w_i matches a vertex with rank θ_i , since w_i chooses v in $M(y_v = \theta_i, \vec{y}_{-v})$ but not in $M(y_v = \theta_i, \vec{y}_{-v})$.

Lemma B.2 *For all $y \in [\theta, \tau_i)$, if v is unmatched in $M(y_v = y, \vec{y}_{-v})$, then v is the victim of w_i .*

Proof: Let $\vec{y}_1 = (y_v = y, \vec{y}_{-v})$, where $y \in [\theta, \tau_i)$. Trivially, v is an unmatched neighbor of w_i in $M(\vec{y}_1)$. To prove that v is the victim of w_i , it suffices to show that w_i is active in $M(\vec{y}_1)$ and v becomes matched when we remove w_i from the graph.

Let $\vec{y}_2 = (y_v = \theta_i, \vec{y}_{-v})$. By definition, v is passively matched by w_i and u actively matches z_i with $y_{z_i} = \tau_i$ in $M(\vec{y}_2)$. It is easy to see that w_i is also active in $M(\vec{y}_1)$, as otherwise, w_i should still be passive in $M(\vec{y}_2)$ given that v does not affect any decisions before w_i 's deadline.

Let \vec{y}_3 be the ranks by removing the w_i -th entry from \vec{y}_1 . Assume for contrary that v is unmatched in $M(\vec{y}_3)$. Then we have $M(\vec{y}_3) = M(\vec{y}_2) \setminus \{(w_i, v)\}$. This implies that u actively matches z_i in $M(\vec{y}_3)$ while v (with rank $y_v < y_{z_i} = \tau_i$) is unmatched, which is a contradiction. ■

Equipped with Lemma B.2, we first give a lower bound on the expected gain of α_v . For notational convenience, we define a new function $\phi : [0, 1] \rightarrow [0, 1]$ such that $\phi(y) := 1 - g(y) - h(y)$. Recall by definition of g and h , ϕ is a non-increasing function with $\phi(1) = 0$.

Lemma B.3 $\mathbf{E}_{y_v} [\alpha_v] \geq \int_0^\theta g(y_v) dy_v + \min\{(1 - \theta)\phi(1^-), \sum_{i=0}^m (\tau_i - \theta)h(\theta_i)\}$.

Proof: By definition, v is passive when $y_v < \theta$. Hence we have $\mathbf{E}_{y_v} [\alpha_v \cdot \mathbb{1}(y_v < \theta)] = \int_0^\theta g(y_v) dy_v$, which corresponds to the first term on the RHS.

For all $y_v \geq \theta$, v is either active or unmatched. In the first case, let p be matched passively by v in $M(y_v, \vec{y}_{-v})$. We know that v gains $1 - g(y_p)$ during the gain sharing phase and gives away $h(y_p)$ to its victim (if any), which implies $\alpha_v \geq 1 - g(y_z) - h(y_z) = \phi(y_z) \geq \phi(1^-)$. Hence we have $\mathbf{E}_{y_v} [\alpha_v \cdot \mathbb{1}(y_v \geq \theta)] \geq (1 - \theta)\phi(1^-)$.

In the second case, by Lemma B.2, v is the victim of w_i when $y_v \in [\theta, \tau_i)$. Hence v gains $h(\theta_i)$ from w_i in the compensation phase (recall that w_i matches a vertex with rank θ_i when $y_v \in [\theta, \tau_i)$). Putting all compensation (from w_1, \dots, w_m) together, we get $\mathbf{E}_{y_v} [\alpha_v \cdot \mathbb{1}(y_v \geq \theta)] \geq \sum_{i=0}^m (\tau_i - \theta)h(\theta_i)$, which corresponds to the second term on the RHS. ■

Lemma B.4 $\mathbf{E}_{y_v} [\alpha_u \cdot \mathbb{1}(y_v < \theta)] \geq \sum_{i=0}^m (\theta_i - \theta_{i+1}) \min\{g(y_u), \phi(\tau_i)\}$.

Proof: We partition the interval $[0, \theta)$ into $m + 1$ segments: $[\theta_{i+1}, \theta_i)$, for $0 \leq i \leq m$. Fix any i , and consider $y_v \in [\theta_{i+1}, \theta_i)$. If u is passive in $M(y_v, \vec{y}_{-v})$, then we have $\alpha_u \geq g(y_u)$. Otherwise, we know that u actively matches a vertex z with $y_z \leq \tau_i$ (by Fact B.1). Hence u gains $1 - g(y_z)$ during the gain sharing phase and gives away $h(y_z)$ to its victim (if any), i.e., we have $\alpha_u \geq \phi(y_z) \geq \phi(\tau_i)$. Summing up the gain from the $m + 1$ segments concludes the proof. ■

Observe that for lower bounding $\mathbf{E}_{y_v} [\alpha_u + \alpha_v]$, we shall consider the total gain of $\alpha_u + \alpha_v$. We may omit the compensation from u to v , since it does not change the summation. For analysis convenience, we assume v is never a victim of u .

Lemma B.5 $\mathbf{E}_{y_v} [\alpha_u \cdot \mathbb{1}(y_v \geq \theta)] \geq (1 - \theta) \min\{g(y_u), \phi(\theta)\}$. Moreover, if $\tau_m = 1$, then we have $\mathbf{E}_{y_v} [\alpha_u \cdot \mathbb{1}(y_v \geq \theta)] \geq (1 - \theta) \min\{g(y_u), 1 - g(\theta)\}$.

Proof: For all $y_v \geq \theta$, u is either passive or actively matches a vertex with rank at most θ . Therefore, $\alpha_u \geq \min\{g(y_u), \phi(\theta)\}$. Integrating y_v from θ to 1 gives the first statement.

When $\tau_m = 1$, we know that u is unmatched when $y_v = \theta_m$. Fix any $\vec{y} = (y_v, \vec{y}_{-v})$, where $y_v \geq \theta$. We show that u does not have any unmatched neighbor other than v in $M(\vec{y})$, which implies that u does not have a victim and hence $\alpha_u \geq \min\{g(y_u), 1 - g(\theta)\}$.

Suppose otherwise, let $z \neq v$ be the unmatched neighbor of u in $M(\vec{y})$.

Let $\vec{y}_1 = (y_v = \theta_m, \vec{y}_{-v})$. We know that u is matched in $M(\vec{y})$ and unmatched in $M(\vec{y}_1)$. Consider the partial matchings produced right after u 's deadline when Ranking is run with \vec{y} and

\vec{y}_1 , respectively. We denote the matchings by $M_u(\vec{y})$ and $M_u(\vec{y}_1)$, respectively. It is easy to see that the symmetric difference between $M_u(\vec{y})$ and $M_u(\vec{y}_1)$ is an alternating path, with u being one endpoint. Observe that v is matched in $M_u(\vec{y}_1)$ (as u is unmatched), and is unmatched in $M_u(\vec{y})$ (as it is not passive in $M(\vec{y})$). Hence v is the other end point of the alternating path. Consequently, we know that z is unmatched in $M_u(\vec{y}_1)$ (as it is unmatched in $M_u(\vec{y})$), which is a contradiction as its neighbor u is also unmatched in $M_u(\vec{y}_1)$. ■

The next technical lemma shows that the worst case is achieved when there exists only one threshold $\theta_m = \theta$, i.e. u matches some vertex with rank $\tau > \theta$ in $M(y_v = \theta, \vec{y}_{-v})$, and matches a vertex with rank at most τ for all $y_v < \theta$.

Lemma B.6 *Given that $h(y)/y$ is a non-increasing function, we have*

$$\begin{aligned} & \min \left\{ (1 - \theta)\phi(1^-), \sum_{i=0}^m (\tau_i - \theta)h(\theta_i) \right\} + \sum_{i=0}^m (\theta_i - \theta_{i+1}) \min\{g(y_u), \phi(\tau_i)\} \\ & \geq \min_i \left\{ \min \{(1 - \theta)\phi(1^-), (\tau_i - \theta)h(\theta)\} + \theta \min \{g(y_u), \phi(\tau_i)\} \right\}. \end{aligned}$$

Proof: Consider the first term of LHS. If $(1 - \theta)\phi(1^-) < \sum_{i=0}^m (\tau_i - \theta)h(\theta_i)$, we have

$$\begin{aligned} \text{LHS} & \geq (1 - \theta)\phi(1^-) + \sum_{i=0}^m (\theta_i - \theta_{i+1}) \min \{g(y_u), \phi(\tau_m)\} \\ & = (1 - \theta)\phi(1^-) + \theta \min \{g(y_u), \phi(\tau_m)\} \geq \text{RHS}. \end{aligned}$$

If $(1 - \theta)\phi(1^-) \geq \sum_{i=0}^m (\tau_i - \theta)h(\theta_i)$, we have

$$\begin{aligned} \text{LHS} & = \sum_{i=0}^m (\tau_i - \theta)h(\theta_i) + \sum_{i=0}^m (\theta_i - \theta_{i+1}) \min\{g(y_u), \phi(\tau_i)\} \\ & = \sum_{i=0}^m (\theta_i - \theta_{i+1}) \left(\frac{\tau_i - \theta}{\theta_i - \theta_{i+1}} h(\theta_i) + \min\{g(y_u), \phi(\tau_i)\} \right) \\ & \geq \sum_{i=0}^m (\theta_i - \theta_{i+1}) \left(\frac{\tau_i - \theta}{\theta} h(\theta) + \min\{g(y_u), \phi(\tau_i)\} \right) \\ & \geq \sum_{i=0}^m (\theta_i - \theta_{i+1}) \min_j \left\{ \frac{\tau_j - \theta}{\theta} h(\theta) + \min\{g(y_u), \phi(\tau_j)\} \right\} \\ & = \min_j \left\{ (\tau_j - \theta)h(\theta) + \theta \min\{g(y_u), \phi(\tau_j)\} \right\} \geq \text{RHS}, \end{aligned}$$

where the first inequality follows from $\frac{\tau_i - \theta}{\theta_i - \theta_{i+1}} h(\theta_i) \geq (\tau_i - \theta) \frac{h(\theta_i)}{\theta_i} \geq (\tau_i - \theta) \frac{h(\theta)}{\theta}$. ■

Proof of Lemma B.1: By Lemma B.3 and B.4, we have

$$\begin{aligned} & \mathbf{E}_{y_v} [\alpha_v] + \mathbf{E}_{y_v} [\alpha_u \cdot \mathbb{1}(y_v < \theta)] \\ & \geq \int_0^\theta g(y_v) dy_v + \min \left\{ (1 - \theta)\phi(1^-), \sum_{i=0}^m (\tau_i - \theta)h(\theta_i) \right\} + \sum_{i=0}^m (\theta_i - \theta_{i+1}) \min\{g(y_u), \phi(\tau_i)\} \\ & \geq \int_0^\theta g(y_v) dy_v + \min_i \left\{ \min\{(1 - \theta)\phi(1^-), (\tau_i - \theta)h(\theta)\} + \theta \min\{g(y_u), \phi(\tau_i)\} \right\}, \end{aligned}$$

where the last inequality comes from Lemma B.6.

Combining with Lemma B.5 (which gives different lower bounds for $\mathbf{E}_{y_v} [\alpha_u \cdot \mathbb{1}(y_v \geq \theta)]$ depending on whether $\tau_m = 1$), we prove Lemma B.1 for two cases, depending on whether $\tau_m = 1$.

If $\tau_m < 1$, we have (recall that $\phi(\theta) = 1 - g(\theta) - h(\theta) \leq 1 - g(\theta)$)

$$\begin{aligned} \mathbf{E}_{y_v} [\alpha_u + \alpha_v] &= \mathbf{E}_{y_v} [\alpha_v] + \mathbf{E}_{y_v} [\alpha_u \cdot \mathbb{1}(y_v < \theta)] + \mathbf{E}_{y_v} [\alpha_u \cdot \mathbb{1}(y_v \geq \theta)] \\ &\geq \int_0^\theta g(y_v) dy_v + \min_{\tau \in [\theta, 1]} \{ \min\{(1 - \theta)\phi(1^-), (\tau - \theta)h(\theta)\} + \theta \min\{g(y_u), \phi(\tau)\} \} \\ &\quad + (1 - \theta) \min\{g(y_u), \phi(\theta)\}, \end{aligned}$$

which corresponds to the first term of the outer most min in the expression of Lemma B.1.

If $\tau_m = 1$, we have

$$\begin{aligned} \mathbf{E}_{y_v} [\alpha_u + \alpha_v] &\geq \int_0^\theta g(y_v) dy_v + \min_{\tau \in [\theta, 1]} \{ \min\{(1 - \theta)\phi(1^-), (\tau - \theta)h(\theta)\} + \theta \min\{g(y_u), \phi(\tau)\} \} \\ &\quad + (1 - \theta) \min\{g(y_u), 1 - g(\theta)\} \\ &\geq \min \left\{ \min_{\tau \in [\theta, 1]} \left\{ \int_0^\theta g(y_v) dy_v + \min\{(1 - \theta)\phi(1^-), (\tau - \theta)h(\theta)\} + \theta \min\{g(y_u), \phi(\tau)\} \right. \right. \\ &\quad \left. \left. + (1 - \theta) \min\{g(y_u), \phi(\theta)\} \right\}, \right. \\ &\quad \left. \int_0^\theta g(y_v) dy_v + (1 - \theta) \min\{\phi(1^-), h(\theta)\} + (1 - \theta) \min\{g(y_u), 1 - g(\theta)\} \right\}. \end{aligned}$$

Taking the minimum over all possible θ 's concludes the proof. ■

B.2 Lower Bound of the Competitive Ratio

Recall that

$$\begin{aligned} f(y_u) &\stackrel{\text{def}}{=} \min_{\theta} \left\{ \min_{\tau \in [\theta, 1]} \left\{ \int_0^\theta g(y_v) dy_v + \min\{(1 - \theta)(1 - g(1^-) - h(1^-)), (\tau - \theta)h(\theta)\} \right. \right. \\ &\quad \left. \left. + (1 - \theta) \min\{g(y_u), 1 - g(\theta) - h(\theta)\} + \theta \min\{g(y_u), 1 - g(\tau) - h(\tau)\} \right\}, \right. \\ &\quad \left. \int_0^\theta g(y_v) dy_v + (1 - \theta) \min\{1 - g(1^-) - h(1^-), h(\theta)\} + (1 - \theta) \min\{g(y_u), 1 - g(\theta)\} \right\}. \end{aligned}$$

We construct the functions g, h explicitly as follows (refer to Figure B.1).

$$g(x) = \begin{cases} k_g^1 x + b, & x \in [0, t] \\ k_g^2(x - t) + k_g^1 t + b, & x \in (t, 1) \\ 1 & x = 1 \end{cases}, \quad h(x) = \begin{cases} k_h^1 x, & x \in [0, t] \\ k_h^2(x - t) + k_h^1 t, & x \in (t, 1) \\ 0 & x = 1 \end{cases}$$

where $t = 0.3, k_g^1 = 0.21, k_g^2 = 0.1, b = 0.46, k_h^1 = 0.26, k_h^2 = 0.17$. It is easy to see that both g, h are non-decreasing in $[0, 1)$ and $h(x)/x$ is non-increasing.




Figure B.1: $g(x)$ and $h(x)$

We first simplify the expression of $f(y_u)$. Recall that we define $\phi(y) = 1 - g(y) - h(y)$, which is a decreasing function with $\phi(1) = 0$. By definition of g, h stated above, we have $\phi(1^-) = 0.21$.

Observe that $(1 - \theta)\phi(1^-) > (1 - \theta) \cdot 0.197 = (1 - \theta)h(1^-) \geq (\tau - \theta)h(\theta)$ for all τ, θ . Hence,

$$f(y_u) = \min \left\{ \min_{\theta \leq \tau < 1} \psi_1(y_u, \theta, \tau), \min_{\theta \leq 1} \psi_2(y_u, \theta) \right\},$$

where

$$\begin{aligned} \psi_1(y_u, \theta, \tau) &\stackrel{\text{def}}{=} \int_0^\theta g(y_v) dy_v + (\tau - \theta)h(\theta) + (1 - \theta) \min \{g(y_u), \phi(\theta)\} + \theta \min \{g(y_u), \phi(\tau)\}, \\ \psi_2(y_u, \theta) &\stackrel{\text{def}}{=} \int_0^\theta g(y_v) dy_v + (1 - \theta)h(\theta) + (1 - \theta) \min \{g(y_u), 1 - g(\theta)\}. \end{aligned}$$

The following lemma implies that Ranking is 0.5211-competitive.

Lemma B.7 $\int_0^1 f(y_u) dy_u > 0.5211$.

Proof: As we will show in Lemma B.8 and Lemma B.9 below, we have

$$f(y_u) = \min \left\{ \min_{\theta \leq \tau < 1} \psi_1(y_u, \theta, \tau), \min_{\theta \leq 1} \psi_2(y_u, \theta) \right\} \geq \min \{g(y_u), 0.5349\}.$$

Let y_u^* be such that $g(y_u^*) = 0.5349$, we have

$$\int_0^1 f(y_u) dy_u = \int_0^{y_u^*} f(y_u) dy_u + \int_{y_u^*}^1 f(y_u) dy_u = \int_0^{y_u^*} g(y_u) dy_u + (1 - y_u^*) \cdot 0.5349 > 0.5211. \quad \blacksquare$$

We first consider the easier one, ψ_2 .

Lemma B.8 For any $\theta \leq 1$, we have $\psi_2(y_u, \theta) > \min \{g(y_u), 0.5349\}$.

Proof: First, if $\theta = 1$, we have $\psi_2(y_u, \theta) = \int_0^1 g(y_v) dy_v \approx 0.5381 > 0.5349$. Now consider $\theta < 1$. If $g(y_u) < 1 - g(\theta)$, we have $\psi_2(y_u, \theta) = \int_0^\theta g(y_v) dy_v + (1 - \theta)h(\theta) + (1 - \theta)g(y_u)$. Thus

$$\begin{aligned} \frac{\partial \psi_2}{\partial \theta} &= g(\theta) + (1 - \theta)h'(\theta) - h(\theta) - g(y_u), \\ \frac{\partial^2 \psi_2}{\partial \theta^2} &= g'(\theta) - 2 \cdot h'(\theta) = \begin{cases} k_g^1 - 2k_h^2 = -0.31 & \theta < t \\ k_g^2 - 2k_h^2 = -0.24 & \theta > t \end{cases}. \end{aligned}$$

Hence the minimum of $\psi_2(y_u, \theta)$ is achieved at $\arg \min_{\theta < 1} \{\psi_2(y_u, \theta)\} \in \{0, t, 1\}$. Note that

$$\begin{aligned} \psi_2(y_u, 0) &= h(0) + g(y_u) \geq g(y_u), \\ \psi_2(y_u, t) &= \int_0^t g(y_v) dy_v + (1 - t)h(t) + (1 - t)g(y_u) > t \cdot 0.673 + (1 - t) \cdot g(y_u) > g(y_u), \\ \psi_2(y_u, 1^-) &= \int_0^1 g(y_v) dy_v \approx 0.5381 > 0.5349. \end{aligned}$$

We have $\min_{\theta < 1} \{\psi_2(y_u, \theta)\} \geq \min\{g(y_u), 0.5349\}$, as claimed.

If $g(y_u) \geq 1 - g(\theta)$, we have $\psi_2(y_u, \theta) = \int_0^\theta g(y_v) dy_v + (1 - \theta)h(\theta) + (1 - \theta)(1 - g(\theta))$. Thus

$$\begin{aligned} \frac{\partial \psi_2}{\partial \theta} &= g(\theta) + (1 - \theta)h'(\theta) - h(\theta) - (1 - g(\theta)) - (1 - \theta)g'(\theta), \\ \frac{\partial^2 \psi_2}{\partial \theta^2} &= 3 \cdot g'(\theta) - 2 \cdot h'(\theta) = \begin{cases} 0.11 & \theta < t \\ -0.04 & \theta > t \end{cases}. \end{aligned}$$

We calculate the zero point of $\frac{\partial \psi_2}{\partial \theta}$ in $[0, t)$, i.e., let $\frac{\partial \psi_2}{\partial \theta} = 0$, we have solution

$$\theta^* = \frac{1 + k_g^1 - k_h^1 - 2b}{3k_g^1 - 2k_h^1} \approx 0.273.$$

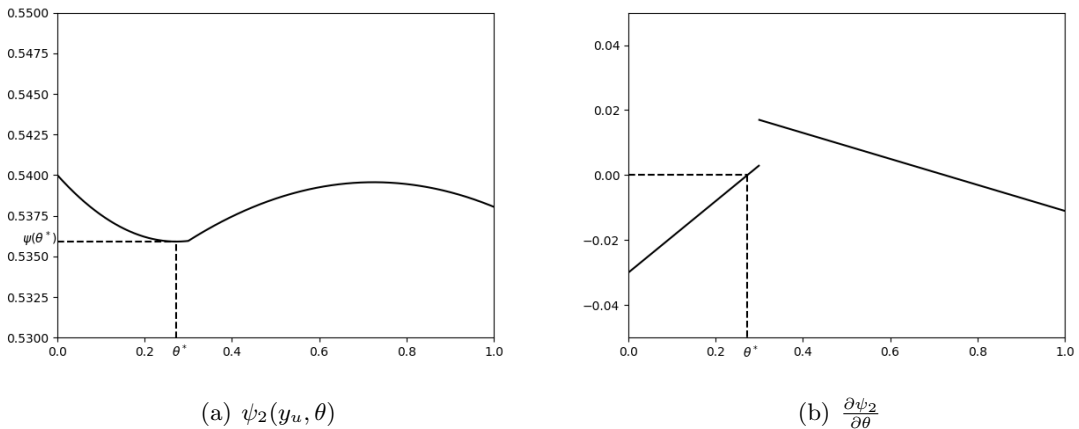


Figure B.2: $\psi_2(y_u, \theta)$ and $\frac{\partial \psi_2}{\partial \theta}$

Thus, for any fixed y_u , $\psi_2(y_u, \theta)$ is decreasing in $[0, \theta^*]$ and increasing in (θ^*, t) . So the minimum of $\psi_2(y_u, \theta)$ in $[0, t]$ is achieved at θ^* . Also, since $\frac{\partial^2 \psi_2}{\partial \theta^2} < 0$ in $(t, 1)$, the minimum of $\psi_2(y_u, \theta)$ in

$[t, 1)$ is achieved at either t or 1^- . To sum up, the overall minimum is achieved at either θ^* or 1^- :

$$\begin{aligned}\psi_2(y_u, \theta^*) &= \int_0^{\theta^*} g(y_v) dy_v + (1 - \theta^*)h(\theta^*) + (1 - \theta^*)(1 - g(\theta^*)) \approx 0.5359 > 0.5349 \\ \psi_2(y_u, 1^-) &= \int_0^1 g(y_v) dy_v \approx 0.5381 > 0.5349,\end{aligned}$$

Thus we have $\min_{\theta < 1} \psi_2(y_u, \theta) = \min\{\psi_2(y_u, \theta^*), \psi_2(y_u, 1^-)\} > 0.5349$, as required. \blacksquare

Next we consider ψ_1 .

Lemma B.9 *For all $\theta \leq \tau < 1$, we have $\psi_1(y_u, \theta, \tau) > \min\{g(y_u), 0.5349\}$.*

Proof: If $g(y_u) \leq \phi(\tau)$, then we have $\psi_1(y_u, \theta, \tau) = \int_0^\theta g(y_v) dy_v + (\tau - \theta)h(\theta) + (1 - \theta)g(y_u) + \theta g(y_u) \geq g(y_u)$, as required. Now consider $g(y_u) > \phi(\tau)$. Observe that

$$\frac{\partial \psi_1}{\partial \tau} = h(\theta) - \theta(g'(\tau) + h'(\tau)) \leq 0,$$

where the last inequality holds since $h(\theta) \leq k_h^1 \theta \leq (k_g^2 + k_h^2) \theta \leq (g'(\tau) + h'(\tau)) \theta$ for all τ . Thus, for all τ we have $\psi_1(y_u, \theta, \tau) \geq \psi_1(y_u, \theta, 1^-)$, i.e., the minimum is achieved when $\tau = 1^-$.

Depending on whether $g(y_u) \geq \phi(\theta)$, we consider two cases. If $g(y_u) < \phi(\theta)$, we have

$$\psi_1(y_u, \theta, \tau) = \int_0^\theta g(y_v) dy_v + (\tau - \theta) \cdot h(\theta) + (1 - \theta) \cdot g(y_u) + \theta \cdot \phi(\tau).$$

For any fixed θ , the minimum is achieved when $\tau = 1^-$, which is

$$\begin{aligned}\psi_1(y_u, \theta, 1^-) &= \int_0^\theta g(y_v) dy_v + (1^- - \theta)h(\theta) + (1 - \theta)g(y_u) + \theta \cdot \phi(1^-) \\ &\geq \int_0^\theta g(y_v) dy_v + (1 - \theta)h(\theta) + (1 - \theta)g(y_u) \geq \psi_2(y_u, \theta) \geq \min\{g(y_u), 0.5349\},\end{aligned}$$

where the last inequality follows from Lemma B.8. If $g(y_u) \geq \phi(\theta)$, we have

$$\psi_1(y_u, \theta, \tau) = \int_0^\theta g(y_v) dy_v + (\tau - \theta) \cdot h(\theta) + (1 - \theta) \cdot \phi(\theta) + \theta \cdot \phi(\tau),$$

the minimum of which is achieved when $\tau = 1^-$. Define $\psi(\theta)$ to be the minimum:

$$\begin{aligned}\psi(\theta) &\stackrel{\text{def}}{=} \psi_1(y_u, \theta, 1^-) = \int_0^\theta g(y_v) dy_v + (1^- - \theta) \cdot h(\theta) + (1 - \theta) \cdot \phi(\theta) + \theta \cdot \phi(1^-) \\ &= \int_0^\theta g(y_v) dy_v + (1 - \theta) \cdot (1 - g(\theta)) + \theta \cdot (1 - g(1^-) - h(1^-)).\end{aligned}$$

By the following, we have $\frac{\partial \psi}{\partial \theta} > 0$ for all $\theta \in (t, 1)$, i.e., $\psi(\theta)$ is increasing when $\theta \in (t, 1)$.

$$\begin{aligned}\frac{\partial \psi}{\partial \theta} &= g(\theta) + (1 - \theta)(-g'(\theta)) - (1 - g(\theta)) + 1 - g(1^-) - h(1^-). \\ \frac{\partial^2 \psi}{\partial \theta^2} &= 3g'(\theta) > 0, \quad \forall \theta \in [0, t) \cup (t, 1) \\ \left. \frac{\partial \psi}{\partial \theta} \right|_{\theta=0} &= 2g(0) - k_g^1 - g(1^-) - h(1^-) \approx -0.08 < 0 \\ \left. \frac{\partial \psi}{\partial \theta} \right|_{\theta=t} &= 2g(t) - (1 - t)k_g^2 - g(1^-) - h(1^-) = 0.186 > 0.\end{aligned}$$




Figure B.3: $\psi(\theta)$ and $\frac{\partial \psi}{\partial \theta}$

Hence the minimum of $\psi(\theta)$ is achieved when $\theta \in (0, t]$. Let $\frac{\partial \psi}{\partial \theta} = 0$, we have

$$\begin{aligned}
 \frac{\partial \psi}{\partial \theta} &= g(\theta) - (1 - \theta)g'(\theta) - (1 - g(\theta)) + 1 - g(1^-) - h(1^-) = 0 \\
 \iff & 3k_g^1 \theta + 2b - g(1^-) - h(1^-) - k_g^1 = 0. \\
 \iff & \theta^* = \frac{k_g^1 + g(1^-) + h(1^-) - 2b}{3k_g^1} \approx 0.127.
 \end{aligned}$$

Thus for all $\theta \leq \tau < 1$, we have $\psi_1(y_u, \theta, \tau) \geq \psi(\theta) \geq \psi(\theta^*) \approx 0.5349$, as claimed. ■

C Hardness Results

Proof of Theorem 1.3: Consider the following hard instance. Let k, h be integer parameters, and $n := \sum_{i=0}^h k^i = \frac{k^{h+1}-1}{k-1}$ be the number of vertices on each side of a bipartite graph. In the following, we construct a bipartite graph on vertices $U \cup V$, where $U = \{u_1, \dots, u_n\}$ and $V = \{v_1, \dots, v_{n-k^h}, b_1, \dots, b_{k^h}\}$. It is easy to check by our construction that the graph is bipartite, but U, V does not correspond to the two sides of the bipartite graph.

Hard Instance. Refer to Figure C.1 (an illustrating example with $k = 3$ and $h = 2$). At the beginning, vertex u_1 arrives, together with all its $k + 1$ neighbors (children). Let the deadline of u_1 be reached immediately. Then we choose uniformly at random k vertices from the $k + 1$ neighbors of u_1 to be u_2, \dots, u_{k+1} . Let the remaining vertex be v_1 . We repeat the procedure for vertices u_2, \dots, u_{k+1} , i.e., each vertex u_i has $k + 1$ children, among which k vertices are chosen to be $u_{(i-1)k+2}, \dots, u_{ik+1}$ while the remaining one becomes v_i , and let the deadline of u_i be reached immediately. We continue building the tree for h levels. Note that in level i , there are k^i vertices (excluding the v vertices). Hence the tree with h levels has n vertices.

At last, we pick a random permutation $A = (a_1, a_2, \dots, a_{k^h})$ of the k^h vertices $\{u_{n-k^h+1}, \dots, u_n\}$ at level h . Let vertices $B = \{b_1, \dots, b_{k^h}\}$ arrive (b_1 arrives first and b_{k^h} last), such that b_i is connected to vertices a_i, \dots, a_n , and the deadline of b_i is reached immediately when it arrives.

Let the deadlines of vertices in $A \cup (V \setminus B)$ be reached at the end.




Figure C.1: Hard instance for any online algorithm: an illustrating example with $k = 3$ and $h = 2$.

Competitive Ratio. First observe that graph G has a perfect matching, by matching u_i to v_i (for all $i \in [n - k^h]$) and a_i to b_i (for all $i \in [k^h]$). Now we consider any online algorithm. Note that when the deadline of u_i is reached, it is not worse to match u_i if it has an unmatched neighbor: if we do not match u_i , then the symmetric difference is an alternating path, thus the number of vertices matched does not increase. Hence we assume w.l.o.g. that all vertices in $U \setminus A$ will be matched eventually.

Let p_i be the probability that a vertex u_x from level i is matched when the deadline of its parent u_y in the tree is reached. Note that p_i is also the probability that v_y is matched, as v_y is chosen uniformly at random among the $k + 1$ children of u_y . Observe that we have $p_i = \frac{1 - p_{i-1}}{k + 1}$, where $p_0 = 1$. It is easy to check (by induction) that for all $i \leq h$, we have

$$p_i = \frac{1}{k + 2} \left(1 - \left(\frac{-1}{k + 1} \right)^i \right)$$

Hence before vertex b_1 arrives, each vertex from A is matched with probability p_h . By the standard water-filling algorithm, it is easy to see that the expected number of matched vertices from B (at the end of the algorithm) is t such that

$$\frac{1}{k^h} + \frac{1}{k^h - 1} + \dots + \frac{1}{k^h - t + 1} = 1 - p_h = \frac{k + 1}{k + 2} + \frac{1}{k + 2} \left(\frac{-1}{k + 1} \right)^h.$$

When h tends to infinity, we have $t \approx (1 - e^{-\frac{k+1}{k+2}}) \cdot k^h$. Hence the competitive ratio is

$$\frac{2t + |U \setminus A| + \frac{1}{k+2}|A| + \frac{1}{k+2}|V \setminus B|}{2n} = \frac{2(1 - e^{-\frac{k+1}{k+2}}) \cdot k^h \cdot (k - 1) + k^h - 1}{2(k^{h+1} - 1)} + \frac{1}{2(k + 2)},$$

which tends to $\frac{k-1}{k}(1 - e^{-\frac{k+1}{k+2}}) + \frac{1}{2k} + \frac{1}{2(k+2)}$ when h tends to infinity. For $k = 7$, the ratio is $\frac{62}{63} - \frac{6}{7} \cdot e^{-\frac{8}{9}} \approx 0.631745$. \blacksquare

Proof of Theorem 1.4: Consider the following hard instance. Let k, h be integer parameters, and $n := k \cdot h$ be the number of vertices on each side of a bipartite graph G . In the following, we construct a bipartite graph on vertices $U \cup V$, where $U = \{u_1, \dots, u_n\}$ and $V = \{v_1, \dots, v_n\}$. As before, it is easy to check by our construction that the graph is bipartite, but U, V does not correspond to the two sides of the bipartite graph.

Hard Instance. Refer to Figure C.2. For all $i \in [n]$, let u_i be the only neighbor of v_i . We group every k consecutive vertices in U as a group, i.e., let $U = \cup_{i \in [h]} U_i$, where the i -th group $U_i = \{u_{(i-1)k+1}, u_{(i-1)k+2}, \dots, u_{ik}\}$. Let there be an edge between u_i and u_j if they are from two consecutive groups, respectively. In other words, we form a complete bipartite graph between any two consecutive groups U_i and U_{i+1} .




Figure C.2: Hard instance for Ranking: an illustrating example with $k = 3$ and $h = 4$: vertices in U and V are represented by the solid black circles and white circles, respectively.

Let the deadline of vertex u_1 be reached first, then u_2 's deadline, u_3 's deadline, etc.

Competitive Ratio of Ranking. It is easy to see that graph G has a perfect matching, by matching u_i to v_i for each $i \in [n]$. Recall that in the **Ranking** algorithm, each vertex u is assigned a random rank $y_u \in [0, 1)$. At the deadline of an unmatched vertex, it is matched to its unmatched neighbor v (if any) with the smallest y_v . Observe that in our instance, all vertices from U will be matched eventually, while each $v_i \in V$ will be matched only if at the deadline of u_i , u_i is unmatched and y_{v_i} is smaller than the ranks of all unmatched vertices from the next group $U_{\lceil \frac{i}{k} \rceil + 1}$.

For all $i \in [h]$, let $X_i \in \{0, 1, \dots, k\}$ be the number of unmatched vertices in U_i right before the deadline of the first vertex $u_{(i-1)k+1}$ in U_i is reached. It is easy to see that X_{i+1} is a random variable that depends only on X_i . Hence the sequence X_1, X_2, \dots, X_h forms a Markov chain (with $k+1$ states) with initial state $X_1 = k$. Observe that all vertices in U are matched, and the number of vertices matched in V_i equals $X_i + X_{i+1} - k$. Hence the competitive ratio of **Ranking** $\approx \frac{\sum_{i \in [h]} X_i}{kh}$.

We say phase i begin when the deadline of the first vertex of U_i is reached, and end after the deadline of the last vertex of U_i . Fix any phase i , where $i < h$. Recall that initially X_i vertices of U_i are unmatched. Let $Z(t)$ be the number of unmatched vertices in U_{i+1} , when the deadlines of exactly t unmatched vertices in U_i have passed. We have $Z(0) = k$ and $X_{i+1} = Z(X_i)$. Let $y_1 \leq y_2 \leq \dots \leq y_k$ be the ranks of vertices in U_{i+1} . It is easy to see that $\mathbf{E}[Z(t+1)] = Z(t) - 1 + y_{Z(t)}$. Taking expectation over all y_i 's, we have $\mathbf{E}[Z(t+1)] = \mathbf{E}[Z(t)] - 1 + \frac{\mathbf{E}[Z(t)]}{k+1}$. Let $z(\frac{t}{k}) \stackrel{\text{def}}{=} \frac{Z(t)}{k}$. It is easy to see when $k \rightarrow \infty$, $z(\frac{t}{k}) \rightarrow e^{-\frac{t}{k}}$. This is saying, given X_i , $\mathbf{E}[X_{i+1}] = k \cdot e^{-\frac{X_i}{k}}$ when k tends to infinity.

Finally, note that all vertices in U_{i+1} are symmetric. Hence, each of them is unmatched at the end of phase i with probability $\frac{\mathbf{E}[X_{i+1}]}{k} = e^{-\frac{X_i}{k}}$. Moreover, for any two vertices $u_a, u_b \in U_{i+1}$, the probability that u_a is unmatched at the end of phase i is negatively correlated with the probability of u_b : conditioned on u_b being unmatched at the end of phase i , the probability of u_a being unmatched is smaller. Thus we have measure concentration bound on X_{i+1} , by standard argument

using moment generation function. In other words, the stationary distribution (when k and h tends to infinity) converges to a single point mass with,

$$\frac{X}{k} = e^{-\frac{x}{k}},$$

which implies $\frac{X}{k} \approx 0.56714$, the Omega constant, which is also the competitive ratio of Ranking. ■