

Tight Competitive Ratios of Classic Matching Algorithms in the Fully Online Model*

Zhiyi Huang[†] Binghui Peng[‡] Zhihao Gavin Tang[†] Runzhou Tao[‡] Xiaowei Wu[§]
Yuhao Zhang[†]

Abstract

Huang et al. (STOC 2018) introduced the fully online matching problem, a generalization of the classic online bipartite matching problem in that it allows all vertices to arrive online and considers general graphs. They showed that the ranking algorithm by Karp et al. (STOC 1990) is strictly better than 0.5-competitive and the problem is strictly harder than the online bipartite matching problem in that no algorithms can be $(1 - 1/e)$ -competitive.

This paper pins down two tight competitive ratios of classic algorithms for the fully online matching problem. For the fractional version of the problem, we show that a natural instantiation of the water-filling algorithm is $2 - \sqrt{2} \approx 0.585$ -competitive, together with a matching hardness result. Interestingly, our hardness result applies to arbitrary algorithms in the edge-arrival models of the online matching problem, improving the state-of-art $\frac{1}{1+\ln 2} \approx 0.5906$ upper bound. For integral algorithms, we show a tight competitive ratio of ≈ 0.567 for the ranking algorithm on bipartite graphs, matching a hardness result by Huang et al. (STOC 2018).

1 Introduction

Following the seminal work by Karp et al. [KVV90] that initiated the study of the Online Bipartite Matching problem by proposing the Ranking algorithm, online matching problems have drawn a lot of attentions in the online algorithm literature. These problems have found numerous real-life applications, notably, in online

advertising. They are also the driving-force behind many important techniques for designing and analyzing online algorithms, including the randomized primal dual technique by Devanur et al. [DJK13].

Recently, Huang et al. [HKT⁺18] proposed a generalization of the Online Bipartite Matching problem called Fully Online Matching. The generalization considers general graphs and allows all vertices to arrive online. It captures a much wider family of real-life scenarios, including the ride-sharing problem. Concretely, consider an undirected graph $G = (V, E)$. Each step is either the arrival or the deadline of a vertex. At a vertex v 's arrival, all the edges between v and those that arrive before v are revealed. At its deadline, on the other hand, the algorithm must irrevocably either match it to an unmatched neighbor (if it is not matched already) or leave it unmatched. The model assumes that all neighbors of a vertex v arrive before v 's deadline. This turns out to be a natural condition when it comes to concrete scenarios such as ride-sharing.

Further, Huang et al. [HKT⁺18] showed that the Fully Online Matching problem is quite intriguing from an algorithmic viewpoint in that 1) it takes a number of novel ideas to show that the Ranking algorithm by Karp et al. [KVV90] is strictly better than 0.5-competitive even in the fully online setting, and 2) the fully online setting, even on bipartite graphs, is strictly harder than the original Online Bipartite Matching problem in that no algorithms can be $1 - 1/e \approx 0.632$ -competitive.

1.1 Our Contributions and Techniques. We develop better understandings on the Fully Online Matching problem by establishing two tight competitive ratios. The first result considers the fractional version of the problem, where we are allowed to fractionally match each vertex to multiple neighbors so long as the total mass sum to at most one. We show that the Water-Filling algorithm, which at each vertex's deadline matches its unmatched portion fractionally to all neighbors with smallest matched portion (i.e., the lowest water-level), gets a competitive ratio of $2 - \sqrt{2} \approx 0.585$. We also construct a matching hard instance for Water-

*Supported in part by the Hong Kong RGC grant 17202115E. Zhihao Gavin Tang is partially supported by his supervisor Hubert Chan's Hong Kong RGC grant 17202715. Xiaowei Wu is partially sponsored by Project 11771365 supported by NSFC.

[†]Department of Computer Science, The University of Hong Kong. Email: {zhiyi, zhtang, yhzhang2}@cs.hku.hk.

[‡]Institute for Interdisciplinary Information Sciences, Tsinghua University. Part of the work was done while Binghui and Runzhou were visiting the University of Hong Kong. Email: {pbh15, trz15}@mails.tsinghua.edu.cn.

[§]Department of Computer Science, City University of Hong Kong, Email: wxw0711@gmail.com.

Filling. The hardness result applies to arbitrary algorithms if we consider edge arrival models [BST17], even when preemptions are allowed [ELSW13, McG05], improving the best known bounds in these models. The second result focuses on the integral problem and the **Ranking** algorithm. We prove that its competitive ratio is exactly the Ω constant¹ ≈ 0.567 on bipartite graphs, improving the previous bound of ≈ 0.554 and matching the previous hardness result by Huang et al. [HKT⁺18]. See Figure 1 for where our results sit compared with the previous works.

Competitive Analysis of Water-Filling. The analysis of the Water-Filling algorithm is the relatively easy part of the paper. We follow the online primal dual framework by Buchbinder et al. [BJN07], building on the notions of passive and active vertices by Huang et al. [HKT⁺18].

When a vertex u matches another vertex v at u 's deadline, Huang et al. [HKT⁺18] referred to u as the active vertex and v as the passive vertex. Intuitively, when edge (u, v) is of concern, v plays a role similar to an offline vertex in the Online Bipartite Matching problem since it sits back and allows u to make the matching decision, while u plays a role similar to an online vertex. Following the same principle, for every vertex v , we refer to the portion that is matched before its deadline as the passive portion, and the portion that is matched at its deadline as the active portion.

When a small portion p of edge (u, v) is chosen into the fractional matching, we shall split the gain of p between the endpoints according to the current water-level x_v of the passive vertex v (i.e., the one with a later deadline). For some function g to be chosen in the analysis, u shall get $(1 - g(x_v)) \cdot p$ while v shall get $g(x_v) \cdot p$. Then, by an appropriate argument, we can lower bound the total gain of u and v by:

$$(1.1) \int_0^{p_u} g(x)dx + (1 - p_u)(1 - g(x_v)) + \int_0^{x_v} g(x)dx.$$

Here, p_u is the passive portion of u , and x_v is the passive portion of v after u 's deadline. The first term is the gain of u due to its passive portion. The second term lower bounds the gain of u due to its active portion. The third term lower bounds of the gain of v due to its passive portion.

It remains to choose g to maximize the above lower bound against the worst p_u and x_v . Unlike in the primal dual analysis of some other online matching problems, this is not exactly a standard ODE. Nonetheless, we observe that it is almost symmetric w.r.t. p_u and x_v . Indeed, choosing g to be an appropriate linear function

¹This is the solution of $\Omega \cdot e^\Omega = 1$.

makes it symmetric and yields the optimal $2 - \sqrt{2}$ bound.

Matching Hardness for Water-Filling. Constructing a hard instance to show a matching $2 - \sqrt{2}$ upper bound on the competitive ratio of Water-Filling presents some technical obstacles beyond the existing techniques. The construction is driven by Eqn. (1.1). By our choice of g , Eqn. (1.1) is equal to the lower bound $2 - \sqrt{2}$ only if p_u and x_v sum to precisely $2 - \sqrt{2}$. Further, the performance of the algorithm is equal to the gain of the endpoints summing over all edges in the optimal matching in hindsight. Therefore, a matching hard instance must satisfy that before the matching decision is made for an edge (u, v) in the optimal matching, the water-levels of the two endpoints are prepared in advance so that the sum equals $2 - \sqrt{2}$. This suggests that a tight instance for Water-Filling must look very different from the existing hard instances in the previous works (e.g., [KVV90, DJ12, HKT⁺18]), where for every edge (u, v) in the optimal matching, one of the two endpoints simply shows up with zero water-level and a matching decision is made for the edge.²

Our construction prepares the water-level of the vertices via a dynamic as follows. It maintains at all time a set of vertices with some number of vertices at each water-level x for $0 \leq x \leq 2 - \sqrt{2}$. At each step, pick a vertex u with an appropriate water-level x_u and let it be u 's deadline. Vertex u connects to a subset of the vertices with water-level $2 - \sqrt{2} - x_u$, among which one vertex v is u 's partner in the optimal matching. After the step, u and v will be removed from the pool; new vertices (with zero water-level) will arrive to refill the pool if needed. The matching decision of u "pumps up" the water-level of all its neighbors to $2 - \sqrt{2} - x_u + \epsilon$. Some of them will serve as the active endpoints with this water-level of some edges in the optimal matching; some of them will serve as the passive counterparts; the water-level of the remaining will be further "pumped up" by some vertex with water-level $x_u - \epsilon$. We show how to maintain such a dynamic so that, in the long run, the endpoints of any edge in the optimal matching will have a total water-level close to $2 - \sqrt{2}$ when a matching decision is made for the edge.

Competitive Analysis of Ranking on Bipartite Graphs. We first explain why the previous analysis of Huang et al. [HKT⁺18] is not tight on their hard instance. Consider an edge (u, v) in the optimal matching where u has an earlier deadline. The previous analysis is tight only if there is a threshold θ such that whenever v 's rank is larger than θ , u is matched and v is

²It is easy to show that one cannot maintain at all time that some vertices have $2 - \sqrt{2}$ water-level while the other have 0 with the Water-Filling algorithm.

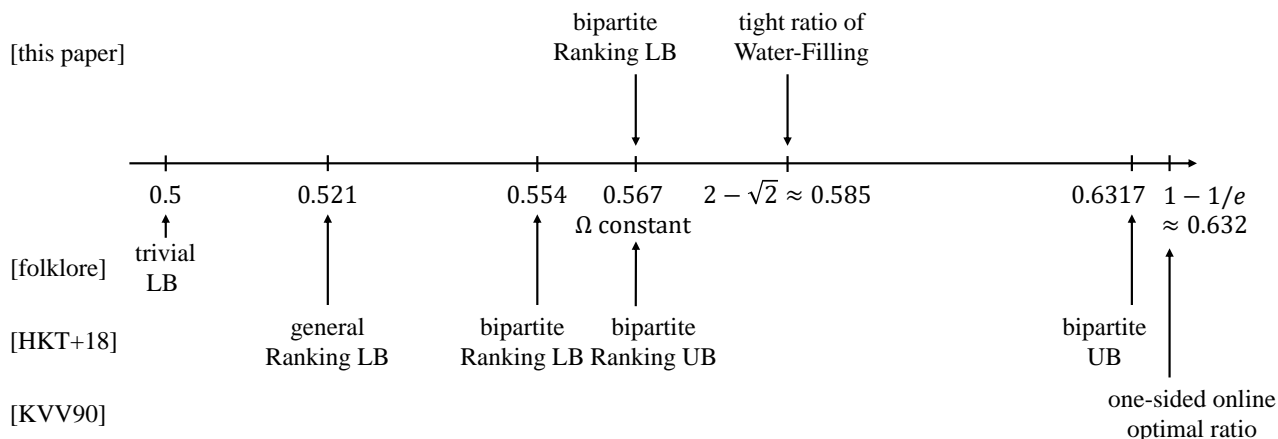


Figure 1: A comparison of the results in this paper and those in previous work. LB means lower bound (algorithmic results), and UB means upper bound (hardness results).

unmatched and, more importantly, whenever v 's rank is smaller than θ , v is passively matched and u matches to the same vertex as in the previous case. In the hard instance, however, u is v 's only neighbor. Therefore, if u 's own rank is sufficiently large such that u matches actively, u and v will match each other when v 's rank is smaller than θ . Taking this extra gain into account gives the optimal ratio of ≈ 0.567 for the hard instance.

Of course, we cannot naïvely assume that one of the endpoints of any edge in the optimal matching will have only one neighbor. The point is the previous approach that tries to characterize the matching status of u and v using a single threshold of v cannot possibly capture the above extra gain. We show that a good enough characterization in general takes three thresholds, a threshold of u and two thresholds of v , one with u in the graph and one without. As a result, we get a new lower bound on the total expected gain of the endpoints that is strictly better than the previous one in all but a few bottleneck cases. Then, we design a different gain sharing function that focuses on these bottleneck cases to obtain a tight analysis.

Finally, we remark that the three thresholds pin down when u and v match each other. Previous works on online matching usually omit the gain from this case, which indeed happens with negligible probability in the worst case of those models (with a recent exception of [HTWZ18]). Our analysis shows it was just a lucky coincident that we do not need to consider the case when the endpoints match each other in those problems. It becomes critical in a more general online matching model.

1.2 Other Related Works Following Karp et al. [KVV90], a series of works study different variants of the problem, including b -matching [KP00], ad-words [MSVV07, BJN07, DJ12], vertex-weighted matching [AGKM11] and the random arrival model [KMT11, MY11, HTWZ18]. Besides, the analysis of Ranking has been simplified in a series of papers [GM08, BM08, DJK13].

The Water-Filling algorithm has been studied to tackle several versions of the Online Bipartite Matching problems [BJN07, KP00]. Devanur et al. [DHK⁺13] considered the whole page optimization problem and extended the Water-Filling algorithm to use a carefully designed “level function” instead of a single water-level. Wang and Wong [WW15] considered an alternative model of Online Bipartite Matching that allows both sides of vertices to arrive online. They showed a 0.526-competitive algorithm for a fractional version of the problem. Both analysis of [DHK⁺13, WW15] are based on the online primal dual framework by [BJN07]. This paper further illustrates the power of this framework for studying online fractional matching problems.

The hardness result in this paper improves the bounds for the following online matching models. In online preemptive matching [ELSW13, McG05], each edge arrives online and the algorithm must immediately decide whether to add the edge to the matching and to dispose of previously selected edges if needed. A harder edge-arrival model [BST17] forbids edge disposals. For both problems, the best previous bound stands at $\frac{1}{1+\ln 2} \approx 0.5906$ [ELSW13].

Very recently, weighted variants of Fully Online Matching have been studied by [ABJS18, DS18], both

considering the “windowed” version of the problem, motivated by the ride-sharing applications.

2 Preliminaries

We study both the fractional and the integral versions of Fully Online Matching. When the underlying graph is bipartite, we refer to the problem as Fully Online Bipartite Matching. Consider the following standard linear program formulation of the matching problem and its dual.

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

Fractional Matching. In this setting, we may match edges fractionally. Let $x_{uv} \in [0, 1]$ be the fraction of edge (u, v) in the matching. Assuming u has an earlier deadline than v , this variable increases only at u 's deadline. We refer to it as Fully Online Fractional Matching and study the classic Water-Filling algorithm (e.g., [BJN07]) in this setting. We give a formal definition of the algorithm below, in which the dual variables are updated as well. Note that the dual variables are used only in the analysis. We fix an increasing function $g : [0, 1] \rightarrow [0, 1]$ to be specified later and use $x_u \stackrel{\text{def}}{=} \sum_{v:(u,v) \in E} x_{uv}$ to keep track of the water-level (i.e. total fractional mass) of u at all time.

Algorithm 1 The Water-Filling Algorithm

Initialize all x_{uv} 's and α_u 's to be zero.

When the deadline of vertex u is reached:

Let $p_u = x_u$ be the water-level collected before u 's deadline. $\triangleright p_u$: passive water-level of u .

Let $N(u)$ be the set of neighbors of u whose deadlines are not reached.

while $x_u < 1$ and $\min_{v \in N(u)} \{x_v\} < 1$ **do**

Allocate a dx amount to each x_{uv} for $v \in \arg \min_{v \in N(u)} \{x_v\}$.

If x_{uv} increases by dx , increase α_u and α_v respectively by

$$d\alpha_u = (1 - g(x_v))dx \quad \text{and} \quad d\alpha_v = g(x_v)dx.$$

We call the vertices in $N(u)$ the *available* neighbors of u at u 's deadline. We further import the notions of active and passive vertices from [HKT⁺18] and define them for both fractional and integral algorithms.

DEFINITION 2.1. (ACTIVE, PASSIVE) For any edge (u, v) that is (fractionally) matched by an algorithm at u 's deadline, we say that u is active and v is passive (w.r.t. edge (u, v)).

Integral Matching. In this setting, x_{uv} 's must have binary values. We will analyze the Ranking algorithm in Section 4 when the underlying graph G is bipartite. Recall the definition of Ranking and some important notions from [HKT⁺18].

Algorithm 2 The Ranking Algorithm [HKT⁺18]

- (1) a vertex v arrives:
 - pick $y_v \in [0, 1]$ uniformly at random.
 - (2) a vertex v 's deadline is reached:
 - 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 $\arg \min_{u \in N(v)} \{y_u\}$.
-

Let $M(\vec{y})$ denote the matching produced when Ranking is run with \vec{y} as the ranks.

DEFINITION 2.2. (MARGINAL RANK [HKT⁺18]) For any u and any ranks \vec{y}_{-u} of other vertices, the marginal rank θ of u w.r.t. \vec{y}_{-u} is the largest value such that u is passive in $M(y_u = \theta^-, \vec{y}_{-u})$.

The following is a restatement of Lemma 2.5 from [HKT⁺18] when restricted to bipartite graphs.

LEMMA 2.1. In a bipartite graph, if u is matched in \vec{y} , then from $M(\vec{y})$ to $M(\vec{y}_{-u})$, all neighbors of u do not 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 smaller rank is better.

We set primal variables according to Ranking. The randomized primal dual technique [DJK13] allows us to prove competitive ratio bounds through the following.

LEMMA 2.2. ([HKT⁺18], LEMMA 2.6) 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) $\mathbf{E}_{\vec{y}}[\alpha_u + \alpha_v] \geq F$ for all $(u, v) \in E$.

3 Tight Competitive Ratio of Water-Filling

In this section, we give a tight analysis on the competitive ratio of the Water-Filling algorithm for the Fully Online Fractional Matching problem.

3.1 Lower Bound on the Competitive Ratio We first prove that the competitive ratio of **Water-Filling** is at least $2 - \sqrt{2}$. Our approach is based on a primal dual analysis.

THEOREM 3.1. *Water-Filling is $(2 - \sqrt{2})$ -competitive.*

Proof. Recall that we update the primal variables according to **Water-Filling** and dual variables in a way that the dual objective always equals the primal objective. Using the standard primal dual technique, in order to prove that **Water-Filling** is $(2 - \sqrt{2})$ -competitive, it suffices to show that $\alpha_u + \alpha_v \geq 2 - \sqrt{2}$ for all pairs of neighbors u and v .

Let $g(x) = \frac{\sqrt{2}}{2}x + 1 - \frac{\sqrt{2}}{2}$ be the function we used for defining dual variables.

Fix any pair of neighbors u, v where u has an earlier deadline than v . Consider the moment right **after** u 's deadline. It must be that either $x_u = 1$ or $x_v = 1$ (otherwise x_u will further increase). As v can only be matched passively, if $x_v = 1$, we have

$$\alpha_u + \alpha_v \geq \int_0^1 g(x)dx = 1 - \frac{\sqrt{2}}{4} \geq 2 - \sqrt{2}.$$

Now suppose $x_u = 1$ and $x_v < 1$. Then, we have $\alpha_v = \int_0^{x_v} g(x)dx$. Next, consider the value of α_u . Before u 's deadline, we have $\alpha_u = \int_0^{p_u} g(x)dx$ (recall that p_u is the passive water-level of u). Since $x_u = 1$, and the water-level of v after the deadline of u is $x_v < 1$, at any moment when the water-level of u is increased from p_u to 1, the neighbor that u matches has a water-level at most x_v . Hence, we have

$$\begin{aligned} \alpha_u &\geq \int_0^{p_u} g(x)dx + \int_{p_u}^1 (1 - g(x_v))dx \\ &= \int_0^{p_u} g(x)dx + (1 - p_u)(1 - g(x_v)). \end{aligned}$$

Summing the lower bounds on the two dual variables and by the definition of g , we have

$$\begin{aligned} \alpha_u + \alpha_v &\geq \int_0^{p_u} g(x)dx + (1 - p_u)(1 - g(x_v)) \\ &\quad + \int_0^{x_v} g(x)dx \\ &= \frac{\sqrt{2}}{4}(p_u^2 + x_v^2) + (1 - \frac{\sqrt{2}}{2})(p_u + x_v) \\ &\quad + (1 - p_u)(\frac{\sqrt{2}}{2} - \frac{\sqrt{2}}{2}x_v) \\ &= \frac{\sqrt{2}}{4} \left((p_u + x_v) - (2 - \sqrt{2}) \right)^2 + 2 - \sqrt{2} \\ &\geq 2 - \sqrt{2}. \end{aligned}$$

Hence, in both cases we have $\alpha_u + \alpha_v \geq 2 - \sqrt{2}$, which gives the $2 - \sqrt{2}$ lower bound on the competitive ratio of **Water-Filling**.

3.2 Upper Bound on the Competitive Ratio In this section we explicitly construct a hard instance, for which **Water-Filling** gives a solution of value $(2 - \sqrt{2}) \cdot \text{OPT}$.

Hard Instance. Let there be $2k \cdot m$ vertices, which are partitioned into m groups of size $2k$. For all $t \in [m]$, let the vertices in the t -th group be $U_t \cup V_t$, where $U_t = \{u_{t,1}, \dots, u_{t,k}\}$ and $V_t = \{v_{t,1}, \dots, v_{t,k}\}$. Let $h : [0, 1] \rightarrow [0, 1]$ be a decreasing function³ (to be determined later) with $h(0) = 1$ and $h(1) = 0$. There are two types of edges in the graph (refer to Figure 2):

Upper triangle edges between U_t and V_t : $\forall t \in [m], i \in [k]$ and $j \geq i, (u_{t,i}, v_{t,j}) \in E$;

h -induced edges between U_t and U_{t+1} : $\forall t \in [m - 1], i \in [k]$ and $j \leq \lfloor k \cdot h(\frac{i-1}{k}) \rfloor, (u_{t,i}, u_{t+1,j}) \in E$.




Figure 2: Subgraph induced by $U_t \cup V_t \cup U_{t+1}$: illustrating example with $t = 1$ and $k = 4$

Finally, let the deadlines of the u vertices be reached first, following the lexicographical order on (t, i) . Then let the deadlines of the v vertices be reached, i.e., after the deadline of $u_{m,k}$.⁴

It is easy to see that the hard instance is bipartite, where $(U_1 \cup U_3 \cup \dots) \cup (V_2 \cup V_4 \cup \dots)$ and $(U_2 \cup U_4 \cup \dots) \cup (V_1 \cup V_3 \cup \dots)$ are the two sides of vertices. This graph admits a perfect matching, in which $u_{t,i}$ matches $v_{t,i}$ for all $t \in [m], i \in [k]$ and hence, $\text{OPT} = km$.

We first construct the function h and prove the following technical lemma. Let $c = 2 - \sqrt{2}$, and function

³When $h(x) \equiv 1$, our instance becomes the $\frac{1}{1+\ln 2} \approx 0.5906$ hard instance by [ELSW13] for the edge arrival model.

⁴The relative order of the deadlines of v vertices does not matter, as long as $v_{t,i}$'s deadline is after $u_{t,i}$'s deadline.

$f : [0, c] \rightarrow [0, 1]$ be defined as

$$f(x) \stackrel{\text{def}}{=} \frac{1}{2} (\ln(1-x) + \ln(1-c+x)) + \frac{1}{\sqrt{2}(x-1)} + \frac{2 + \sqrt{2} - \ln(1-c)}{2}.$$

Let $\tau(x) \stackrel{\text{def}}{=} f^{-1}(x)$ and $h(x) \stackrel{\text{def}}{=} f(c - f^{-1}(x))$. It is not difficult to see that f is strictly decreasing. Hence, functions $h : [0, 1] \rightarrow [0, 1]$ and $\tau : [0, 1] \rightarrow [0, c]$ are well defined. Moreover, since $f(0) = 1$ and $f(c) = 0$, we have that h is decreasing, $h(0) = 1$ and $h(1) = 0$, as required in the construction of the hard instance. These functions might seem mysteries at this point, we will show a connection between the functions h and g via duality in Appendix A, where g is the gain sharing function that we used to define the dual variables in Water-Filling.

LEMMA 3.1. For all $x \in [0, 1]$ we have

$$\int_0^x \frac{1 - \tau(y)}{1 - y + h(y)} dy = c - \tau(x), \quad \int_0^1 \tau(y) dy = 1 - c$$

and

$$\int_0^1 \frac{1}{1 - y + h(y)} dy < 1.$$

Proof. First we show the first equation, i.e., for all $x \in [0, 1]$ we have $\int_0^x \frac{1 - \tau(y)}{1 - y + h(y)} dy = c - \tau(x)$. Note that $\tau(0) = c$ and, thus, both sides equal 0 when $x = 0$. It suffices to check that for all $x \in [0, 1]$, $\frac{1 - \tau(x)}{1 - x + h(x)} = -\tau'(x)$. Let $\phi = \tau(x) \in [0, c]$, we have $f(\phi) = x$ and $h(x) = f(c - f^{-1}(x)) = f(c - \phi)$. Then, we only need to check that

$$\frac{1 - \phi}{1 - f(\phi) + f(c - \phi)} = -\tau'(x) = -\frac{1}{f'(\phi)},$$

which is true as f is defined such that for all $\phi \in [0, c]$,

$$1 - f(\phi) + f(c - \phi) + (1 - \phi)f'(\phi) = 0.$$

Taking integration from 0 to c , the contributions of the 2nd and the 3rd terms cancel. We have

$$0 = c + \int_0^c (1 - x)f'(x) dx = c - 1 + \int_0^c f(x) dx,$$

which implies the second equation because $\int_0^1 \tau(y) dy = \int_0^c f(x) dx = 1 - c$, where the first equality follows because $\tau = f^{-1}$, f is strictly decreasing, $f(0) = 0$, and $f(c) = 0$.

Now we prove the last equation, i.e., $\int_0^1 \frac{1}{1 - y + h(y)} dy < 1$.

Observe that both $1 - \tau(y)$ and $\frac{1}{1 - y + h(y)}$ are increasing in terms of y . Hence we have

$$\begin{aligned} c = c - \tau(1) &= \int_0^1 \frac{1 - \tau(y)}{1 - y + h(y)} dy \\ &> \int_0^1 (1 - \tau(y)) dy \cdot \int_0^1 \frac{1}{1 - y + h(y)} dy \\ &= c \cdot \int_0^1 \frac{1}{1 - y + h(y)} dy. \end{aligned}$$

Dividing both sides by c proves the last equation.

Now we analyze the performance of Water-Filling on this instance. We first prove that by running Water-Filling on the hard instance, the passive water-levels of almost all vertices are strictly smaller than 1.

LEMMA 3.2. For large enough k , Water-Filling produces a fractional matching with $p_{u_{t,i}} < 1$ for all $t \in [m], i \in [k]$ and $p_{v_{t,i}} < 1$ for all $t \in [m - 1], i \in [k]$.

Proof. Observe that at the deadline of each $u_{t,i}$, where $t \in [m - 1]$, it has $|N(u_{t,i}) \cap V_t| + |N(u_{t,i}) \cap U_{t+1}| = k - i + 1 + \lfloor k \cdot h(\frac{i-1}{k}) \rfloor$ neighbors whose deadlines are not reached. Moreover, as h is decreasing, it is easy to see (by induction) that at the deadline of $u_{t,i}$, all available neighbors of $u_{t,i}$ have the same water-level. Hence, Water-Filling increases the water-level of the available neighbors of $u_{t,i}$ at the same rate until $\min_{v \in N(u_{t,i})} x_v = 1$ or $x_{u_{t,i}} = 1$.

Since $u_{t+1,1}$ is a neighbor of every vertex in U_t , we have $p_{u_{t+1,1}} = \max_{j \in [k]} \{p_{u_{t+1,j}}, p_{v_{t,j}}\}$. Therefore, it suffices to show that $p_{u_{t+1,1}}$ is smaller than 1. Note that each vertex $u_{t,i}$ has at most 1 unit of unmatched portion that is distributed among $k - i + 1 + \lfloor k \cdot h(\frac{i-1}{k}) \rfloor$ available neighbors and, thus, it increases the water-level of $u_{t+1,1}$ by at most $\frac{1}{k - i + 1 + \lfloor k \cdot h(\frac{i-1}{k}) \rfloor}$. Hence, when $k \rightarrow \infty$, we have

$$\begin{aligned} p_{u_{t+1,1}} &\leq \sum_{i=1}^k \frac{1}{k - i + 1 + \lfloor k \cdot h(\frac{i-1}{k}) \rfloor} \\ &\approx \int_0^1 \frac{1}{1 - y + h(y)} dy < 1, \end{aligned}$$

where the last inequality follows from Lemma 3.1. This finishes the proof.

Lemma 3.2 implies that, for large enough k , we can guarantee that when running Water-Filling on the hard instance, after the deadline of every $u_{t,i}$, where $t \in [m - 1]$, we must have $x_{u_{t,i}} = 1$, as none of its neighbors with a later deadline has a water-level that reaches 1.

COROLLARY 3.1. *For all $t \in [m - 1]$, we have $x_{u_{t,i}} = 1$ after $u_{t,i}$'s deadline.*

Now we are ready to prove the main theorem of this section.

THEOREM 3.2. *Water-Filling is at most $(2 - \sqrt{2})$ -competitive.*

Proof. Let $\mathbf{p}_t = (p_{u_{t,1}}, p_{u_{t,2}}, \dots, p_{u_{t,k}})^T$ denote the passive water-level vector of U_t . Since the increment of matching at $u_{t,i}$'s deadline is at most $1 - p_{u_{t,i}}$, the solution given by Water-Filling is

$$\sum_{(u,v) \in E} x_{uv} \leq \sum_{t,i} (1 - p_{u_{t,i}}) = \sum_t (k - \|\mathbf{p}_t\|_1).$$

Indeed, by Corollary 3.1, for all $t \in [m - 1]$, the increment of matching at $u_{t,i}$'s deadline is exactly $1 - p_{u_{t,i}}$. Recall that in the hard instance, $u_{t+1,i}$ is a neighbor of $u_{t,j}$ iff $\frac{i}{k} \leq h(\frac{j-1}{k})$. Hence we have

$$\begin{aligned} p_{u_{t+1,i}} &= \sum_{j=1}^{\lfloor k \cdot h^{-1}(\frac{i}{k}) + 1 \rfloor} \frac{1 - p_{u_{t,j}}}{k - j + 1 + \lfloor k \cdot h(\frac{j-1}{k}) \rfloor} \\ &= \sum_{j=1}^{\lfloor k \cdot h^{-1}(\frac{i}{k}) + 1 \rfloor} (1 - p_{u_{t,j}}) \cdot a_j, \end{aligned}$$

where $a_j = \frac{1}{k-j+1+\lfloor k \cdot h(\frac{j-1}{k}) \rfloor}$ is independent of t . In other words, there exists a $k \times k$ matrix \mathbf{M} such that for all $t \in [m - 1]$, $\mathbf{p}_{t+1} = \mathbf{M}(\mathbf{1} - \mathbf{p}_t)$. More precisely, we have $M_{i,j} = a_j$ if $j \leq \lfloor k \cdot h^{-1}(\frac{i}{k}) + 1 \rfloor$, $M_{i,j} = 0$ otherwise. Hence, for any $i \in [k]$, by Lemma 3.1, we have

$$\sum_{j \in [k]} M_{i,j} \leq \sum_{j \in [k]} a_j < 1.$$

That is, \mathbf{M} is a contraction matrix and the above mapping from \mathbf{p}_t to \mathbf{p}_{t+1} has a unique stationary vector \mathbf{p}^* , i.e. $\mathbf{p}^* = \mathbf{M}(\mathbf{1} - \mathbf{p}^*)$. Moreover, $\lim_{t \rightarrow \infty} \mathbf{p}_t = \mathbf{p}^*$ ⁵. Thus, for any fixed k , when $m \rightarrow \infty$, the ratio between the matching size of Water-Filling and the optimal is

$$\lim_{m \rightarrow \infty} \frac{\sum_t (k - \|\mathbf{p}_t\|_1)}{m} = 1 - \frac{1}{k} \cdot \|\mathbf{p}^*\|_1.$$

Finally, we consider when $k \rightarrow \infty$ and calculate the stationary vector. In this case, \mathbf{p}^* becomes a function $p : [0, 1] \rightarrow [0, 1]$ and the linear equation $\mathbf{p}^* = \mathbf{M}(\mathbf{1} - \mathbf{p}^*)$ becomes the following

$$\int_0^{h^{-1}(x)} \frac{1 - p(y)}{1 - y + h(y)} dy = p(x), \quad \forall x \in [0, 1].$$

⁵Observe that $(\mathbf{p}_{t+1} - \mathbf{p}^*) = \mathbf{M}(\mathbf{p}_t - \mathbf{p}^*)$ and \mathbf{M} is a contraction matrix.

We verify that $p = \tau$ is a solution to this system of equations by Lemma 3.1. For all x , we have

$$\begin{aligned} \int_0^{h^{-1}(x)} \frac{1 - \tau(y)}{1 - y + h(y)} dy &= c - \tau(h^{-1}(x)) \\ &= \tau(f(c - \tau(h^{-1}(x)))) \\ &= \tau(h(h^{-1}(x))) = \tau(x). \end{aligned}$$

Thus, the ratio between Water-Filling and OPT is $1 - \int_0^1 \tau(y) dy = c = 2 - \sqrt{2}$.

Interestingly, we show that our hardness result applies to the edge-arrival models of the online matching problems. In the Online Edge Arrival Matching problem [BST17], at each step, an edge arrives online and the algorithm must irrevocably decide whether to add the edge to the matching; in the preemptive setting (Online Preemptive Matching [ELSW13, McG05]), instead, we are allowed to dispose of edges in the matching before accepting a new edge.

COROLLARY 3.2. *No algorithm can be better than $(2 - \sqrt{2})$ -competitive for Online Edge Arrival Matching and Online Preemptive Matching, even if fractional matching is allowed.*

Proof. Since the edge arrival model (resp. integral matching) is strictly harder than the preemptive model (resp. fractional matching), it suffices to consider the second model with fractional matching. Consider the previous hard instance with the following modifications. The underlying graph remains the same and each vertex is associated with the same deadline as before. At $u_{t,i}$'s deadline, its incident edges with available neighbors are revealed one by one. In this way, all available neighbors of $u_{t,i}$ are indistinguishable at this moment, i.e. they share the same set of neighbors. Thus by assigning random identities to these vertices, the available neighbors of $u_{t,i}$ have the same expected increment in matched fraction. Moreover, since no edge incident to each vertex comes after its deadline, it is not beneficial for an algorithm to dispose of previously chosen edges. Therefore, no algorithm can do better than Water-Filling in expectation and the lower bound $2 - \sqrt{2}$ applies.

4 Tight Competitive Ratio of Ranking on Bipartite Graphs

Let $\Omega \approx 0.5671$ denote the Omega constant, which is the solution for the equation $\Omega \cdot e^\Omega = 1$. In this section, we prove that Ranking is Ω -competitive for the Fully Online Bipartite Matching, matching the Ω hardness result given by Huang et al. [HKT⁺18].

THEOREM 4.1. *Ranking is Ω -competitive for Fully On-line Bipartite Matching.*

We adopt the randomized primal dual analysis from [HKT⁺18]. Recall the dual assignment that distributes the gain of each matched edge between its two endpoints as follows.

- *Gain Sharing:* Whenever a pair (u, v) is matched with u being active and v being passive, let $\alpha_u = 1 - g(y_v)$ and $\alpha_v = g(y_v)$, where $g : [0, 1] \rightarrow [0, 1]$ is non-decreasing, and $g(1) = 1$.

By Lemma 2.2, it suffices to prove that $\mathbf{E}_{\vec{y}}[\alpha_u + \alpha_v] \geq \Omega$ for all pairs of neighbors u, v . Suppose u has an earlier deadline than v and \vec{y}_{-v} is the rank vector of all vertices excluding v . Let θ be the marginal rank of v . The following lemma lies in the central of the proof by [HKT⁺18].

FACT 4.1. ([HKT⁺18], LEMMA 3.2) *For any arbitrarily fixed \vec{y}_{-v} , we have*

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

Our main technical contribution is an improved version of the above lower bound. Indeed, using Fact 4.1 as a lower bound, one cannot achieve a competitive ratio greater than 0.56 by just optimizing g .⁶ In the following, we will first illustrate how this lower bound can be improved for the hard instance given in [HKT⁺18]. Then, we show in Section 4.2 how to prove the Ω competitive ratio for general instances.

4.1 Better Competitive Ratio for the Hard Instance Recall the following hard instance for Ranking that is given by [HKT⁺18]. In the instance (refer to Figure 3), the vertices are organized into (infinitely many) groups of size $2k$, where each group $U_t \cup V_t$ induces a perfect matching. For all $t \in [m-1]$, the vertices U_t and U_{t+1} are connected by a complete bipartite graph. The deadline of every $u_{t,i}$ is earlier than $v_{t,i}$, and deadlines of $u_{t,i}$ follow the lexicographic order on (t, i) .

It is shown in [HKT⁺18] that when running Ranking on the above instance, at the deadline of the first vertex of each group, e.g., $u_{t,1}$, the expected fraction of unmatched vertices in U_t (which is also the competitive ratio of Ranking) is given by the equation $x = e^{-x}$. In




Figure 3: Hard instance of Ranking: illustrating example with $k = 4$.

other words, the competitive ratio of Ranking is Ω on the above instance (when $k \rightarrow \infty$).

In the following, we show that the competitive ratio of Ranking is Ω , using the randomized primal dual framework, and explain what is missing in the previous analysis. Fix any pair of neighbors u, v in the same group s.t. u has an earlier deadline than v . Next, we fix the ranks of all vertices but v arbitrarily, and lower bound $\mathbf{E}_{y_v} [\alpha_u + \alpha_v]$ for any edge (u, v) that appears in the perfect matching⁷.

Observe that u is the only neighbor of v . If u is passive, then v is unmatched regardless of y_v , which implies $\mathbf{E}_{y_v} [\alpha_u + \alpha_v] = g(y_u)$. Otherwise, let θ be the marginal rank of v . By definition, when $y_v > \theta$, u matches a vertex with rank θ and hence $\alpha_u = 1 - g(\theta)$. For the case when $y_v < \theta$, it is shown in [HKT⁺18] (using Lemma 2.1) that u does not get worse: u either is passive, or actively matches a vertex with rank at most θ . That is, $\alpha_u \geq \min\{g(y_u), 1 - g(\theta)\}$ when $y_v < \theta$. However, for the specific hard instance given in Figure 3, u is v 's only neighbor. Hence, u and v will match each other when $y_v < \theta$. Therefore, we have

$$\begin{aligned} \mathbf{E}_{y_v} [\alpha_u + \alpha_v] &= \int_0^\theta (\alpha_u + \alpha_v) dy_v + \int_\theta^1 \alpha_u dy_v \\ &= \theta + (1 - \theta) \cdot (1 - g(\theta)). \end{aligned}$$

Together with the case when u is passive, we have that

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

This bound is strictly stronger than Fact 4.1, as we fully characterize the gain of α_u when y_v is smaller than

⁶The function g is not optimized in [HKT⁺18] with respect to their lower bound. However, the ratio is less than 0.56 with the optimal g function.

⁷Note that the competitive ratio equals $\sum_u \mathbf{E} [\alpha_u] = \sum_{(u,v) \text{ appears in the perfect matching}} \mathbf{E} [\alpha_u + \alpha_v]$.

its marginal rank, rather than the loose lower bound $\min\{g(y_u), 1 - g(\theta)\}$ given in [HKT⁺18]. By taking expectation over y_u and optimizing the function $g(\cdot)$ (see Section 4.2), the above lower bound implies that Ranking is Ω -competitive on the hard instance.

In general, v does not necessarily match u when $y_v < \theta$. However, when this fails to happen, we are able to retrieve extra gain of α_v when u is passive. (Recall that in the hard instance, v is unmatched when u is passive.) The complete analysis involves a more careful treatment that considers the randomness of y_u, y_v at the same time, when deriving the lower bound.

4.2 Proof of Theorem 4.1 Consider any neighboring vertices u and v . In the following, we fix an arbitrary assignment of ranks to all vertices but u, v . We denote this assignment of ranks by \vec{y}_{-uv} . Unless otherwise specified, we use $\mathbf{E}[\cdot]$ to denote the expectation taken over the randomness of y_u and y_v .

Instead of using a single threshold θ of v as in the previous analysis, we will make use of multiple thresholds to give a good enough characterization of the matching status of u and v in order to derive the tight competitive ratio. We introduce the first two below.

DEFINITION 4.1. (τ AND γ) Consider the graph $G - \{v\}$ with v removed. Let τ be the marginal rank of u w.r.t. \vec{y}_{-uv} . In other words, u is passive iff $M(y_u < \tau, \vec{y}_{-uv})$. Similarly, let γ be the marginal rank of v w.r.t. \vec{y}_{-uv} in graph $G - \{u\}$, i.e., with u removed.

LEMMA 4.1. $\mathbf{E}[\alpha_u \cdot \mathbf{1}(y_u < \tau) + \alpha_v \cdot \mathbf{1}(y_v < \gamma)] = \int_0^\tau g(y_u) dy_u + \int_0^\gamma g(y_v) dy_v$

Proof. Consider $y_u = y < \tau$. By the definition of τ , we know that for all $y_v \in [0, 1]$, u is passive in $M(y_u = y, y_v, \vec{y}_{-uv})$, because inserting v (with any rank) to the graph cannot make u worse (by Lemma 2.1). Thus, for all $y_u < \tau$ and $y_v \in [0, 1]$, we have $\alpha_u = g(y_u)$, which correspond to the first term of the RHS. For the same reason, for all $y_u \in [0, 1]$, v is passive in $M(y_v < \gamma, y_u, \vec{y}_{-uv})$, which gives $\alpha_v = g(y_v)$, and the second term of the RHS.

For all $y_u \in [0, 1]$, let $\theta(y_u)$ be the marginal rank of v w.r.t. $\vec{y}_{-v} = (y_u, \vec{y}_{-uv})$. Recall v is always passive (regardless of y_u) when $y_v < \gamma$. Hence, we have $\theta(y_u) \geq \gamma$ for all $y_u \in [0, 1]$.

LEMMA 4.2. For any fixed $y_u > \tau$, we have

$$\mathbf{E}_{y_v} [\alpha_u + \alpha_v \cdot \mathbf{1}(y_v > \gamma)] \geq 1 - \gamma - (1 - \theta(y_u)) \cdot g(\theta(y_u)) + \gamma \cdot \min\{g(y_u), 1 - g(\theta(y_u))\}.$$

Proof. By the definition of $\theta(y_u)$, we know that when $y_v = \theta(y_u)^+$ (slightly larger than $\theta(y_u)$), v is not passive. Thus, u must be matched. Moreover, u must be active. Otherwise u should remain passive when v is removed, because the deadline of v is later than u , which contradicts the definition of τ (recall that we fix some $y_u > \tau$). Hence, when $y_v = \theta(y_u)^+$, u actively matches some vertex with rank at most $\theta(y_u)$. As increasing the rank of v does not create any difference to the final matching, for all $y_v > \theta(y_u)$, we have $\alpha_u \geq 1 - g(\theta(y_u))$.

Note that it is possible that $\theta(y_u) = 1$, i.e., v is passive for all rank $y_v \in [0, 1]$, in which case the above lower bound still holds. Since the graph is bipartite, by Lemma 2.1, for all $y_v < \theta(y_u)$, we have $\alpha_u \geq \min\{g(y_u), 1 - g(\theta(y_u))\}$.

Finally, we show that for any $y_v \in (\gamma, \theta(y_u))$, we have $\alpha_u + \alpha_v \geq 1$. Fix any $y_v \in (\gamma, \theta(y_u))$. By definition v is passive. Consider the first moment when one of u, v is matched.

Suppose at this moment, v is matched (passively) by some vertex z . Then, we show that $z = u$, which gives $\alpha_u + \alpha_v = 1$. Otherwise, z must have an earlier deadline than u . Then, we know that v remains passive with u removed, which contradicts the definition of γ .

Suppose at this moment, u is matched. Then we know that u must active, as otherwise u remains passive with v removed, which contradicts the definition of τ . Suppose u matches some vertex z . Since v is not matched at this moment, the rank of z is no more than y_v , which implies $\alpha_u \geq 1 - g(y_v) = 1 - \alpha_v$, as required.

To sum up, for any fixed $y_u > \tau$, we have

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

as claimed.

Combing the two lemmas, we have the following lower bound. Observe that the following bound degrades to the one we derived for the hard instance in Subsection 4.1, when $\gamma = 0$.

LEMMA 4.3. For any neighbor u of v that has an earlier

deadline than v , and for any \vec{y}_{-uv} , we have

$$\begin{aligned} & \mathbf{E} [\alpha_u + \alpha_v] \\ & \geq \min_{\tau, 0 \leq \gamma \leq \theta \leq 1} \left\{ \int_0^\tau g(y_u) dy_u + \int_0^\gamma g(y_v) dy_v \right. \\ & \quad + (1 - \tau)(1 - \gamma - (1 - \theta) \cdot g(\theta)) \\ & \quad \left. + \gamma \cdot \int_\tau^1 \min\{1 - g(\theta), g(y_u)\} dy_u \right\} \end{aligned}$$

Proof. First, we show that there exists θ such that $\theta(y_u) = \theta$ for all $y_u > \tau$. Consider the graph with v removed, and let $y_u = \tau^+$. By the definition of τ , u is not passive.

1. If u is unmatched, then we know that after inserting v with any $y_v \in [0, 1]$, v is passive, as otherwise u will be matched with v removed. Hence, we have $\theta(y_u) = 1$ for all $y_u > \tau$;
2. Otherwise, u is active. Let $\theta = \theta(\tau^+)$. Then, we know that v is not passive when inserted to the graph with $y_v = \theta^+$. Moreover, we know that u is active after the insertion: if u is passive, then u remains passive with v removed, which contradicts the definition of τ . Since increasing y_u does not change the matching, we have $\theta(y_u) \leq \theta$ for all $y_u > \tau$. On the other hand, when $y_v = \theta$ and $y_u = \tau^+$, u is active and v is passive. Since increasing y_u does not change the matching, we have $\theta(y_u) \geq \theta$ for all $y_u > \tau$. The sandwiching bounds imply that $\theta(y_u) = \theta$ for all $y_u > \tau$.

Hence, combining Lemma 4.1 and 4.2, we have

$$\begin{aligned} \mathbf{E} [\alpha_u + \alpha_v] & \geq \mathbf{E} [\alpha_u \cdot \mathbf{1}(y_u < \tau) + \alpha_v \cdot \mathbf{1}(y_v < \gamma)] \\ & \quad + \int_\tau^1 \mathbf{E}_{y_v} [\alpha_u + \alpha_v \cdot \mathbf{1}(y_v > \gamma)] dy_u \\ & = \int_0^\tau g(y_u) dy_u + \int_0^\gamma g(y_v) dy_v \\ & \quad + \int_\tau^1 \left(1 - \gamma - (1 - \theta) \cdot g(\theta) \right. \\ & \quad \left. + \gamma \cdot \min\{g(y_u), 1 - g(\theta)\} \right) dy_u \\ & = \int_0^\tau g(y_u) dy_u + \int_0^\gamma g(y_v) dy_v \\ & \quad + (1 - \tau) \cdot (1 - \gamma - (1 - \theta) \cdot g(\theta)) \\ & \quad + \gamma \cdot \int_\tau^1 \min\{g(y_u), 1 - g(\theta)\} dy_u. \end{aligned}$$

Taking minimum over τ and $\gamma \leq \theta$ gives Lemma 4.3.

Proof of Theorem 4.1: Fix the non-decreasing function g as follows:

$$g(y) = \begin{cases} \frac{c}{1-y}, & \text{when } y < \frac{1-2c}{1-c}, \\ 1-c, & \text{when } \frac{1-2c}{1-c} \leq y < 1, \\ 1, & \text{when } y = 1, \end{cases}$$

where $c = \frac{1}{1+e^\Omega} \approx 0.3619$. Let $f(\tau, \gamma, \theta)$ denote the expression to be minimized on the RHS of Lemma 4.3. Then, we have

$$\begin{aligned} f(\tau, \gamma, \theta) & = \int_0^\tau g(y_u) dy_u + \int_0^\gamma g(y_v) dy_v \\ & \quad + (1 - \tau) \left(1 - \gamma - (1 - \theta) \cdot g(\theta) \right) \\ & \quad + \gamma \cdot \int_\tau^1 \min\{g(y_u), 1 - g(\theta)\} dy_u. \end{aligned}$$

Fix any γ and θ , and suppose $g(\tau) < 1 - g(\theta)$, then observe that

$$\begin{aligned} \frac{\partial f(\tau, \gamma, \theta)}{\partial \tau} & = g(\tau) - (1 - \gamma - (1 - \theta) \cdot g(\theta)) \\ & \quad - \gamma \cdot \min\{g(\tau), 1 - g(\theta)\} \\ & = (1 - \gamma) \cdot g(\tau) - (1 - \gamma) + (1 - \theta) \cdot g(\theta) \\ & < (1 - \gamma)(1 - g(\theta)) - (1 - \gamma) + (1 - \theta) \cdot g(\theta) \\ & = (\gamma - \theta) \cdot g(\theta) \leq 0. \end{aligned}$$

Here, the last inequality holds because we have $\theta \geq \gamma$ by their definitions.

Thus, the minimum of $f(\tau, \gamma, \theta)$ over $\tau \in [0, 1]$, $0 \leq \gamma \leq \theta \leq 1$ must be obtained when $g(\tau) \geq 1 - g(\theta)$. As a result, we get that

$$\begin{aligned} f(\tau, \gamma, \theta) & = \int_0^\tau g(y_u) dy_u + \int_0^\gamma g(y_v) dy_v \\ & \quad + (1 - \tau) \left(1 - (1 - \theta + \gamma) \cdot g(\theta) \right). \end{aligned}$$

If we relax the constraint that $\theta \geq \gamma$, then the maximum of $(1 - \theta + \gamma)g(\theta)$ is achieved when $\theta^* = \frac{1-2c}{1-c}$ (for which $g(\theta^*) = 1 - c$). Note that the maximum is $(\frac{c}{1-c} + \gamma) \cdot (1 - c) = (1 - c) \cdot \gamma + c$, which is greater than the value of expression when $\theta = 1$, i.e., γ . Thus, we have

$$\begin{aligned} f(\tau, \gamma, \theta) & \geq f(\tau, \gamma, \theta^*) = \int_0^\tau g(y_u) dy_u + \int_0^\gamma g(y_v) dy_v \\ & \quad + (1 - \tau) \cdot (1 - \gamma) \cdot (1 - c). \end{aligned}$$

It is easy to see that the minimum of $f(\tau, \gamma, \theta^*)$ must be achieved when $\gamma < \theta^* = \frac{1-2c}{1-c}$ (for which $g(\gamma) < 1 - c$), as otherwise the partial derivative

$$\frac{\partial f(\tau, \gamma, \theta^*)}{\partial \gamma} = g(\gamma) - (1 - \tau) \cdot (1 - c) \geq 0.$$

Since $f(\tau, \gamma, \theta^*)$ is symmetric for τ and γ , the same conclusion holds for τ , which means

$$\begin{aligned} f(\tau, \gamma, \theta^*) &= \int_0^\tau \frac{c}{1-x} dx + \int_0^\gamma \frac{c}{1-x} dx \\ &\quad + (1-\tau) \cdot (1-\gamma) \cdot (1-c) \\ &= -c \ln(1-\tau) - c \ln(1-\gamma) \\ &\quad + (1-\tau) \cdot (1-\gamma) \cdot (1-c) \\ &\geq c - c \cdot \ln\left(\frac{c}{1-c}\right) = \frac{1+\Omega}{1+e^\Omega} = \Omega, \end{aligned}$$

where the inequality comes from the fact that (take $(1-\tau) \cdot (1-\gamma)$ as the variable) function $(1-c) \cdot x - c \cdot \ln(x)$ achieves its minimum when $x = \frac{c}{1-c}$. ■

References

- [ABJS18] Itai Ashlagi, Maximilien Burq, Patrick Jaillet, and Amin Saberi. Maximizing efficiency in dynamic matching markets. *CoRR*, abs/1803.01285, 2018.
- [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.
- [DHK⁺13] Nikhil R. Devanur, Zhiyi Huang, Nitish Korula, Vahab S. Mirrokni, and Qiqi Yan. Whole-page optimization and submodular welfare maximization with online bidders. In *EC*, pages 305–322. ACM, 2013.
- [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.
- [DS18] Chinmoy Dutta and Chris Sholley. Online matching in a ride-sharing platform. *CoRR*, abs/1806.10327, 2018.
- [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.
- [HKT⁺18] Zhiyi Huang, Ning Kang, Zhihao Gavin Tang, Xiaowei Wu, Yuhao Zhang, and Xue Zhu. How to match when all vertices arrive online. In *STOC*, pages 17–29. ACM, 2018.
- [HTWZ18] Zhiyi Huang, Zhihao Gavin Tang, Xiaowei Wu, and Yuhao Zhang. Online vertex-weighted bipartite matching: Beating $1-1/e$ with random arrivals. In *ICALP*, volume 107 of *LIPICs*, pages 79:1–79:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018.
- [KMT11] Chinmay Karande, Aranyak Mehta, and Pushkar Tripathi. Online bipartite matching with unknown distributions. In *STOC*, pages 587–596, 2011.
- [KP00] Bala Kalyanasundaram and Kirk Pruhs. An optimal deterministic algorithm for online b-matching. *Theor. Comput. Sci.*, 233(1-2):319–325, 2000.
- [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.
- [WW15] Yajun Wang and Sam Chiu-wai Wong. Two-sided online bipartite matching and vertex cover: Beating the greedy algorithm. In *ICALP (1)*, volume 9134 of *Lecture Notes in Computer Science*, pages 1070–1081. Springer, 2015.

A Primal-Dual Connection between the Upper and Lower Bounds

We provide an interesting primal-dual connection between the primal dual analysis in Section 3.1 and the hard instance in Section 3.2, which inspires us to find the correct h function in Section 3.2.

Recall the following lower bound established in the proof of Theorem 3.1,

$$\alpha_u + \alpha_v \geq \min \left\{ \int_0^1 g(x) dx, \min_{p_u, x_v} \left\{ \int_0^{p_u} g(x) dx + (1-p_u)(1-g(x_v)) + \int_0^{x_v} g(x) dx \right\} \right\}.$$

We are left to optimize function g using the following

linear program:

$$\begin{aligned} \max_g : \quad & r \\ \text{s.t.} \quad & r \leq \int_0^x g(y)dy + \int_0^z g(y)dy \\ & + (1-x)(1-g(z)), \quad \forall x, z \in [0, 1]. \end{aligned}$$

After solving it, we remove redundant constraints with slacks and consider the following program:

$$(P) \quad \begin{aligned} \max_g : \quad & r \\ \text{s.t.} \quad & r \leq \int_0^x g(y)dy + \int_0^{c-x} g(y)dy \\ & + (1-x)(1-g(c-x)), \quad \forall x \in [0, c], \end{aligned}$$

where $c = 2 - \sqrt{2}$. We know that the above two programs have the same optimal value. Moreover, as the program suggests, in order to construct a tight hard instance, all pairs u, v matched in OPT must satisfy $p_u + x_v = c$ when Water-Filling is run⁸. According to the instance structure and argument in Section 3.2, it suffices to find a function $h : [0, 1] \rightarrow [0, 1]$ so that

$$\int_0^x \frac{1 - \tau(y)}{1 - y + h(y)} dy = \tau(h(x)) = c - \tau(x), \quad \forall x \in [0, 1].$$

Here, $\tau : [0, 1] \rightarrow [0, c]$ corresponds to the stationary water level and gives the first equation. Moreover, the perfect partner corresponding to x also has water level $\tau(h(x))$ and we require it to be $c - \tau(x)$, which gives the second equation. Therefore, $h(x) = \tau^{-1}(c - \tau(x))$. Let $f(x) = \tau^{-1}(x)$ and taking derivative over the above equation, it suffices to prove the existence of f, τ, h so that

$$\begin{aligned} \frac{1 - \tau(x)}{1 - x + h(x)} = -\tau'(x) &\Leftrightarrow \frac{1 - \phi}{1 - f(\phi) + f(c - \phi)} = -\frac{1}{f'(\phi)} \\ \Leftrightarrow 1 - f(\phi) + f(c - \phi) + (1 - \phi)f'(\phi) &= 0, \quad \forall \phi \in [0, c]. \end{aligned}$$

Now, consider the dual program of P :

$$\begin{aligned} \min_q : \quad & \int_0^c (1-x)q(x)dx \\ \text{s.t.} \quad & 1 - \int_0^c q(x)dx \leq 0, \\ & \int_0^x q(y)dy + \int_{c-x}^c q(y)dy - (1-x)q(x) \leq 0, \forall x \in [0, c]. \end{aligned}$$

According to primal dual theory, we know that the optimal dual solution $q(x)$ satisfies $\int_0^c q(x)dx = 1$ and

⁸Constraints must be tight almost everywhere. As otherwise, our primal dual analysis proves the competitive ratio of Water-Filling is strictly greater than c on the specific instance.

$\int_0^x q(y)dy + \int_{c-x}^c q(y)dy + (1-x)q(x) = 0$. Let $Q(x) = 1 - \int_0^x q(y)dy = \int_x^c q(y)dy$, we have

$$1 - Q(x) + Q(c-x) + (1-x)Q'(x) = 0, \quad \forall x \in [0, c],$$

which is exactly the same equation we required for f .