

Online Matching with General Arrivals

Buddhima Gamlath
EPFL

buddhima.gamlath@epfl.ch

Michael Kapralov*
EPFL

michael.kapralov@epfl.ch

Andreas Maggiori
EPFL

andreas.maggiori@epfl.ch

Ola Svensson
EPFL

ola.svensson@epfl.ch

David Wajc†
CMU

dwajc@cs.cmu.edu

Abstract

The online matching problem was introduced by Karp, Vazirani and Vazirani nearly three decades ago. In that seminal work, they studied this problem in bipartite graphs with vertices arriving only on one side, and presented optimal deterministic and randomized algorithms for this setting. In comparison, more general arrival models, such as edge arrivals and general vertex arrivals, have proven more challenging and positive results are known only for various relaxations of the problem. In particular, even the basic question of whether randomization allows one to beat the trivially-optimal deterministic competitive ratio of $1/2$ for either of these models was open. In this paper, we resolve this question for both these natural arrival models, and show the following.

1. For edge arrivals, randomization does not help — no randomized algorithm is better than $1/2$ competitive.
2. For general vertex arrivals, randomization helps — there exists a randomized $(1/2 + \Omega(1))$ -competitive online matching algorithm.

*Supported in part by ERC Starting Grant 759471.

†Work done in part while the author was visiting EPFL. Supported in part by NSF grants CCF-1618280, CCF-1814603, CCF-1527110, NSF CAREER award CCF-1750808, and a Sloan Research Fellowship.

1 Introduction

Matching theory has played a prominent role in the area of combinatorial optimization, with many applications [21, 23]. Moreover, many fundamental techniques and concepts in combinatorial optimization can trace their origins to its study, including the primal-dual framework [19], proofs of polytopes’ integrality beyond total unimodularity [8], and even the equation of efficiency with polytime computability [9].

Given the prominence of matching theory in combinatorial optimization, it comes as little surprise that the maximum matching problem was one of the first problems studied from the point of view of online algorithms and competitive analysis. In 1990, Karp, Vazirani, and Vazirani [18] introduced the online matching problem, and studied it under one-sided bipartite arrivals. For such arrivals, Karp et al. noted that the trivial $1/2$ -competitive greedy algorithm (which matches any arriving vertex to an arbitrary unmatched neighbor, if one exists) is optimal among deterministic algorithms for this problem. More interestingly, they provided an elegant randomized online algorithm for this problem, called RANKING, which achieves an optimal $(1 - 1/e)$ competitive ratio. (This bound has been re-proven many times over the years [2, 6, 7, 11, 13].) Online matching and many extensions of this problem under one-sided bipartite vertex arrivals were widely studied over the years, both under adversarial and stochastic arrival models. See recent work [5, 15, 16, 17] and the excellent survey of Mehta [22] for further references on this rich literature.

Despite our increasingly better understanding of one-sided online bipartite matching and its extensions, the problem of online matching under more general arrival models, including edge arrivals and general vertex arrivals, has remained staunchly defiant, resisting attacks. In particular, the basic questions of whether the trivial $1/2$ competitive ratio is optimal for the adversarial edge-arrival and general vertex-arrival models have remained tantalizing open questions in the online algorithms literature. In this paper, we answer both of these questions.

1.1 Prior Work and Our Results

Here we outline the most relevant prior work, as well as our contributions. Throughout, we say an algorithm (either randomized or fractional) has *competitive ratio* α , or equivalently is α -*competitive*, if the ratio of the algorithm’s value (e.g., expected matching size, or overall value, $\sum_e x_e$) to OPT is at least $\alpha \leq 1$ for all inputs and arrival orders. As is standard in the online algorithms literature on maximization problems, we use upper bounds (on α) to refer to hardness results, and lower bounds to positive results.

Edge Arrivals. Arguably the most natural, and the least restricted, arrival model for online matching is the edge arrival model. In this model, edges are revealed one by one, and an online matching algorithm must decide immediately and irrevocably whether to match the edge on arrival, or whether to leave both endpoints free to be possibly matched later.

On the hardness front, the problem is known to be strictly harder than the one-sided vertex arrival model of Karp et al. [18], which admits a competitive ratio of $1 - 1/e \approx 0.632$. In particular, Epstein et al. [10] gave an upper bound of $\frac{1}{1+\ln 2} \approx 0.591$ for this problem, recently improved by Huang et al. [16] to $2 - \sqrt{2} \approx 0.585$. (Both bounds apply even to online algorithms with preemption; i.e., allowing edges to be removed from the matching in favor of a newly-arrived edge.) On the positive side, as pointed out by Buchbinder et al. [3], the edge arrival model has proven challenging, and results beating the $1/2$ competitive ratio were only achieved under various relaxations, including: random order edge arrival [14], bounded number of arrival batches [20], on trees, either with or without preemption [3, 24], and for bounded-degree graphs [3]. The above papers all asked whether

there exists a randomized $(1/2 + \Omega(1))$ -competitive algorithm for adversarial edge arrivals (see also Open Question 17 in Mehta’s survey [22]).

In this work, we answer this open question, providing it with a strong negative answer. In particular, we show that no online algorithm for *fractional* matching (i.e., an algorithm which immediately and irrevocably assigns values x_e to edge e upon arrival such that \vec{x} is in the fractional matching polytope $\mathcal{P} = \{\vec{x} \geq \vec{0} \mid \sum_{e \ni v} x_e \leq 1 \forall v \in V\}$) is better than $1/2$ competitive. As any randomized algorithm induces a fractional algorithm with the same competitive ratio, this rules out any randomized online matching algorithm which is better than deterministic algorithms.

Theorem 1.1. *No fractional online algorithm is $1/2 + \Omega(1)$ competitive for online matching under adversarial edge arrivals, even in bipartite graphs.*

This result shows that the study of relaxed variants of online matching under edge arrivals is not only justified by the difficulty of beating the trivial bound for this problem, but rather by its *impossibility*.

General Vertex Arrivals. In the online matching problem under vertex arrivals, vertices are revealed one at a time, together with their edges to their previously-revealed neighbors. An online matching algorithm must decide immediately and irrevocably upon arrival of a vertex whether to match it (or keep it free for later), and if so, who to match it to. The one-sided bipartite problem studied by Karp et al. [18] is precisely this problem when all vertices of one side of a bipartite graph arrive first. As discussed above, for this one-sided arrival model, the problem is thoroughly understood (even down to lower-order error terms [11]). Wang and Wong [25] proved that general vertex arrivals are strictly harder than one-sided bipartite arrivals, providing an upper bound of $0.625 < 1 - 1/e$ for the more general problem, later improved by Buchbinder et al. [3] to $\frac{2}{3+\phi^2} \approx 0.593$. Clearly, the general vertex arrival model is no harder than the online edge arrival model but is it *easier*? The answer is “yes” for *fractional* algorithms, as shown by combining our Theorem 1.1 with the 0.526-competitive fractional online matching algorithm under general vertex arrivals of Wang and Wong [25]. For *integral* online matching, however, the problem has proven challenging, and the only positive results for this problem, too, are for various relaxations, such as restriction to trees, either with or without preemption [3, 4, 24], for bounded-degree graphs [3], or (recently) allowing vertices to be matched during some known time interval [15, 16].

We elaborate on the last relaxation above. In the model recently studied by Huang et al. [15, 16] vertices have both arrival and departure times, and edges can be matched whenever both their endpoints are present. (One-sided vertex arrivals is a special case of this model with all online vertices departing immediately after arrival and offline vertices departing at ∞ .) We note that any α -competitive online matching under general vertex arrivals is α -competitive in the less restrictive model of Huang et al. As observed by Huang et al., for their model an optimal approach might as well be greedy; i.e., an unmatched vertex v should always be matched at its departure time if possible. In particular, Huang et al. [15, 16], showed that the RANKING algorithm of Karp et al. achieves a competitive ratio of ≈ 0.567 . For general vertex arrivals, however, RANKING (and indeed any maximal matching algorithm) is no better than $1/2$ competitive, as is readily shown by a path on three edges with the internal vertices arriving first. Consequently, new ideas and algorithms are needed.

The natural open question for general vertex arrivals is whether a competitive ratio of $(1/2 + \Omega(1))$ is achievable by an *integral* randomized algorithm, without any assumptions (see e.g., [25]). In this work, we answer this question in the affirmative:

Theorem 1.2. *There exists a $(1/2 + \Omega(1))$ -competitive randomized online matching algorithm for general adversarial vertex arrivals.*

1.2 Our Techniques

Edge Arrivals. All prior upper bounds in the online literature [3, 10, 11, 16, 18] can be rephrased as upper bounds for *fractional* algorithms; i.e., algorithms which immediately and irrevocably assign each edge e a value x_e on arrival, so that \vec{x} is contained in the fractional matching polytope, $\mathcal{P} = \{\vec{x} \geq \vec{0} \mid \sum_{e \ni v} x_e \leq 1 \forall v \in V\}$. With the exception of [3], the core difficulty of these hard instances is uncertainty about “identity” of vertices (in particular, which vertices will neighbor which vertices in the following arrivals). Our hardness instances rely on uncertainty about the “time horizon”. In particular, the underlying graph, vertex identifiers, and even arrival order are known to the algorithm, but the number of edges of the graph to be revealed (to arrive) is uncertain. Consequently, an α -competitive algorithm must accrue high enough value up to each arrival time to guarantee a high competitive ratio at all points in time. As we shall show, for competitive ratio $1/2 + \Omega(1)$, this goal is at odds with the fractional matching constraints, and so such a competitive ratio is impossible. In particular, we provide a family of hard instances and formulate their prefix-competitiveness and matching constraints as linear constraints to obtain a linear program whose objective value bounds the optimal competitive ratio. Solving the obtained LP’s dual, we obtain by weak duality the claimed upper bound on the optimal competitive ratio.

General Vertex Arrivals. Our high-level approach here will be to round online a fractional online matching algorithm’s output, specifically that of Wang and Wong [25]. While this approach sounds simple, there are several obstacles to overcome. First, the fractional matching polytope is not integral in general graphs, where a fractional matching may have value, $\sum_e x_e$, some $3/2$ times larger than the optimal matching size. (For example, in a triangle graph with value $x_e = 1/2$ for each edge e .) Therefore, any general rounding scheme must lose a factor of $3/2$ on the competitive ratio compared to the fractional algorithm’s value, and so to beat a competitive ratio of $1/2$ would require an online fractional matching with competitive ratio $> 3/4 > 1 - 1/e$, which is impossible. To make matters worse, even in bipartite graphs, for which the fractional matching polytope is integral and offline lossless rounding is possible [1, 12], *online* lossless rounding of fractional matchings is impossible, even under one-sided vertex arrivals [5].

Despite these challenges, we show that a slightly better than $1/2$ -competitive fractional matching computed by the algorithm of [25] can be rounded online without incurring too high a loss, yielding $(1/2 + \Omega(1))$ -competitive randomized algorithm for online matching under general vertex arrivals.

To outline our approach, we first consider a simple method to round matchings online. When vertex v arrives, we pick an edge $\{u, v\}$ with probability $z_u = x_{uv} / \Pr[u \text{ free when } v \text{ arrives}]$, and add it to our matching if u is free.

If $\sum_u z_u \leq 1$, this allows us to pick at most one edge per vertex and have each edge $e = \{u, v\}$ be in our matching with the right marginal probability, x_e , resulting in a lossless rounding. Unfortunately, we know of no better-than- $1/2$ -competitive fractional algorithm for which this rounding guarantees $\sum_u z_u \leq 1$.

However, we observe that, for the correct set of parameters, the fractional matching algorithm of Wang and Wong [25] makes $\sum_u z_u$ close to one, while still ensuring a better-than- $1/2$ -competitive fractional solution. Namely, as we elaborate later in Section 3.3, we set the parameters of their algorithm so that $\sum_u z_u \leq 1 + O(\varepsilon)$, while retaining a competitive ratio of $1/2 + O(\varepsilon)$. Now consider the same rounding algorithm with normalized probabilities: I.e., on v ’s arrival, sample a neighbor u with probability $z'_u = z_u / \max\{1, \sum_u z_u\}$ and match if u is free. As the sum of z_u ’s is slightly above one in the worst case, this approach does not drastically reduce the competitive ratio. But the normalization factor is still too significant compared to the competitive ratio of the fractional solution, driving the competitive ratio of the rounding algorithm slightly below $1/2$.

To account for this minor yet significant loss, we therefore augment the simple algorithm by allowing it, with small probability (e.g., say $\sqrt{\varepsilon}$), to sample a second neighbor u_2 for each arriving vertex v , again with probabilities proportional to z'_{u_2} : If the first sampled choice, u_1 , is free, we match v to u_1 . Otherwise, if the second choice, u_2 , is free, we match v to u_2 . What is the marginal probability that such an approach matches an incoming vertex v to a given neighbor u ? Letting F_u denote the event that u is free when v arrives, this probability is precisely

$$\Pr[F_u] \cdot \left(z'_u + z'_u \cdot \sqrt{\varepsilon} \cdot \sum_w z'_w \cdot (1 - \Pr[F_w | F_u]) \right). \quad (1)$$

Here the first term in the parentheses corresponds to the probability that v matches to u via the first choice, and the second term corresponds to the same happening via the second choice (which is only taken when the first choice fails).

Ideally, we would like (1) to be at least x_{uv} for all edges, which would imply a lossless rounding. However, as mentioned earlier, this is difficult and in general impossible to do, even in much more restricted settings including one-sided bipartite vertex arrivals. We therefore settle for showing that (1) is at least $x_{uv} = \Pr[F_u] \cdot z_u$ for *most* edges (weighted by x_{uv}). Even this goal, however, is challenging and requires a nontrivial understanding of the correlation structure of the random events F_u . To see this, note that for example if the F_w events are perfectly positively correlated, i.e., $\Pr[F_w | F_u] = 1$, then the possibility of picking e as a second edge does not increase this edge's probability of being matched *at all* compared to if we only picked a single edge per vertex. This results in e being matched with probability $\Pr[F_u] \cdot z'_u = \Pr[F_u] \cdot z_u / \sum_w z_w = x_{uv} / \sum_w z_w$, which does not lead to any gain over the $1/2$ competitive ratio of greedy. Such problems are easily shown not to arise if all F_u variables are independent or negatively correlated. Unfortunately, positive correlation does arise from this process, and so we need to control these positive correlations.

The core of our analysis is therefore dedicated to showing that even though positive correlations do arise, they are by and large rather weak. Our main technical contribution consists of developing techniques for bounding such positive correlations. The idea behind the analysis is to consider the primary choices and secondary choices of vertices as defining a graph, and showing that after a natural pruning operation that reflects the structure of dependencies, most vertices are most often part of a very small connected component in the graph. The fact that connected components are typically very small is exactly what makes positive correlations weak and results in the required lower bound on (1) for most edges (in terms of x -value), which in turn yields our $1/2 + \Omega(1)$ competitive ratio.

2 Edge Arrivals

In this section we prove the asymptotic optimality of the greedy algorithm for online matching under adversarial edge arrivals. As discussed briefly in Section 1, our main idea will be to provide a “prefix hardness” instance, where an underlying input and the arrival order is known to the online matching algorithm, but the prefix of the input to arrive (or “termination time”) is not. Consequently, the algorithm must accrue high enough value up to each arrival time, to guarantee a high competitive ratio at all points in time. As we show, the fractional matching constraints rule out a competitive ratio of $1/2 + \Omega(1)$ even in this model where the underlying graph is known.

Theorem 2.1. *There exists an infinite family of bipartite graphs with maximum degree n and edge arrival order for which any online matching algorithm is at best $\left(\frac{1}{2} + \frac{1}{2n+2}\right)$ -competitive.*

Proof. We will provide a family of graphs for which no fractional online matching algorithm has better competitive ratio. Since any randomized algorithm induces a fractional matching algorithm, this immediately implies our claim. The n^{th} graph of the family, $G_n = (U, V, E)$, consists of a bipartite graph with $|U| = |V| = n$ vertices on either side. We denote by $u_i \in U$ and $v_i \in V$ the i^{th} node on the left and right side of G_n , respectively. Edges are revealed in n discrete rounds. In round $i = 1, 2, \dots, n$, the edges of a perfect matching between the first i left and right vertices arrive in some order. I.e., a matching of u_1, u_2, \dots, u_i and v_1, v_2, \dots, v_i is revealed. Specifically, edges (u_j, v_{i-j+1}) for all $i \geq j$ arrive. (See Figure 1 for example.) Intuitively, the difficulty for an algorithm attempting to assign much value to edges of OPT is that the (unique) maximum matching OPT changes every round, and no edge ever re-enters OPT .




Figure 1: G_5 , together with arrival order. Edges of current (prior) round are solid (dashed).

Consider some α -competitive fractional algorithm \mathcal{A} . We call the edge of a vertex w in the (unique) maximum matching of the subgraph of G_n following round i the i^{th} edge of w . For $i \geq j$, denote by $x_{i,j}$ the value \mathcal{A} assigns to the i^{th} edge of vertex u_j (and of v_{i-j+1}); i.e., to (u_j, v_{i-j+1}) . By feasibility of the fractional matching output by \mathcal{A} , we immediately have that $x_{i,j} \geq 0$ for all i, j , as well as the following matching constraints for u_j and v_j . (For the latter, note that the i^{th} edge of v_{i-j+1} is assigned value $x_{i,j} = x_{i, i-(i-j+1)+1}$ and so the i^{th} edge of v_j is assigned value $x_{i, i-j+1}$).

$$\sum_{i=j}^n x_{i,j} \leq 1. \quad (u_j \text{ matching constraint}) \quad (2)$$

$$\sum_{i=j}^n x_{i, i-j+1} \leq 1. \quad (v_j \text{ matching constraint}) \quad (3)$$

On the other hand, as \mathcal{A} is α -competitive, we have that after some k^{th} round – when the maximum matching has cardinality k – algorithm \mathcal{A} 's fractional matching must have value at least $\alpha \cdot k$. (Else an adversary can stop the input after this round, leaving \mathcal{A} with a worse than α -competitive matching.) Consequently, we have the following competitiveness constraints.

$$\sum_{i=1}^k \sum_{j=1}^i x_{i,j} \geq \alpha \cdot k \quad \forall k \in [n]. \quad (4)$$

Combining constraints (2), (3) and (4) together with the non-negativity of the $x_{i,k}$ yields the following linear program, $LP(n)$, whose optimal value upper bounds any fractional online matching algorithm's competitiveness on G_n , by the above.

$$\begin{aligned}
& \text{maximize} && \alpha \\
& \text{subject to:} && \sum_{i=j}^n x_{i,j} \leq 1 && \forall j \in [n] \\
& && \sum_{i=j}^n x_{i,i-j+1} \leq 1 && \forall j \in [n] \\
& && \sum_{i=1}^k \sum_{j=1}^i x_{i,j} \geq \alpha \cdot k && \forall k \in [n] \\
& && x_{i,j} \geq 0 && \forall i, j \in [n].
\end{aligned}$$

To bound the optimal value of $\text{LP}(n)$, we provide a feasible solution its LP dual, which we denote by $\text{Dual}(n)$. By weak duality, any dual feasible solution's value upper bounds the optimal value of $\text{LP}(n)$, which in turn upper bounds the optimal competitive ratio. Using the dual variables ℓ_j, r_j for the degree constraints of the j^{th} left and right vertices respectively (u_j and v_j) and dual variable c_k for the competitiveness constraint of the k^{th} round, we get the following dual linear program. Recall here again that $x_{i,i-j+1}$ appears in the matching constraint of v_j , with dual variable r_j , and so $x_{i,j} = x_{i,i-(i-j+1)+1}$ appears in the same constraint for v_{i-j+1} .

$$\begin{aligned}
& \text{minimize} && \sum_{j=1}^n (\ell_j + r_j) \\
& \text{subject to:} && \sum_{k=1}^n k \cdot c_k \geq 1 \\
& && \ell_j + r_{i-j+1} - \sum_{k=i}^n c_k \geq 0 && \forall i \in [n], j \in [i] \\
& && \ell_j, r_j, c_k \geq 0 && \forall j, k \in [n].
\end{aligned}$$

We provide the following dual solution.

$$\begin{aligned}
c_k &= \frac{2}{n(n+1)} && \forall k \in [n] \\
\ell_j = r_j &= \begin{cases} \frac{n-2(j-1)}{n(n+1)} & \text{if } j \leq n/2 + 1 \\ 0 & \text{if } n/2 + 1 < j \leq n. \end{cases}
\end{aligned}$$

We start by proving feasibility of this solution. The first constraint is satisfied with equality. For the second constraint, as $\sum_{k=i}^n c_k = \frac{2(n-i+1)}{n(n+1)}$ it suffices to show that $\ell_j + r_{i-j+1} \geq \frac{2(n-i+1)}{n(n+1)}$ for all $i \in [n], j \in [i]$. Note that if $j > n/2 + 1$, then $\ell_j = r_j = 0 > \frac{n-2(j-1)}{n(n+1)}$. So, for all j we have $\ell_j = r_j \geq \frac{n-2(j-1)}{n(n+1)}$. Consequently, $\ell_j + r_{i-j+1} \geq \frac{n-2(j-1)}{n(n+1)} + \frac{n-2(i-j+1-1)}{n(n+1)} = \frac{2(n-i+1)}{n(n+1)}$ for all $i \in [n], j \in [i]$. Non-negativity of the ℓ_j, r_j, c_k variables is trivial, and so we conclude that the above is a feasible dual solution.

It remains to calculate this dual feasible solution's value. We do so for n even,¹ for which

$$\sum_{j=1}^n (\ell_j + r_j) = 2 \cdot \sum_{j=1}^n \ell_j = 2 \cdot \sum_{j=1}^{n/2+1} \frac{n-2(j-1)}{n(n+1)} = \frac{1}{2} + \frac{1}{2n+2},$$

completing the proof. \square

Remark 1. Recall that Buchbinder et al. [3] and Lee and Singla [20] presented better-than- $1/2$ -competitive algorithms for bounded-degree graphs and bounded number of arrival batches. Our upper bound above shows that a deterioration of the competitive guarantees as the maximum degree and number of arrival batches increase (as in the algorithms of [3, 20]) is inevitable.

Remark 2. Recall that the *asymptotic* competitive ratio of an algorithm is the maximum c such that the algorithm always guarantees value at least $ALG \geq c \cdot OPT - b$ for some fixed $b > 0$. Our proof extends to this weaker notion of competitiveness easily, by revealing multiple copies of the hard family of Theorem 2.1 and letting x_{ik} denote the average of its counterparts over all copies.

¹The case of n odd is similar. As it is unnecessary to establish the result of this theorem, we omit it.

3 General Vertex Arrivals

In this section we present a $(1/2 + \Omega(1))$ -competitive randomized algorithm for online matching under general arrivals. As discussed in the introduction, our approach will be to round (online) a *fractional* online matching algorithm's output. Specifically, this will be an algorithm from the family of fractional algorithms introduced in [25]. In Section 3.1 we describe this family of algorithms. To motivate our rounding approach, in Section 3.2 we first present a simple lossless rounding method for a $1/2$ -competitive algorithm in this family. In Section 3.3 we then describe our rounding algorithm for a better-than- $1/2$ -competitive algorithm in this family. Finally, in Section 3.4 we analyze this rounding scheme, and show that it yields a $(1/2 + \Omega(1))$ -competitive algorithm.

3.1 Finding a fractional solution

In this section we revisit the algorithm of Wang and Wong [25], which beats the $1/2$ competitiveness barrier for online fractional matching under general vertex arrivals. Their algorithm (technically, family of algorithms) applies the primal-dual method to compute both a fractional matching and a fractional vertex cover – the dual of the fractional matching relaxation. The LPs defining these dual problems are as follows.

Primal-Matching	Dual-Vertex Cover
maximize $\sum_{e \in E} x_e$ subject to: $\sum_{u \in N(v)} x_{uv} \leq 1 \quad \forall u \in V$ $x_e \geq 0 \quad \forall e \in E$	minimize $\sum_{u \in V} y_u$ subject to: $y_u + y_v \geq 1 \quad \forall e = \{u, v\} \in E$ $y_u \geq 0 \quad \forall u \in V$

Before introducing the algorithm of [25], we begin by defining the fractional online vertex cover problem for vertex arrivals. When a vertex v arrives, if $N_v(v)$ denotes the previously-arrived neighbors of v , then for each $u \in N_v(v)$, a new constraint $y_u + y_v \geq 1$ is revealed, which an online algorithm should satisfy by possibly increasing y_u or y_v . Suppose v has its dual value set to $y_v = 1 - \theta$. Then all of its neighbors should have their dual increased to at least θ . Indeed, an algorithm may as well increase y_u to $\max\{y_u, \theta\}$. The choice of θ therefore determines an online fractional vertex cover algorithm. The increase of potential due to the newly-arrived vertex v is thus $1 - \theta + \sum_{u \in N_v(v)} (\theta - y_u)^+$.² In [25] θ is chosen to upper bound this term by $1 - \theta + f(\theta)$ for some function $f(\cdot)$. The primal solution (fractional matching) assigns values x_{uv} so as to guarantee feasibility of \vec{x} and a ratio of β between the primal and dual values of \vec{x} and \vec{y} , implying $\frac{1}{\beta}$ -competitiveness of this online fractional matching algorithm, by feasibility of \vec{y} and weak duality. The algorithm, parameterized by a function $f(\cdot)$ and parameter β to be discussed below, is given formally in Algorithm 1. In the subsequent discussion, $N_v(u)$ denotes the set of neighbors of u that arrive before v .

Algorithm 1 is parameterized by a function f and a constant β . The family of functions considered by [25] are as follows.

Definition 3.1. Let $f_\kappa(\theta) := \left(\frac{1+\kappa}{2} - \theta\right)^{\frac{1+\kappa}{2\kappa}} \left(\theta + \frac{\kappa-1}{2}\right)^{\frac{\kappa-1}{2\kappa}}$. We define $W := \{f_\kappa \mid \kappa \geq 1\}$.

As we will see, choices of β guaranteeing feasibility of \vec{x} are related to the following quantity.

Definition 3.2. For a given $f : [0, 1] \rightarrow \mathbb{R}_+$ let $\beta^*(f) := \max_{\theta \in [0, 1]} 1 + f(1 - \theta) + \int_\theta^1 \frac{1-t}{f(t)} dt$.

²Here and throughout the paper, we let $x^+ := \max\{0, x\}$ for all $x \in \mathbb{R}$.

Input : A stream of vertices v_1, v_2, \dots, v_n . At step i , vertex v_i and $N_{v_i}(v_i)$ are revealed.

Output: A fractional vertex cover solution \vec{y} and a fractional matching \vec{x} .

```

1 Let  $y_u \leftarrow 0$  for all  $u$ , let  $x_{uv} \leftarrow 0$  for all  $u, v$ .
2 foreach  $v$  in the stream do
    maximize  $\theta$ 
3     subject to:  $\theta \leq 1$ 
                 $\sum_{u \in N_v(v)} (\theta - y_u)^+ \leq f(\theta)$ 
4     foreach  $u \in N_v(v)$  do
5          $x_{uv} \leftarrow \frac{(\theta - y_u)^+}{\beta} \left(1 + \frac{1 - \theta}{f(\theta)}\right)$ .
6          $y_u \leftarrow \max\{y_u, \theta\}$ .
7      $y_v \leftarrow 1 - \theta$ .

```

Algorithm 1: Online general vertex arrival fractional matching and vertex cover

For functions $f \in W$ this definition of $\beta^*(f)$ can be simplified to $\beta^*(f) = 1 + f(0)$, due to the observation (see [25, Lemmas 4,5]) that all functions $f \in W$ satisfy

$$\beta^*(f) = 1 + f(1 - \theta) + \int_{\theta}^1 \frac{1 - \theta}{f(\theta)} d\theta \quad \forall \theta \in [0, 1]. \quad (5)$$

As mentioned above, the competitiveness of Algorithm 1 for appropriate choices of f and β is obtained by relating the overall primal and dual values, $\sum_e x_e$ and $\sum_v y_v$. As we show (and rely on later), one can even bound individual vertices' contributions to these sums. In particular, for any vertex v 's arrival time, each vertex u 's contribution to $\sum_e x_e$, which we refer to as its *fractional degree*, $x_u := \sum_{w \in N_v(u)} x_{uw}$, can be bounded in terms of its dual value by this point, y_u , as follows.

Lemma 3.3. *For any vertex $u, v \in V$, let y_u be the potential of u prior to arrival of v . Then the fractional degree just before v arrives, $x_u := \sum_{w \in N_v(u)} x_{uw}$, is bounded as follows:*

$$\frac{y_u}{\beta} \leq x_u \leq \frac{y_u + f(1 - y_u)}{\beta}.$$

Broadly, the lower bound on x_u is obtained by lower bounding the increase x_u by the increase to y_u/β after each vertex arrival, while the upper bound follows from a simplification of a bound given in [25, Invariant 1] (implying feasibility of the primal solution), which we simplify using (5). See Appendix B for a full proof.

Another observation we will need regarding the functions $f \in W$ is that they are decreasing.

Observation 3.4. *Every function $f \in W$ is non-increasing in its argument in the range $[0, 1]$.*

Proof. As observed in [25], differentiating (5) with respect to z yields $-f'(1 - z) - \frac{1 - z}{f(z)} = 0$, from which we obtain $f(z) \cdot f'(1 - z) = z - 1$. Replacing z by $1 - z$, we get $f(1 - z) \cdot f'(z) = -z$, or $f'(z) = -\frac{z}{f(1 - z)}$. As $f(z)$ is positive for all $z \in [0, 1]$, we have that $f'(z) < 0$ for all $z \in [0, 1]$. \square

The next lemma of [25] characterizes the achievable competitiveness of Algorithm 1.

Lemma 3.5 ([25]). *Algorithm 1 with function $f \in W$ and $\beta \geq \beta^*(f) = 1 + f(0)$ is $\frac{1}{\beta}$ competitive.*

Wang and Wong [25] showed that taking $\kappa \approx 1.1997$ and $\beta = \beta^*(f_\kappa)$, Algorithm 1 is ≈ 0.526 competitive. In later sections we show how to round the output of Algorithm 1 with f_κ with $\kappa = 1 + 2\varepsilon$ for some small constant ε and $\beta = 2 - \varepsilon$ to obtain a $(1/2 + \Omega(1))$ -competitive algorithm. But first, as a warm up, we show how to round this algorithm with $\kappa = 1$ and $\beta = \beta^*(f_1) = 2$.

3.2 Warmup: a $1/2$ -competitive randomized algorithm

In this section we will round the $1/2$ -competitive fractional algorithm obtained by running Algorithm 1 with function $f(\theta) = f_1(\theta) = 1 - \theta$ and $\beta = \beta^*(f) = 2$. We will devise a lossless rounding of this fractional matching algorithm, by including each edge e in the final matching with a probability equal to the fractional value x_e assigned to it by Algorithm 1. Note that if v arrives after u , then if F_u denotes the event that u is free when v arrives, then edge $\{u, v\}$ is matched by an online algorithm with probability $\Pr[\{u, v\} \in M] = \Pr[\{u, v\} \in M \mid F_u] \cdot \Pr[F_u]$. Therefore, to match each edge $\{u, v\}$ with probability x_{uv} , we need $\Pr[\{u, v\} \in M \mid F_u] = x_{uv} / \Pr[F_u]$. That is, we must match $\{u, v\}$ with probability $z_u = x_{uv} / \Pr[F_u]$ conditioned on u being free. The simplest way of doing so (if possible) is to pick an edge $\{u, v\}$ with the above probability z_u always, and to match it only if u is free. Algorithm 2 below does just this, achieving a lossless rounding of this fractional algorithm. As before, $N_v(u)$ denotes the set of neighbors of u that arrive before v .

Input : A stream of vertices v_1, v_2, \dots, v_n . At step i , vertex v_i and $N_{v_i}(v_i)$ are revealed.
Output: A matching M .

- 1 Let $y_u \leftarrow 0$ for all u , let $x_{uv} \leftarrow 0$ for all u, v .
- 2 Let $M \leftarrow \emptyset$.
- 3 **foreach** v in the stream **do**
- 4 Update y_u 's and x_{uv} 's using Algorithm 1 with $\beta = 2$ and $f = f_1$.
- 5 **foreach** $u \in N_v(v)$ **do**
- 6 $z_u \leftarrow \frac{x_{uv}}{\Pr[u \text{ is free when } v \text{ arrives}]}$. // z_u is $x_{uv}/(1 - y_u)$ as shown later
- 7 Sample (at most) one neighbor $u \in N_v(v)$ according to z_u .
- 8 **if** a free neighbor u is sampled **then**
- 9 Add $\{u, v\}$ to M .

Algorithm 2: Online vertex arrival warmup randomized fractional matching

Algorithm 2 is well defined if for each vertex v 's arrival, z is a probability distribution; i.e., $\sum_{u \in N_v(v)} z_u \leq 1$. The following lemma asserts precisely that. Moreover, it asserts that Algorithm 2 matches each edge with the desired probability.

Lemma 3.6. *Algorithm 2 is well defined, since for every vertex v on arrival, z is a valid probability distribution. Moreover, for each v and $u \in N_v(v)$, it matches edge $\{u, v\}$ with probability x_e .*

Proof. We prove both claims in tandem for each v , by induction on the number of arrivals. For the base case (v is the first arrival), the set $N_v(v)$ is empty and thus both claims are trivial. Consider the arrival of a later vertex v . By the inductive hypothesis we have that each vertex $u \in N_v(v)$ is previously matched with probability $\sum_{w \in N_v(u)} x_{uw}$. But by our choice of $f(\theta) = f_1(\theta) = 1 - \theta$ and $\beta = 2$, if w arrives after u , then y_u and θ at arrival of w satisfy $x_{uw} = \frac{(\theta - y_u)^+}{\beta} \cdot \left(1 + \frac{1 - \theta}{f(\theta)}\right) = (\theta - y_u)^+$. That is, x_{uw} is precisely the increase in y_u following arrival of w . On the other hand, when u arrived we have that its dual value y_u increased by $1 - \theta = \sum_{v' \in N_u(u)} (\theta - y_{v'})^+ = \sum_{v' \in N_u(u)} x_{uv'}$. To see this last step, we recall first that by definition of Algorithm 1 and our choice of $f(\theta) = 1 - \theta$, the value θ on arrival of v is chosen to be the largest $\theta \leq 1$ satisfying

$$\sum_{\forall u \in N_v(v)} (\theta - y_u)^+ \leq 1 - \theta. \quad (6)$$

But the inequality (6) is an equality whether or not $\theta = 1$ (if $\theta = 1$, both sides are zero). We conclude that $y_u = \sum_{v' \in N_v(u)} x_{uv'}$ just prior to arrival of v . But then, by the inductive hypothesis,

this implies that $\Pr[u \text{ free when } v \text{ arrives}] = 1 - y_u$ (yielding an easily-computable formula for z_u). Consequently, by (6) we have that when v arrives z is a probability distribution, as

$$\sum_{u \in N_v(v)} z_u = \sum_{u \in N_v(v)} \frac{(\theta - y_u)^+}{1 - y_u} \leq \sum_{u \in N_v(v): y_u \leq \theta} \frac{(\theta - y_u)^+}{1 - \theta} = \sum_{u \in N_v(v)} \frac{(\theta - y_u)^+}{1 - \theta} \leq 1.$$

Finally, for u to be matched to a latter-arriving neighbor v , it must be picked and free when v arrives, and so $\{u, v\}$ is indeed matched with probability

$$\Pr[\{u, v\} \in M] = \frac{x_{uv}}{\Pr[u \text{ is free when } v \text{ arrives}]} \cdot \Pr[u \text{ is free when } v \text{ arrives}] = x_{uv}. \quad \square$$

In the next section we present an algorithm which allows to round better-than-1/2-competitive algorithms derived from Algorithm 1.

3.3 An improved algorithm

In this section, we build on Algorithm 2 and show how to improve it to get a $(1/2 + \Omega(1))$ competitive ratio.

There are two concerns when modifying Algorithm 2 to work for a general function from the family W . The first is how to compute the probability that a vertex u is free when vertex v arrives, in Line 6. In the simpler version, we inductively showed that this probability is simply $1 - y_u$, where y_u is the dual value of u as of v 's arrival (see the proof of Lemma 3.6). With a general function f , this probability is no longer given by a simple formula. Nevertheless, it is easily fixable: We can either use Monte Carlo sampling to estimate the probability of u being free at v 's arrival to a given inverse polynomial accuracy, or we can in fact exactly compute these probabilities by maintaining their marginal values as the algorithm progresses. In what follows, we therefore assume that our algorithm can compute these probabilities exactly.

The second and more important issue is with the sampling step in Line 7. In the simpler algorithm, this step is well-defined as the sampling probabilities indeed form a valid distribution: I.e., $\sum_{u \in N_v(v)} z_u \leq 1$ for all vertices v . However, with a general function f , this sum can exceed one, rendering the sampling step in Line 7 impossible. Intuitively, we can normalize the probabilities to make it a proper distribution, but by doing so, we end up losing some amount from the approximation guarantee. We hope to recover this loss using a second sampling step, as we mentioned in Section 1.2 and elaborate below.

Suppose that, instead of $\beta = 2$ and $f = f_1$ (i.e., the function $f(\theta) = 1 - \theta$), we use $f = f_{1+2\varepsilon}$ and $\beta = 2 - \varepsilon$ to define x_{uv} and y_u values. As we show later in this section, for an ε sufficiently small, we then have $\sum_{u \in N_v(v)} z_u \leq 1 + O(\varepsilon)$, implying that the normalization factor is at most $1 + O(\varepsilon)$. However, since the approximation factor of the fractional solution is only $1/2 + O(\varepsilon)$ for such a solution, (i.e., $\sum_{\{u,v\} \in E} x_{uv} \geq (1/\beta) \cdot \sum_{u \in V} y_u$), the loss due to normalization is too significant to ignore.

Now suppose that we allow arriving vertices to sample a second edge with a small (i.e., $\sqrt{\varepsilon}$) probability and match that second edge if the endpoint of the first sampled edge is already matched. Consider the arrival of a fixed vertex v such that $\sum_{u \in N_v(v)} z_u > 1$, and let z'_u denote the normalized z_u values. Further let F_w denote the event that vertex w is free (i.e., unmatched) at the arrival of v . Then the probability that v matches u for some $u \in N_v(v)$ using either of the two sampled edges is

$$\Pr[F_u] \cdot \left(z'_u + z'_u \sqrt{\varepsilon} \cdot \sum_{w \in N_v(v)} z'_w \cdot (1 - \Pr[F_w | F_u]) \right), \quad (7)$$

which is the same expression from (1) from Section 1.2, restated here for quick reference. Recall that the first term inside the parentheses accounts for the probability that v matches u via the first sampled edges, and the second term accounts for the probability that the same happens via the second sampled edge. Note that the second sampled edge is used only when the first one is incident to an already matched vertex and the other endpoint of the second edge is free. Hence we have the summation of conditional probabilities in the second term, where the events are conditioned on the other endpoint, u , being free. If the probability given in (7) is x_{uv} for all $\{u, v\} \in E$, we would have the same guarantee as the fractional solution x_{uv} , and the rounding would be lossless. This seems unlikely, yet we can show that the quantity in (7) is at least $(1 - \varepsilon^2) \cdot x_{uv}$ for most (not by number, but by the total fractional value of x_{uv} 's) of the edges in the graph, showing that our rounding is *almost* lossless. We postpone further discussion of the analysis to Section 3.4 where we highlight the main ideas before proceeding with the formal proof.

Input : A stream of vertices v_1, v_2, \dots, v_n . At step i , vertex v_i and $N_{v_i}(v)$ are revealed.
Output: A matching M .

- 1 Let $y_u \leftarrow 0$ for all u , let $x_{uv} \leftarrow 0$ for all u, v .
- 2 Let $M \leftarrow \emptyset$.
- 3 **foreach** v in the stream **do**
- 4 Update y_u 's and x_{uv} 's using Algorithm 1 with $\beta = 2 - \varepsilon$ and $f = f_{1+2\varepsilon}$.
- 5 **foreach** $u \in N_v(v)$ **do**
- 6 /* Compute $\Pr[u \text{ is free when } v \text{ arrives}]$ as explained in Section 3.3 */
- 6 $z_u \leftarrow \frac{x_{uv}}{\Pr[u \text{ is free when } v \text{ arrives}]}$.
- 7 **foreach** $u \in N_v(v)$ **do**
- 8 $z'_u \leftarrow z_u / \max \left\{ 1, \sum_{u \in N_v(v)} z_u \right\}$.
- 9 Pick (at most) one $u_1 \in N_v(v)$ with probability z'_{u_1} .
- 10 **if** $\sum_{u \in N_v(v)} z_u > 1$ **then**
- 11 With probability $\sqrt{\varepsilon}$, pick (at most) one $u_2 \in N_v(v)$ with probability z'_{u_2} .
- 11 /* Probability of dropping edge $\{u, v\}$ can be computed using (7). */
- 12 Drop u_2 with minimal probability ensuring $\{u_2, v\}$ is matched with probability at most x_{u_2v} .
- 13 **if** a free neighbor u_1 is sampled **then**
- 14 Add $\{u_1, v\}$ to M .
- 15 **else if** a free neighbor u_2 is sampled **then**
- 16 Add $\{u_2, v\}$ to M .

Algorithm 3: A randomized online matching algorithm under general vertex arrivals.

Our improved algorithm is outlined in Algorithm 3. Up until Line 6, it is similar to Algorithm 2 except that it uses $\beta = 2 - \varepsilon$ and $f = f_{1+2\varepsilon}$ where we choose $\varepsilon > 0$ to be any constant small enough such that the results in the analysis hold. In Line 8, if the sum of z_u 's exceeds one we normalize the z_u to obtain a valid probability distribution z'_u . In Line 9, we sample the first edge incident to an arriving vertex v . In Line 11, we sample a second edge incident to the same vertex with probability $\sqrt{\varepsilon}$ if we had to scale down z_u 's in Line 8. Then in Line 12, we drop the sampled second edge with the minimal probability to ensure that no edge $\{u, v\}$ is matched with probability more than x_{uv} . Since (7) gives the exact probability of $\{u, v\}$ being matched, this probability of dropping an edge $\{u, v\}$ can be computed by the algorithm. However, to compute this, we need the

conditional probabilities $\Pr[F_w \mid F_u]$, which again can be estimated using Monte Carlo sampling³. In the subsequent lines, we match v to a chosen free neighbor (if any) among its chosen neighbors, prioritizing its first choice.

For the purpose of analysis we view Algorithm 3 as constructing a greedy matching on a directed acyclic graph (DAG) H_τ defined in the following two definitions.

Definition 3.7 (Non-adaptive selection graph G_τ). *Let τ denote the random choices made by the vertices of G . Let G_τ be the DAG defined by all the arcs $(v, u_1), (v, u_2)$ for all vertices $v \in V$. We call the arcs (v, u_1) primary arcs, and the arcs (v, u_2) the secondary arcs.*

Definition 3.8 (Pruned selection graph H_τ). *Now construct H_τ from G_τ by removing all arcs (v, u) (primary or secondary) such that there exists a primary arc (v', u) with v' arriving before v . We further remove a secondary arc (v, u) if there is a primary arc (v, u) ; i.e., if a vertex u has at least one incoming primary arc, remove all incoming primary arcs that came after the first primary arc and all secondary arcs that came after or from the same vertex as the first primary arc.*

It is easy to see that the matching constructed by Algorithm 3 is a greedy matching constructed on H_τ based on order of arrival and prioritizing primary arcs. The following lemma shows that the set of matched vertices obtained by this greedy matching does not change much for any change in the random choices of a single vertex v , which will prove useful later on. It can be proven rather directly by an inductive argument showing the size of the symmetric difference in matched vertices in G_τ and $G_{\tau'}$ does not increase after each arrival besides the arrival of v , whose arrival clearly increases this symmetric difference by at most two. See Appendix A for details.

Lemma 3.9. *Let G_τ and $G_{\tau'}$ be two realizations of the random digraph where all the vertices in the two graphs make the same choices except for one vertex v . Then the number of vertices that have different matched status (free/matched) in the matchings computed in H_τ and $H_{\tau'}$ at any point of time is at most two.*

3.4 Analysis

In this section, we analyze the competitive ratio of Algorithm 3. We start with an outline of the analysis where we highlight the main ideas.

3.4.1 High-Level Description of Analysis

As described in Section 3.3, the main difference compared to the simpler $1/2$ -competitive algorithm is the change of the construction of the fractional solution, which in turn makes the rounding more complex. In particular, we may have at the arrival of a vertex v that $\sum_{u \in N_v(v)} z_u > 1$. The majority of the analysis is therefore devoted to such “problematic” vertices since otherwise, if $\sum_{u \in N_v(v)} z_u \leq 1$, the rounding is lossless due to the same reasons as described in the simpler setting of Section 3.2. We now outline the main ideas in analyzing a vertex v with $\sum_{u \in N_v(v)} z_u > 1$. Let F_w be the event that vertex w is free (i.e., unmatched) at the arrival of v . Then, as described in Section 3.3, the probability that we select edge $\{u, v\}$ in our matching is the minimum of x_{uv} (because of the pruning in Line 12), and

$$\Pr[F_u] \cdot \left(z'_u + z'_u \sqrt{\varepsilon} \cdot \sum_{w \in N_v(v)} z'_w \cdot (1 - \Pr[F_w \mid F_u]) \right).$$

³It is also possible to compute them exactly if we allow the algorithm to take exponential time.

By definition, $\Pr[F_u] \cdot z_u = x_{uv}$, and the expression inside the parentheses is at least z_u (implying $\Pr[\{u, v\} \in M] = x_{uv}$) if

$$1 + \sqrt{\varepsilon} \cdot \sum_{w \in N_v(v)} z'_w \cdot (1 - \Pr[F_w | F_u]) \geq \frac{z_u}{z'_u}. \quad (8)$$

To analyze this inequality, we first use the structure of the selected function $f = f_{1+2\varepsilon}$ and the selection of $\beta = 2 - \varepsilon$ to show that if $\sum_{u \in N_v(v)} z_u > 1$ then several structural properties hold (see Lemma 3.10 and Corollary 3.11 in Section 3.4.2). In particular, there are absolute constants $0 < c < 1$ and $C > 1$ (both independent of ε) such that

1. $\sum_{u \in N_v(v)} z_u \leq 1 + C\varepsilon$;
2. $z_u \leq C\sqrt{\varepsilon}$ for every $u \in N_v(v)$; and
3. $c \leq \Pr[F_w] \leq 1 - c$ for every $w \in N_v(v)$.

The first property implies that the right-hand-side of (8) is at most $1 + C\varepsilon$; and the second property implies that v has at least $\Omega(1/\sqrt{\varepsilon})$ neighbors and that each neighbor u satisfies $z'_u \leq z_u \leq C\sqrt{\varepsilon}$.

For simplicity of notation, we assume further in the high-level overview that v has exactly $1/\sqrt{\varepsilon}$ neighbors and each $u \in N_v(v)$ satisfies $z'_u = \sqrt{\varepsilon}$. Inequality (8) would then be implied by

$$\sum_{w \in N_v(v)} (1 - \Pr[F_w | F_u]) \geq C. \quad (9)$$

To get an intuition why we would expect the above inequality to hold, it is instructive to consider the unconditional version:

$$\sum_{w \in N_v(v)} (1 - \Pr[F_w]) \geq c|N_v(v)| = c/\sqrt{\varepsilon} \gg C,$$

where the first inequality is from the fact that $\Pr[F_w] \leq 1 - c$ for any neighbor $w \in N_v(v)$. The large slack in the last inequality, obtained by selecting $\varepsilon > 0$ to be a sufficiently small constant, is used to bound the impact of conditioning on the event F_u . Indeed, due to the large slack, we have that (9) is satisfied if the quantity $\sum_{w \in N_v(v)} \Pr[F_w | F_u]$ is not too far away from the same summation with unconditional probabilities, i.e., $\sum_{w \in N_v(v)} \Pr[F_w]$. Specifically, it is sufficient to show

$$\sum_{w \in N_v(v)} (\Pr[F_w | F_u] - \Pr[F_w]) \leq c/\sqrt{\varepsilon} - C. \quad (10)$$

We do so by bounding the correlation between the events F_u and F_w in a highly non-trivial manner, which constitutes the heart of our analysis. The main challenges are that events F_u and F_w can be positively correlated and that, by conditioning on F_u , the primary and secondary choices of different vertices are no longer independent.

We overcome the last difficulty by replacing the conditioning on F_u by a conditioning on the component in H_τ (at the time of v 's arrival) that includes u . As explained in Section 3.3, the matching output by our algorithm is equivalent to the greedy matching constructed in H_τ and so the component containing u (at the time of v 's arrival) determines F_u . But how can this component look like, assuming the event F_u ? First, u cannot have any incoming primary arc since then u would be matched (and so the event F_u would be false). However, u could have incoming secondary arcs,




Figure 2: Two examples of the component of H_τ containing u . Vertices are depicted from right to left in the arrival order. Primary and secondary arcs are solid and dashed, respectively. The edges that take part in the matching are thick.

assuming that the tails of those arcs are matched using their primary arcs. Furthermore, u can have an outgoing primary and possibly a secondary arc if the selected neighbors are already matched. These neighbors can in turn have incoming secondary arcs, at most one incoming primary arc (due to the pruning in the definition of H_τ), and outgoing primary and secondary arcs; and so on. In Figure 2, we give two examples of the possible structure, when conditioning on F_u , of u 's component in H_τ (at the time of v 's arrival). The left example contains secondary arcs, whereas the component on the right is arguably simpler and only contains primary arcs.

An important step in our proof is to prove that, for most vertices u , the component is of the simple form depicted to the right with probability almost one. That is, it is a path P consisting of primary arcs, referred to as a primary path (see Definition 3.13) that further satisfies:

- (i) it has length $O(\ln(1/\varepsilon))$; and
- (ii) the total z -value of the arcs in the blocking set of P is $O(\ln(1/\varepsilon))$. The blocking set is defined in Definition 3.14. Informally, it contains those arcs that if appearing as primary arcs in G_τ would cause arcs of P to be pruned (or blocked) from H_τ .

Let \mathcal{P} be the primary paths of above type that appear with positive probability as u 's component in H_τ . Further let EQ_P be the event that u 's component equals P . Then we show (for most vertices) that $\sum_{P \in \mathcal{P}} \Pr[\text{EQ}_P \mid F_u]$ is almost one. For simplicity, let us assume here that the sum is equal to one. Then by the law of total probability and since $\sum_{P \in \mathcal{P}} \Pr[\text{EQ}_P \mid F_u] = 1$,

$$\begin{aligned} \sum_{w \in N_v(v)} (\Pr[F_w \mid F_u] - \Pr[F_w]) &= \sum_{P \in \mathcal{P}} \Pr[\text{EQ}_P \mid F_u] \left(\sum_{w \in N_v(v)} (\Pr[F_w \mid F_u, \text{EQ}_P] - \Pr[F_w]) \right) \\ &= \sum_{P \in \mathcal{P}} \Pr[\text{EQ}_P \mid F_u] \left(\sum_{w \in N_v(v)} (\Pr[F_w \mid \text{EQ}_P] - \Pr[F_w]) \right), \end{aligned}$$

where the last equality is because the component P determines F_u . The proof is then completed by analyzing the term inside the parentheses for each primary path $P \in \mathcal{P}$ separately. As we prove in Lemma 3.15, the independence of primary and secondary arc choices of vertices is maintained after conditioning on EQ_P .⁴ Furthermore, we show that there is a bijection between the outcomes of the unconditional and the conditional distributions, so that the expected number of vertices

⁴To be precise, conditioning on a primary path P with a so-called termination certificate T , see Definition 3.13. In the overview, we omit this detail and consider the event $\text{EQ}_{P,T}$ (instead of EQ_P) in the formal proof.

that make different choices under this pairing can be upper bounded by roughly the length of the path plus the z -value of the edges in the blocking set. So, for a path P as above, we have that the expected number of vertices that make different choices in the paired outcomes is $O(\ln(1/\varepsilon))$ which, by Lemma 3.9, implies that the expected number of vertices that change matched status is also upper bounded by $O(\ln(1/\varepsilon))$. In other words, we have for every $P \in \mathcal{P}$ that

$$\sum_{w \in N_v(v)} (\Pr[F_w | \text{EQ}_P] - \Pr[F_w]) \leq \sum_{w \in V} (\Pr[F_w | \text{EQ}_P] - \Pr[F_w]) = O(\ln(1/\varepsilon)),$$

which implies (10) for a small enough choice of ε . This completes the overview of the main steps in the analysis. The main difference in the formal proof is that not all vertices satisfy that their component is a short primary path with probability close to 1. To that end, we define the notion of *good* vertices in Section 3.4.4, which are the vertices that are very unlikely to have long directed paths of primary arcs rooted at them. These are exactly the vertices v for which we can perform the above analysis for most neighbors u (in the proof of the “key lemma”) implying that the rounding is almost lossless for v . Then, in Section 3.4.6, we show using a rather simple charging scheme that most of the vertices in the graph are good. Finally, in Section 3.4.7, we put everything together and prove Theorem 1.2.

3.4.2 Useful Properties of W Functions and Algorithm 3

For the choice of $f = f_{1+2\varepsilon}$ as we choose, we have $f(\theta) = (1 + \varepsilon - \theta) \cdot \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta}\right)^{\frac{\varepsilon}{1+2\varepsilon}}$. In Appendix C we give a more manageable upper bound for $f(\theta)$ which holds for sufficiently small ε . Based on this simple upper bound on f and some basic calculus, we obtain the following useful structural properties for the conditional probabilities, z_u , of Algorithm 3. See Appendix C.

Lemma 3.10. *(Basic bounds on conditional probabilities z_u) There exist absolute constants $c \in (0, 1)$ and $C > 1/c > 1$ and $\varepsilon_0 \in (0, 1)$ such that for every $\varepsilon \in (0, \varepsilon_0)$ the following holds: for every vertex $v \in V$, if y_u is the dual variable of a neighbor $u \in N_v(v)$ before v 's arrival and θ is the value chosen by Algorithm 1 on v 's arrival, then for z_u as defined in Algorithm 3, we have:*

- (1) If $\theta \notin (c, 1 - c)$, then $\sum_{u \in N_v(v)} z_u \leq 1$,
- (2) If $\theta \in [0, 1]$, then $\sum_{u \in N_v(v)} z_u \leq 1 + C\varepsilon$,
- (3) If $\sum_{u \in N_v(v)} z_u > 1$, then $z_u \leq C\sqrt{\varepsilon}$ for every $u \in N_v(v)$,
- (4) If $\sum_{u \in N_v(v)} z_u > 1$, then for every $u \in N_v(v)$ such that $z_u > 0$, one has $y_u \in [c/2, 1 - c/2]$, and
- (5) For all $u \in N_v(v)$, one has $z_u \leq 1/2 + O(\sqrt{\varepsilon})$.

The following corollary will be critical to our analysis:

Corollary 3.11. *There exist absolute constants $c > 0$ and $\varepsilon_0 > 0$ such that for all $\varepsilon \in (0, \varepsilon_0)$, on arrival of any vertex $v \in V$, if z as defined in Algorithm 3 satisfies $\sum_{u \in N_v(v)} z_u > 1$, then for every $u \in N_v(v)$ we have*

$$c \leq \Pr[u \text{ is free when } v \text{ arrives}] \leq 1 - c.$$

Proof. By Lemma 3.10, (1) and (4) we have that if $\sum_{u \in N_v(v)} z_u > 1$, then $\theta \in (c, 1 - c)$ (c is the constant from Lemma 3.10), and for every $u \in N_v(v)$ one has

$$y_u \in [c/2, 1 - c/2]. \quad (11)$$

On the other hand, by Lemma 3.3 one has

$$\frac{y_u}{\beta} \leq x_u \leq \frac{y_u + f(1 - y_u)}{\beta}, \quad (12)$$

where x_u is the fractional degree of u when v arrives.

We now note that by Lemma 3.10, (2), we have that Algorithm 3 matches every vertex u with probability at least $x_u/(1 + C\varepsilon)$ (due to choices of primary arcs), and thus

$$\begin{aligned} \Pr[u \text{ is free when } v \text{ arrives}] &\leq 1 - \frac{x_u}{1 + C\varepsilon} \\ &\leq 1 - \frac{y_u}{\beta(1 + C\varepsilon)} \quad (\text{by (12)}) \\ &\leq 1 - \frac{c/2}{2(1 + C\varepsilon)} \quad (\text{by (11) and the setting } \beta = 2 - \varepsilon \leq 2) \\ &\leq 1 - c/5, \end{aligned}$$

as long as ε is sufficiently small.

For the other bound we will use two facts. The first is that since $f(y)$ is monotone decreasing by Observation 3.4 and since we picked $\beta > \beta^*(f) = 1 + f(0)$, we have that for any $y \leq 1 - c/2 \leq 1$,

$$y + f(1 - y) \leq 1 - c/2 + f(0) < \beta - c/2. \quad (13)$$

Then, using the fact that by Line 12, Algorithm 3 matches every vertex u with probability at most x_u , we obtain the second bound, as follows.

$$\begin{aligned} \Pr[u \text{ is free when } v \text{ arrives}] &\geq 1 - x_u \\ &\geq 1 - \frac{y_u + f(1 - y_u)}{\beta} \quad (\text{by (12)}) \\ &\geq 1 - \frac{\beta - c/2}{\beta} \quad (\text{by (11) and (13)}) \\ &\geq c/5. \quad (\beta = 2 - \varepsilon < 2.5) \end{aligned}$$

Choosing $c/5$ as the constant in the statement of the lemma, we obtain the result. \square

Finally, for our analysis we will rely on the competitive ratio of the fractional solution maintained in Line 4 being $1/\beta$. This follows by Lemma 3.5 and the fact that for our choices of $\beta = 2 - \varepsilon$ and $f = f_{1+2\varepsilon}$ we have that $\beta \geq \beta^*(f)$. See Appendix C for a proof of this fact.

Fact 3.12. *For all sufficiently small $\varepsilon > 0$, we have that $2 - \varepsilon \geq \beta^*(f_{1+2\varepsilon})$.*

3.4.3 Structural Properties of G_τ and H_τ

In our analysis later, we focus on maximal primary paths (directed paths made of primary arcs) in H_τ , in the sense that the last vertex along the primary path has no outgoing primary arc in H_τ . The following definition captures termination certificates of such primary paths.

Definition 3.13 (Certified Primary Path). *A tuple (P, T) is a certified primary path in H_τ if P is a directed path of primary arcs in H_τ and either*

- (a) *the last vertex of P does not have an outgoing primary arc in G_τ and $T = \emptyset$, or*
- (b) *the last vertex u of P has an outgoing primary arc (u, w) in G_τ and $T = (u', w)$ is a primary arc in H_τ such that u' precedes u in the arrival order.*

To elaborate, a certified primary path (P, T) is made of a (directed) path P of primary arcs in H_τ and T is a certificate of P 's termination in H_τ that ensures the last vertex u in P has no outgoing primary arc in H_τ , either due to u not picking a primary arc with $T = \emptyset$, or due to the picked primary arc (u, w) being blocked by another primary arc $T = (u', w)$ which appears in H_τ .

As described, G_τ and H_τ differ in arcs (u, w) that are *blocked* by previous primary arcs to their target vertex w . We generally define sets of arcs which can block an edge, or a path, or a certified path from appearing in H_τ as in the following definition:

Definition 3.14 (Blocking sets). *For an arc (u, w) , define its blocking set*

$$B(u, w) := \{(u', w) \mid \{u', w\} \text{ is an edge and } u' \text{ arrived before } u\}$$

to be those arcs, the appearance of any of which as primary arc in G_τ blocks (u, v) from being in H_τ . In other words, an arc (u, v) is in H_τ as primary or secondary arc if and only if (u, v) is in G_τ and none of the arcs in its blocking set $B(u, v)$ is in G_τ as a primary arc.

The blocking set of a path P is simply the union of its arcs' blocking sets,

$$B(P) := \bigcup_{(u,v) \in P} B(u, v).$$

The blocking set of a certified primary path (P, T) is the union of blocking sets of P and T ,

$$B(P, T) := B(P \cup T).$$

The probability of an edge, or path, or certified primary path appearing in H_τ is governed in part by the probability of arcs in their blocking sets appearing as primary arcs in G_τ . As an arc (v, u) is picked as primary arc by when v arrives with probability roughly z_u (more precisely, $z'_u \in [z_{vu}/(1 + C\varepsilon), z_{vu}]$, by Lemma 3.10), it will be convenient to denote by $z(v, u)$ and $z'(v, u)$ the values z_u and z'_u when v arrives, and by $z(S) = \sum_{s \in S} z(s)$ and $z'(S) = \sum_{s \in S} z'(s)$ the sum of z - and z' -values of arcs in a set of arcs S .

Product distributions. Note that by definition the distribution over primary and secondary arc choices of vertices are product distributions (they are independent). As such, their joint distribution is defined by their marginals. Let p_w and s_w denote the distribution on primary and secondary arc choices of w , respectively. That is, for every $u \in N_w(w)$, $p_w(u)$ is the marginal probability that w selects (w, u) as its primary arc, and $s_w(u)$ is the marginal probability that w selects (w, u) as its secondary arc. Given our target bound (8), it would be useful to show that conditioning on F_u preserves the independence of these arc choices. Unfortunately, conditioning on F_u does not preserve this independence. We will therefore refine our conditioning later on the existence of primary paths in H_τ , which as we show below maintains independence of the arc choices.

Lemma 3.15. *For a certified primary path (P, T) let $\text{EQ}_{(P, T)}$ be the event that the path P equals a maximal connected component in H_τ and the termination of P is certified by T . Then the*

conditional distributions of primary and secondary choices conditioned on $\text{EQ}_{(P,T)}$ are product distributions; i.e., these conditional choices are independent. Moreover, if we let \tilde{p}_w and \tilde{s}_w denote the conditional distribution on primary and secondary choices of w , respectively, then

$$\text{TV}(p_w, \tilde{p}_w) \leq z(R(w)) \quad \text{and} \quad \text{TV}(s_w, \tilde{s}_w) \leq z(R(w)),$$

where $R(w) \subseteq \{w\} \times N_w(w)$ is the set of arcs leaving w whose existence as primary arcs in G_τ is ruled out by conditioning on $\text{EQ}_{(P,T)}$, and the union of these $R(w)$, denoted by $R(P,T)$, satisfies

$$R(P,T) := \bigcup_w R(w) \subseteq B(P,T) \cup \{(w,r) \mid r \text{ is root of } P\} \cup \bigcup_{w \in P \cup \{w: T=(w,w')\}} \{w\} \times N_w(w). \quad (14)$$

Proof. We first bound the total variation distance between the conditional and unconditional distributions. For primary choices, conditioning on $\text{EQ}_{(P,T)}$ rules out the following sets of primary arc choices. For vertex $w \notin P$ arriving before the root r of P this conditioning rules out w picking any edge in $B(P,T)$ as primary arc. For vertices $w \notin P$ with w arriving after the root r of P this conditioning rules out picking arcs (w,r) . Finally, this conditioning rules out some subset of arcs leaving vertices in $P \cup \{w : T = (w, w')\}$. Taking the union over these supersets of $R(w)$, we obtain (14). Now, the probability of each ruled out primary choice $(w, u) \in R(w)$ is zero under \tilde{p}_w and $z'(w, u)$ under p_w , and all other primary choices have their probability increase, with a total increase of $\sum_{(w,u) \in R(w)} z'(w, u)$, from which we conclude that

$$\text{TV}(p_w, \tilde{p}_w) = \frac{1}{2} \sum_{u \in N_w(w)} |p_w(u) - \tilde{p}_w(u)| = z'(R(w)) \leq z(R(w)).$$

The proof for secondary arcs is nearly identical, the only differences being that the sets of ruled out secondary arcs can be smaller (specifically, secondary arcs to w' such that $T = (u, w')$ are not ruled out by this conditioning), and the probability of any arc (w, u) being picked as secondary arc of w is at most $\sqrt{\varepsilon} \cdot z'(w, u) \leq z(w, u)$.

Finally, we note that primary and secondary choices for different vertices are independent. Therefore, conditioning on each vertex w not picking a primary arc in its ruled out set $R(w)$ still yields a product distribution, and similarly for the distributions over secondary choices. \square

It is easy to show that a particular certified primary path (P, T) with high value of $z(B(P, T))$ is unlikely to appear in H_τ , due to the high likelihood of arcs in its breaking set being picked as primary arcs. The following lemma asserts that the probability of a vertex u being the root of *any* primary certified path (P, T) with high $z(B(P, T))$ value is low.

Lemma 3.16. *For any $k \geq 0$ and any vertex u , we have the following*

$$\Pr[H_\tau \text{ contains any certified primary path } (P, T) \text{ with } P \text{ rooted at } u \text{ and } z(B(P, T)) \geq k] \leq e^{-k/2},$$

$$\Pr[H_\tau \text{ contains any primary path } P \text{ rooted at } u \text{ with } z(B(P)) \geq k] \leq e^{-k/2}.$$

Proof. We first prove the bound for certified primary paths. For a certified primary path (P, T) where the last vertex of P is w , define P^* as follows:

$$P^* = \begin{cases} P & \text{if } T = \emptyset \\ P \cup \{(w, w'')\} & \text{if } T = (w', w''). \end{cases}$$

Observe that $z(B(P^*)) \geq k$ whenever $z(B(P, T)) \geq k$. This is trivial when $T = \emptyset$. To see this for the case $T = (w', w'')$, let w be the last vertex of P , and note that $B(w', w'') \subseteq B(w, w'')$, as w arrives after w' . Also note that for (P, T) to be in H_τ , we have that P^* must be in G_τ .

We say a directed primary path $P' = u \rightarrow u_1 \rightarrow \dots \rightarrow u_{\ell-1} \rightarrow u_\ell$ is k -minimal if $z(B(P')) \geq k$ and $z(B(P' \setminus \{(u_{\ell-1}, u_\ell)\})) < k$. For such a path P' , define $B^*(P')$ as follows: Initially set $B^*(P') = B(P' \setminus \{(u_{\ell-1}, u_\ell)\})$. Then from $B(u_{\ell-1}, u_\ell)$, the breaking set of the last arc of P' , add arcs to $B^*(P')$ in reverse order of their sources' arrival until $z(B^*(P')) \geq k$.

Consider a certified primary path (P, T) with P rooted at u . If a k -minimal path rooted at u which is not a prefix of P^* is contained in G_τ , then (P, T) does not appear in G_τ , and therefore it does not appear in H_τ . On the other hand, if $z(B(P, T)) \geq k$ then for (P, T) to appear in H_τ , we must have that the (unique) k -minimal prefix P' of P^* must appear in G_τ , and that none of the edges of $B^*(P')$ appear in G_τ . Moreover, for any certified primary path with $z(B(P, T))$, conditioning on the existence of P' in G_τ does not affect random choices of vertices with outgoing arcs in $B^*(P')$, as these vertices are not in P' . Since by Lemma 3.10 each arc (w, w') appears in G_τ with probability $z'(v, u) \geq z(v, u)/(1 + C\varepsilon) \geq z(v, u)/2$, we conclude that for any k -minimal primary path P' rooted at u , we have

$$\begin{aligned} & \Pr[H_\tau \text{ contains any certified primary path } (P, T) \text{ with } z(B(P, T)) \geq k \mid P' \text{ is in } G_\tau] \\ & \leq \Pr[\text{No edge in } B^*(P') \text{ is in } G_\tau \mid P' \text{ is in } G_\tau] \\ & = \prod_{w \notin P'} (1 - \Pr[\text{Some primary edge in } B^*(P') \cap (\{w\} \times N_w(w)) \text{ is in } G_\tau]) \\ & \leq \prod_{w \notin P'} \exp\left(-\sum_{(w, w') \in B(P, T) \times N_w(w)} z(w, w')/2\right) \\ & \leq \exp(-z(B^*(P'))/2) \leq e^{-k/2}. \end{aligned}$$

Taking total probability \mathcal{P}_u , the set of all k -minimal primary paths P' rooted at u , we get that indeed, since u is the root of at most one k -minimal primary path in any realization of G_τ ,

$$\begin{aligned} & \Pr[H_\tau \text{ contains a certified primary path } (P, T) \text{ rooted at } u \text{ with } z(B(P, T)) \geq k] \\ & \leq \sum_{P' \in \mathcal{P}_u} \underbrace{\Pr[H_\tau \text{ contains a } (P, T) \text{ with } z(B(P, T)) \geq k \mid P' \text{ is in } G_\tau]}_{\leq e^{-k/2}} \cdot \Pr[P' \text{ is in } G_\tau] \leq e^{-k/2}. \end{aligned}$$

The proof for primary path is essentially the same as the above, taking $P^* = P$. \square

3.4.4 Analyzing Good Vertices

Consider the set of vertices that are unlikely to be roots of long directed paths of primary arcs in H_τ . In this section, we show that Algorithm 3 achieves almost lossless rounding for such vertices, and hence we call them *good* vertices. We start with a formal definition:

Definition 3.17 (Good vertices). *We say that a vertex v is good if*

$$\Pr_\tau[H_\tau \text{ has a primary path rooted at } v \text{ of length at least } 2000 \cdot \ln(1/\varepsilon)] \leq \varepsilon^6.$$

Otherwise, we say v is bad.

As the main result of this section, for good vertices, we prove the following:

Theorem 3.18. *Let v be a good vertex. Then*

$$\Pr[v \text{ is matched on arrival}] \geq (1 - \varepsilon^2) \cdot \sum_{u \in N_v(v)} x_{uv}.$$

Notational conventions. Throughout this section, we fix v and let z, z' be as in Algorithm 3. Moreover, for simplicity of notation, we suppose that the stream of vertices ends just before v 's arrival and so quantities, such as G_τ and H_τ , refer to their values when v arrives. For a vertex u , we let F_u denote the event that u is free (i.e., unmatched) when v arrives. In other words, F_u is the event that u is free in the stream that ends just before v 's arrival.

To prove the theorem, first note that it is immediate if $\sum_{u \in N_v(v)} z_u \leq 1$: in that case, we have $z' = z$ and so the probability to match v by a primary edge, by definition of z_u , is simply

$$\sum_{u \in N_v(v)} z_u \cdot \Pr[F_u] = \sum_{u \in N_v(v)} x_{uv}.$$

From now on we therefore assume $\sum_{u \in N_v(v)} z_u > 1$, which implies

$$(I) \sum_{u \in N_v(v)} z'_u = 1,$$

and moreover, by Lemma 3.10 and Corollary 3.11, for every $u \in N_v(v)$:

$$(II) z_u \leq C\sqrt{\varepsilon},$$

$$(III) z_u \leq (1 + C\varepsilon) \cdot z'_u, \text{ and}$$

$$(IV) c \leq \Pr[F_u] \leq 1 - c,$$

where c is the constant of Corollary 3.11 and C is the constant of Lemma 3.10.

We now state the key technical lemma in the proof of Theorem 3.18:

Lemma 3.19. *Consider a neighbor $u \in N_v(v)$ such that*

$$\Pr_\tau[H_\tau \text{ has a primary path rooted at } u \text{ of length at least } 2000 \cdot \ln(1/\varepsilon) \mid F_u] \leq \varepsilon^2. \quad (15)$$

Then,

$$\sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w \mid F_u] - \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w] \leq \varepsilon^{1/3}. \quad (16)$$

Note that the above lemma bounds the quantity $\sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w \mid F_u]$, which will allow us to show that (8) holds and thus the edge $\{u, v\}$ is picked in the matching with probability very close to x_{uv} . Before giving the proof of the lemma, we give the formal argument why the lemma implies the theorem.

Proof of Theorem 3.18. Define S to be the neighbors u in $N_v(v)$ satisfying

$$\Pr_\tau[H_\tau \text{ has a primary path rooted at } u \text{ of length at least } 2000 \cdot \ln(1/\varepsilon) \mid F_u] > \varepsilon^2.$$

In other words, S is the set of neighbors of v that violate (15). As v is good, we have

$$\begin{aligned} \varepsilon^6 &\geq \Pr_\tau[H_\tau \text{ has a primary path rooted at } v \text{ of length at least } 2000 \cdot \ln(1/\varepsilon)] \\ &\geq \sum_{u \in N_v(v)} z'_u \cdot \Pr[F_u] \cdot \Pr_\tau[H_\tau \text{ has a primary path rooted at } u \text{ of length at least } 2000 \cdot \ln(1/\varepsilon) - 1 \mid F_u] \\ &\geq \sum_{u \in N_v(v)} z'_u \cdot \Pr[F_u] \cdot \Pr_\tau[H_\tau \text{ has a primary path rooted at } u \text{ of length at least } 2000 \cdot \ln(1/\varepsilon) \mid F_u] \\ &\geq \sum_{u \in S} z'_u \cdot \Pr[F_u] \cdot \varepsilon^2. \end{aligned}$$

The second inequality holds because v selects the primary arc (u, v) with probability z'_u and, conditioned on F_u , u cannot already have an incoming primary arc, which implies that (u, v) is present in H_τ . The last inequality follows from the choice of S .

By Property (III), $z_u \leq (1 + C\varepsilon) \cdot z'_u$ and so by rewriting we get

$$\sum_{u \in S} x_{uv} = \sum_{u \in S} z_u \cdot \Pr[F_u] \leq (1 + C\varepsilon) \cdot \sum_{u \in S} z'_u \cdot \Pr[F_u] \leq (1 + C\varepsilon) \cdot \varepsilon^4 \leq \varepsilon^3.$$

In other words, the contribution of the neighbors of v in S to $\sum_{u \in N_v(v)} x_{uv}$ is insignificant compared to the contribution of all neighbors,

$$\sum_{u \in N_v(v)} x_{uv} = \sum_{u \in N_v(v)} z_u \cdot \Pr[F_u] \geq c, \quad (17)$$

where the inequality follows by the assumption $\sum_{u \in N_v(v)} z_u \geq 1$ and $\Pr[F_u] \geq c$ by Property (IV).

We proceed to analyze a neighbor $u \in N_v(v) \setminus S$. Recall that it is enough to verify (8) to conclude that edge $\{u, v\}$ is picked in the matching with probability x_{uv} . We have that

$$\begin{aligned} & 1 + \sqrt{\varepsilon} \sum_{w \in N_v(v)} z'_w \cdot (1 - \Pr[F_w | F_u]) \\ \geq & 1 + \sqrt{\varepsilon} \sum_{w \in N_v(v)} z'_w \cdot (1 - \Pr[F_w]) - \sqrt{\varepsilon} \cdot \varepsilon^{1/3} && \text{(by Lemma 3.19)} \\ \geq & 1 + \sqrt{\varepsilon} \sum_{w \in N_v(v)} z'_w \cdot c - \sqrt{\varepsilon} \cdot \varepsilon^{1/3} && (\Pr[F_w] \leq 1 - c \text{ by (IV)}) \\ = & 1 + \sqrt{\varepsilon}c - \sqrt{\varepsilon} \cdot \varepsilon^{1/3} && \left(\sum_{w \in N_v(v)} z'_w = 1 \text{ by (I)} \right) \\ \geq & 1 + C\varepsilon && \text{(for } \varepsilon \text{ small enough)} \\ \geq & z_u/z'_u. && \text{(by (III))} \end{aligned}$$

Therefore, by definition of S and Lemma 3.19, we thus have that for every $u \in N_v(v) \setminus S$, the edge $\{u, v\}$ is taken in the matching with probability x_{uv} . Thus, the probability that v is matched on arrival is, as claimed, at least

$$\sum_{u \in N_v(v) \setminus S} x_{uv} = \sum_{u \in N_v(v)} x_{uv} - \sum_{u \in S} x_{uv} \geq \sum_{u \in N_v(v)} x_{uv} - \varepsilon^3 \geq (1 - \varepsilon^2) \sum_{u \in N_v(v)} x_{uv},$$

where the last inequality holds because we have $\sum_{u \in N_v(v)} x_{uv} \geq c$, as calculated in (17). \square

3.4.5 Proof of the Key Lemma

It remains to prove the key lemma, Lemma 3.19, which we do here.

Proof of Lemma 3.19. For a certified primary path (P, T) let $\text{EQ}_{(P, T)}$ be the event as defined in Lemma 3.15, and let $\text{IN}_{(P, T)}$ be the event that P is a maximal *primary* path in H_τ and the termination of P is certified by T . Further, let

$$\mathcal{C} = \{(P, T) : (P, T) \text{ is a certified primary path rooted at } u \text{ with } \Pr[\text{IN}_{(P, T)}] > 0\}$$

be the set of certified primary paths rooted at u that have a nonzero probability of being maximal in H_τ . Then, by the law of total probability and since $\sum_{(P,T) \in \mathcal{C}} \Pr[\text{IN}_{(P,T)} \mid F_u] = 1$ (since conditioning on F_u implies in particular that u has no incoming primary arc), we can rewrite the expression to bound, $\sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w \mid F_u] - \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w]$, as

$$\sum_{(P,T) \in \mathcal{C}} \Pr[\text{IN}_{(P,T)} \mid F_u] \left(\sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w \mid F_u, \text{IN}_{(P,T)}] - \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w] \right). \quad (18)$$

We analyze this expression in two steps. First, in the next claim, we show that we can focus on the case when the certified path (P, T) is very structured and equals the component of u in H_τ . We then analyze the sum in that structured case.

Claim 3.20. *Let $\mathcal{P} \subseteq \mathcal{C}$ contain those certified primary paths (P, T) of \mathcal{C} that satisfy: P has length less than $2000 \cdot \ln(1/\varepsilon)$ and $z(B(P, T)) \leq 2 \ln(1/\varepsilon)$. Then, we have*

$$(18) \leq \sum_{(P,T) \in \mathcal{P}} \Pr[\text{EQ}_{(P,T)} \mid F_u] \left(\sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w \mid \text{EQ}_{(P,T)}] - \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w] \right) + \varepsilon^{1/3}/2.$$

Proof. Define the following subsets of certified primary paths rooted at u :

$$\begin{aligned} \mathcal{C}_1 &= \{(P, T) \in \mathcal{C} \mid P \text{ is of length at least } 2000 \cdot \ln(1/\varepsilon)\} \\ \mathcal{C}_2 &= \{(P, T) \in \mathcal{C} \setminus \mathcal{C}_1 \mid z(B(P, T)) > 2 \ln(1/\varepsilon)\} \end{aligned}$$

Note that $\mathcal{P} = \mathcal{C} \setminus (\mathcal{C}_1 \cup \mathcal{C}_2)$. Since u satisfies (15), we have that

$$\sum_{(P,T) \in \mathcal{C}_1} \Pr[\text{IN}_{(P,T)} \mid F_u] \leq \varepsilon^2 \leq \varepsilon^{1/3}/6.$$

On the other hand, by Lemma 3.16 and $\Pr[F_u] \geq c$ (by Property (IV)), we have that

$$\sum_{(P,T) \in \mathcal{C}_2} \Pr[\text{IN}_{(P,T)} \mid F_u] \leq c^{-1} \cdot \sum_{(P,T) \in \mathcal{C}_2} \Pr[\text{IN}_{(P,T)}] \leq c^{-1} \cdot \varepsilon \leq \varepsilon^{1/3}/6.$$

In other words, almost all probability mass lies in those outcomes where one of the certified paths $(P, T) \in \mathcal{P}$ is in H_τ . It remains to prove that, in those cases, we almost always have that the component of u in H_τ equals the path P (whose termination is certified by T). Specifically, let $\overline{\text{EQ}}_{(P,T)}$ denote the complement of $\text{EQ}_{(P,T)}$. We show

$$\Pr \left[\overline{\text{EQ}}_{(P,T)} \mid \text{IN}_{(P,T)} \right] \leq \varepsilon^{1/3}/7. \quad (19)$$

To see this, note that by the definition of the event $\text{IN}_{(P,T)}$, if we restrict ourselves to primary edges then the component of u in H_τ equals P . We thus have that for the event $\overline{\text{EQ}}_{(P,T)}$ to be true at least one of the vertices in P must have an incoming or outgoing *secondary* edge. Hence the expression $\Pr \left[\overline{\text{EQ}}_{(P,T)} \mid \text{IN}_{(P,T)} \right]$ can be upper bounded by

$$\Pr[\text{a vertex in } P \text{ has an incoming or outgoing secondary arc in } G_\tau \mid \text{IN}_{(P,T)}] \quad (20)$$

Note that event $\text{IN}_{(P,T)}$ is determined solely by choices of primary arcs. By independence of these choices and choices of secondary arcs, conditioning on $\text{IN}_{(P,T)}$ does not affect the distribution of secondary arcs. So the probability that any of the nodes in P selects a secondary edge is at most $\sqrt{\varepsilon}$. Thus, by union bound, the probability that any of the $|P| \leq 2000 \cdot \ln(1/\varepsilon)$ vertices in P pick a secondary arc is at most $\sqrt{\varepsilon} \cdot 2000 \cdot \ln(1/\varepsilon)$. We now turn our attention to incoming secondary arcs. First, considering the secondary arcs that go into u , we have

$$c \leq \Pr[F_u] \leq \prod_{(w,u) \in B(v,u)} (1 - z(w,u)/2) \leq \exp(-z(B(v,u))/2),$$

because any arc $(w,u) \in B(v,u)$ appears as a primary arc in G_τ independently with probability at least $z(w,u)/2$ and the appearance of such an arc implies that u has an incoming primary arc in H_τ and is therefore matched; i.e., the event F_u is false in this case. We thus have $z(B(v,u)) \leq 2 \ln(1/c)$. Further, since $(P,T) \notin \mathcal{C}_2$, we have $z(B(P)) \leq z(B(P,T)) \leq 2 \ln(1/\varepsilon)$. Again using that the conditioning on $\text{IN}_{(P,T)}$ does not affect the distribution of secondary edges, we have that the probability of an incoming secondary arc to any vertex in P is at most $\sqrt{\varepsilon} \cdot (2 \ln(1/c) + 2 \ln(1/\varepsilon))$. Thus, by union bound, the probability that any vertex in P has an incoming or outgoing secondary arc conditioned on $\text{IN}_{(P,T)}$ is at most

$$\sqrt{\varepsilon} \cdot 2000 \cdot \ln(1/\varepsilon) + \sqrt{\varepsilon} \cdot (2 \ln(1/c) + 2 \ln(1/\varepsilon)) \leq \varepsilon^{1/3}/7,$$

for sufficiently small ε , which implies (19) via (20).

We now show how the above concludes the proof of the claim. We have shown that each one of the two sets $\mathcal{C}_1, \mathcal{C}_2$ contributes at most $\varepsilon^{1/3}/6$ to (18) (where we use that $\sum_{w \in N_v(v)} z'_w = 1$). Hence,

$$(18) \leq \sum_{(P,T) \in \mathcal{P}} \Pr[\text{IN}_{(P,T)} | F_u] \left(\sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w | \text{IN}_{(P,T)}, F_u] - \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w] \right) + 2\varepsilon^{1/3}/6.$$

This intuitively concludes the proof of the claim as (19) says that $\Pr[\text{EQ}_{(P,T)} | \text{IN}_{(P,T)}]$ is almost 1. The formal calculations are as follows. Since the event $\text{EQ}_{(P,T)}$ implies the event $\text{IN}_{(P,T)}$, we have that

$$\Pr[\text{EQ}_{(P,T)}] = \Pr[\text{EQ}_{(P,T)} \wedge \text{IN}_{(P,T)}] = \Pr[\text{IN}_{(P,T)}] - \Pr[\overline{\text{EQ}_{(P,T)}} \wedge \text{IN}_{(P,T)}],$$

which by (19) implies

$$\Pr[\text{EQ}_{(P,T)}] = \Pr[\text{IN}_{(P,T)}] \left(1 - \Pr[\overline{\text{EQ}_{(P,T)}} | \text{IN}_{(P,T)}] \right) \geq \Pr[\text{IN}_{(P,T)}] \left(1 - \varepsilon^{1/3}/7 \right). \quad (21)$$

We use this to rewrite $\Pr[\text{IN}_{(P,T)} | F_u] \left(\sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w | \text{IN}_{(P,T)}, F_u] - \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w] \right)$. Specifically, by law of total probability, it can be rewritten as the sum of the expressions (22) and (23) below:

$$\begin{aligned} & \Pr[\text{EQ}_{(P,T)} \wedge \text{IN}_{(P,T)} | F_u] \left(\sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w | \text{EQ}_{(P,T)}, \text{IN}_{(P,T)}, F_u] - \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w] \right) \\ &= \Pr[\text{EQ}_{(P,T)} | F_u] \left(\sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w | \text{EQ}_{(P,T)}] - \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w] \right) \end{aligned} \quad (22)$$

and

$$\Pr[\overline{\text{EQ}}_{(P,T)} \wedge \text{IN}_{(P,T)} \mid F_u] \left(\sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w \mid \overline{\text{EQ}}_{(P,T)}, \text{IN}_{(P,T)}, F_u] - \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w] \right), \quad (23)$$

where (23) can be upper bounded as follows:

$$\begin{aligned} (23) &\leq \Pr[\overline{\text{EQ}}_{(P,T)} \wedge \text{IN}_{(P,T)} \mid F_u] && \text{(by } \sum_{w \in N_v(v)} z'_w \leq 1) \\ &\leq c^{-1} \cdot \Pr[\overline{\text{EQ}}_{(P,T)} \wedge \text{IN}_{(P,T)}] && \text{(by } c \leq \Pr[F_u]) \\ &= c^{-1} \cdot \Pr[\text{IN}_{(P,T)}] \cdot \Pr[\overline{\text{EQ}}_{(P,T)} \mid \text{IN}_{(P,T)}] \\ &\leq c^{-1} \cdot \frac{\Pr[\text{EQ}_{(P,T)}]}{1 - \varepsilon^{1/3}/7} \cdot (\varepsilon^{1/3}/7) && \text{(by (19) and (21))} \\ &\leq \Pr[\text{EQ}_{(P,T)}] \cdot \varepsilon^{1/3}/6. && \text{(for } \varepsilon \text{ small enough)} \end{aligned}$$

As at most one of the events $\{\text{EQ}_{(P,T)}\}_{(P,T) \in \mathcal{P}}$ is true in any realization of G_τ , we have that $\sum_{(P,T) \in \mathcal{P}} \Pr[\overline{\text{EQ}}_{(P,T)} \wedge \text{IN}_{(P,T)} \mid F_u] \leq \sum_{(P,T) \in \mathcal{P}} (\Pr[\text{EQ}_{(P,T)}] \cdot \varepsilon^{1/3}/6) \leq \varepsilon^{1/3}/6$. Thus, again using that $\sum_{w \in N_v(v)} z_w \leq 1$, we have that

$$\begin{aligned} (18) &\leq \sum_{(P,T) \in \mathcal{P}} \Pr[\text{IN}_{(P,T)} \mid F_u] \left(\sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w \mid \text{IN}_{(P,T)}, F_u] - \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w] \right) + 2\varepsilon^{1/3}/6 \\ &\leq \sum_{(P,T) \in \mathcal{P}} \Pr[\text{EQ}_{(P,T)} \mid F_u] \left(\sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w \mid \text{EQ}_{(P,T)}] - \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w] \right) + 3\varepsilon^{1/3}/6, \end{aligned}$$

as claimed. \square

The previous claim bounded the contribution of certified primary paths in $\mathcal{C} \setminus \mathcal{P}$ to (18). The following claim bounds the contribution of paths in \mathcal{P} .

Claim 3.21. *Let $\mathcal{P} \subseteq \mathcal{C}$ contain those certified primary paths (P, T) of \mathcal{C} that satisfy: P has length less than $2000 \cdot \ln(1/\varepsilon)$ and $z(B(P, T)) \leq 2 \ln(1/\varepsilon)$. Then, we have*

$$\sum_{(P,T) \in \mathcal{P}} \Pr[\text{EQ}_{(P,T)}] \left(\sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w \mid \text{EQ}_{(P,T)}] - \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w] \right) \leq \varepsilon^{1/3}/2.$$

Proof. We prove the claim in two steps: first we construct a chain of distributions that interpolates between the unconditional distribution of H_τ and its conditional distribution, and then bound the expected number of vertices that change their matched status along that chain. For the remainder of the proof we fix the certified primary path (P, T) .

Constructing a chain of distributions. Let $H_\tau^{(0)}$ denote the unconditional distribution of H_τ when v arrives, and let $H_\tau^{(n)}$ denote the distribution of H_τ conditioned on $\text{EQ}_{(P,T)}$ when v arrives. Here $n = |V|$ is the number of vertices in the input graph. For every $w \in V$ let $F_w^{(0)}$ denote the indicator of w being free when v arrives (unconditionally) and let $F_w^{(n)}$ denote the indicator variables of w being free when v arrives conditioned on $\text{EQ}_{(P,T)}$. Note that $F^{(0)}$ is determined by $H_\tau^{(0)}$ and $F^{(n)}$ is determined by $H_\tau^{(n)}$. For $t = 0, \dots, n$, we define distributions $H_\tau^{(t)}$ that interpolate between $H_\tau^{(0)}$ and $H_\tau^{(n)}$ as follows.

As in Lemma 3.15, for every $w \in V$ we denote the unconditional distribution of its primary choice by p_w , and the unconditional distribution of its secondary choice by s_w . Similarly, we denote the conditional distribution given $\text{EQ}_{(P,T)}$ of the primary choice by \tilde{p}_w and the conditional distribution of the secondary choice by \tilde{s}_w . For every $t = 0, \dots, n$ the primary choice of vertices $w_j, j = 1, \dots, t$ are sampled independently from \tilde{p}_{w_j} , and the primary choices of vertices $w_j, j = t+1, \dots, n$ are sampled independently from the unconditional distribution p_{w_t} . Similarly, secondary choices of vertices $w_j, j = 1, \dots, t$ are sampled independently from \tilde{s}_{w_j} and secondary choices of vertices $w_j, j = t+1, \dots, n$ are sampled independently from s_{w_j} . Note that $H_\tau^{(0)}$ is sampled from the unconditional distribution of H_τ , and $H_\tau^{(n)}$ is sampled from the conditional distribution (conditioned on $\text{EQ}_{(P,T)}$), as required, due to the independence of the conditional probabilities \tilde{p}_{w_j} and \tilde{s}_{w_j} , by Lemma 3.15. For $t = 0, \dots, n$ let M_t denote the matching constructed by our algorithm on $H_\tau^{(t)}$, and let $F_w^{(t)}$ be the indicator variable for w being free when v arrives in the DAG sampled from $H_\tau^{(t)}$.

Coupling the distributions of $H_\tau^{(t)}$. We now exhibit a coupling between the $H_\tau^{(t)}, t = 0, \dots, n$. Specifically, we will show that for every such t the following holds.

$$\mathbf{E} \left[\sum_{q \in V} |F_q^{(t+1)} - F_q^{(t)}| \right] \leq 4z(R(w_{t+1})), \quad (24)$$

where $R(w_{t+1})$ is as defined in Lemma 3.15 with regard to the certified primary path $R(P, T)$. Recall that $z(R(w_{t+1}))$ is the total probability assigned to arcs leaving w_{t+1} which are ruled out from being primary arcs in G_τ by conditioning on $\text{EQ}_{(P,T)}$.

We construct the coupling by induction. The base case corresponds to $t = 0$ and is trivial. We now give the inductive step ($t \rightarrow t+1$). We write $w := w_{t+1}$ to simplify notation. Let $Z^p \in N_w(w)$ denote the primary choice of w in $H_\tau^{(t)}$, and let $Z^s \in N_w(w)$ denote the secondary choice of w in $N_w(w)$ (they are sampled according to the unconditional distributions p_w and s_w respectively). Let $\tilde{Z}^p \in N_w(w)$ and $\tilde{Z}^s \in N_w(w)$ be sampled from the conditional distributions \tilde{p}_w and \tilde{s}_w respectively, such that that the joint distributions (Z^p, \tilde{Z}^p) and (Z^s, \tilde{Z}^s) satisfy

$$\Pr[Z^p \neq \tilde{Z}^p] = \text{TV}(p_w, \tilde{p}_w) \quad \text{and} \quad \Pr[Z^s \neq \tilde{Z}^s] = \text{TV}(s_w, \tilde{s}_w). \quad (25)$$

First, we note that if $Z^p = \tilde{Z}^p$ and $Z^s = \tilde{Z}^s$, then $w = w_{t+1}$ is matched to the same neighbor under $H_\tau^{(t)}$ and $H_\tau^{(t+1)}$, and so $M_t = M_{t+1}$, due to the greedy nature of the matching constructed. Otherwise, by Lemma 3.9, at most two vertices have different matched status in M_t and M_{t+1} in the latter case (in the former case every vertex has the same matched status). To summarize, we

have, for $R(w)$ determined by (P, T) as in Lemma 3.15, that

$$\begin{aligned} \mathbf{E} \left[\sum_{q \in V} |F_q^{(t+1)} - F_q^{(t)}| \right] &\leq 2 \cdot \Pr[Z^p \neq \tilde{Z}^p \text{ or } Z^s \neq \tilde{Z}^s] \\ &\leq 2(\text{TV}(p_w, \tilde{p}_w) + \text{TV}(s_w, \tilde{s}_w)) \quad (\text{by (25) and union bound}) \\ &\leq 4z(R(w)). \quad (\text{by Lemma 3.15}) \end{aligned} \tag{26}$$

This concludes the proof of the inductive step, and establishes (24). In particular, we get

$$\begin{aligned} \mathbf{E} \left[\sum_{q \in V} |F_q^{(n)} - F_q^{(0)}| \right] &\leq \sum_{t=0}^{n-1} \mathbf{E} \left[\sum_{q \in V} |F_q^{(t+1)} - F_q^{(t)}| \right] \\ &\leq \sum_{t=0}^{n-1} 4z(R(w_{t+1})) \quad (\text{by (26)}) \\ &= 4z(R(P, T)), \end{aligned} \tag{27}$$

by the definition of $R(P, T) = \bigcup_w R(w)$ in Lemma 3.15.

We now finish the claim. First note that for any (P, T) such that P has length at most $2000 \cdot \ln(1/\varepsilon)$ and $z(B(P, T)) \leq 2 \ln(1/\varepsilon)$ one has $\sum_w z(R(w)) = z(R(P, T)) = O(\ln(1/\varepsilon))$. Indeed, by Lemma 3.15 and linearity of z , recalling that u is the root of P and that no vertex appears after v (and thus $B(v, u) = \{(w, u) \mid w \text{ arrives between } u \text{ and } v\}$), we have

$$z(R(P, T)) \leq z(B(P, T)) + z(B(v, u)) + \sum_{w \in P \cup \{w : T = (w, w')\}} z(\{w\} \times N_w(w)). \tag{28}$$

We now bound the contribution to the above upper bound on $\sum_w z(R(w)) = z(R(P, T))$ in (28). First, we have that $z(B(P, T)) \leq 2 \ln(1/\varepsilon)$ by assumption of the lemma. To bound the contribution of $z(B(v, u))$, we note that by Property IV, we have

$$c \leq \Pr[F_u] = \prod_{e \in B(v, u)} (1 - z_e) \leq \exp \left(- \sum_{(w, u) \in B(v, u)} z(w, u)/2 \right) \leq \exp(-z(B(v, u))/2),$$

because any arc $e = (w, u)$ appears as a primary arc in G_τ with probability $z'(w, u) \geq z(w, u)/2$, independently of other such arcs, and the appearance of any such an edge implies that u has an incoming primary edge in H_τ when v arrives and is therefore matched; i.e., the event F_u is false in this case. We thus have $z(B(v, u)) \leq 2 \ln(1/c)$. Finally, it remains to note that for every one of the at most $2000 \cdot \ln(1/\varepsilon) + 1$ vertices $w \in P \cup \{w : T = (w, w')\}$ the contribution of $z(\{w\} \times N_w(w))$ to the right hand side of (28) is at most $1 + C\varepsilon \leq 2$, by Lemma 3.10, (2). Putting these bounds together, we get that for sufficiently small ε ,

$$z(R(P, T)) \leq 2 \ln(1/\varepsilon) + 2 \ln(1/c) + 2 \cdot 2000 \cdot \ln(1/\varepsilon) + 2 = O(\ln(1/\varepsilon)). \tag{29}$$

The term we wish to upper bound is at most

$$\begin{aligned}
& \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w \mid \text{EQ}_{(P,T)}] - \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w] \\
& \leq \left(\max_{w \in N_v(v)} z'_w \right) \cdot \sum_{w \in N_v(v)} \left| \Pr[F_w \mid \text{EQ}_{(P,T)}] - \Pr[F_w] \right| \\
& \leq C\sqrt{\varepsilon} \cdot \sum_{w \in N_v(v)} \left| \Pr[F_w \mid \text{EQ}_{(P,T)}] - \Pr[F_w] \right| \quad (\text{by Lemma 3.10, (3)}) \\
& = C\sqrt{\varepsilon} \cdot \mathbf{E} \left[\sum_{w \in N_v(v)} |F_w^{(n)} - F_w^{(0)}| \right] \quad (\text{by definition of } F^{(0)} \text{ and } F^{(n)})
\end{aligned}$$

then, using (27) and (29), we find that the term we wish to upper bound is at most

$$\begin{aligned}
& \leq C\sqrt{\varepsilon} \cdot \mathbf{E} \left[\sum_{w \in V} |F_w^{(n)} - F_w^{(0)}| \right] \\
& \leq C\sqrt{\varepsilon} \cdot z(R(P, T)) \quad (\text{by (27)}) \\
& = O(\sqrt{\varepsilon} \cdot \log(1/\varepsilon)) \quad (\text{by (29)}) \\
& \leq \varepsilon^{1/3}/2,
\end{aligned}$$

completing the proof. \square

Finally, we obtain Lemma 3.19 by combining Claim 3.20 and Claim 3.21, to find that, as claimed

$$\begin{aligned}
(18) & \leq \sum_{(P,T) \in \mathcal{P}} \Pr[\text{EQ}_{(P,T)} \mid F_u] \left(\sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w \mid \text{EQ}_{(P,T)}] - \sum_{w \in N_v(v)} z'_w \cdot \Pr[F_w] \right) + \varepsilon^{1/3}/2 \\
& \leq \varepsilon^{1/3}/2 + \varepsilon^{1/3}/2 = \varepsilon^{1/3}. \quad \square
\end{aligned}$$

3.4.6 Bounding the Impact of Bad Vertices

In this section, we show that we can completely ignore the bad vertices without losing too much. From the definition of good vertices, for a bad vertex v , we have that

$$\Pr_{\tau}[H_{\tau} \text{ has a primary path rooted at } v \text{ of length at least } 2000 \cdot \ln(1/\varepsilon)] \geq \varepsilon^6.$$

As the main result of this section, we prove the following theorem:

Theorem 3.22. *The number of bad vertices is at most $\varepsilon^3 \cdot \sum_{e \in E} x_e$.*

To prove this, we first describe a charging mechanism in which, for each bad vertex, a charge of one is distributed among a subset of other vertices. Then, using the following supplementary lemma, we show that the total distributed charge over all vertices in the graph is at most $\varepsilon^3 \cdot \sum_{(u,v) \in E} x_{uv}$.

Lemma 3.23. *We call a primary path P a primary predecessor path of v if it ends at v . That is, $P = v_{\ell} \rightarrow v_{\ell-1} \rightarrow \dots \rightarrow v_1 = v$. We have*

$$\Pr_{\tau}[v \text{ has any primary predecessor path } P \text{ with } z(B(P)) \leq 20 \cdot \ln(1/\varepsilon) \text{ and } |P| \geq 1000 \cdot \ln(1/\varepsilon)] \leq \varepsilon^{10}.$$

Proof. We use the principle of deferred decisions and traverse the path backwards. Let b be the current vertex, which is initially set to v . Consider all incoming arcs to b , say $(a_1, b), \dots, (a_k, b)$ where we index a 's by time of arrival; i.e., a_i arrives before a_j if $i < j$ (and b arrived before any a_i).

First consider the random choice of a_1 and see if it selected the arc (a_1, b) .

- If it does, then the path including b in H_τ will use the arc (a_1, b) .
- Otherwise, if a_1 does not select the arc (a_1, b) , then go on to consider a_2 and so on.

If no a_1, \dots, a_k selects b , then the process stops; i.e., the primary path starts at this vertex since b has no incoming primary arc. Otherwise let i be the first index so that (a_i, b) was selected. Then (a_i, b) is in the primary path ending at v in H_τ . Now, observe that no a_1, \dots, a_{i-1} may be in the path in this case, because these vertices arrived before a_i and after b . Moreover, we have not revealed any randomness regarding a_{i+1}, \dots, a_k that may appear later in the path. We can therefore repeat the above process with b now set to a_i and “fresh” randomness for all vertices we consider, as the random choices of arcs of all vertices are independent. We now show that this process, with good probability, does not result in a long predecessor path P of low $z(B(P))$ value.

Recall from Lemma 3.10, (5), that $z(u, v) \leq 3/5$ for all $(u, v) \in V \times V$. Suppose that $\sum_{i=1}^k z(a_i, b) \geq 4/5$. Let j be the first index such that $\sum_{i=1}^j z(a_i, b) \geq 1/5$. Thus $\sum_{i=1}^j z(a_i, b) \leq 4/5$, and hence the probability that none of the first j vertices select b is at least $\prod_{i=1}^j (1 - z(a_i, b)) \geq 1 - \sum_{i=1}^j z(a_i, b) \geq 1/5$. Consequently, with probability at least $1/5$, vertex b either has no predecessor or the increase to $z(B(P))$ is at least $1/5$.

In the other case, we have $\sum_{i=1}^k z(a_i, b) \leq 4/5$. Then the probability that b has no predecessor is $\prod_{i=1}^k (1 - z(a_i, b)) \geq 1 - \sum_{i=1}^k z(a_i, b) \geq 1/5$.

Therefore, at any step in the above random process, with probability at least $1/5$, we either stop or increase $z(B(P))$ by $1/5$. Let Z_i be an indicator variable for the random process either stopping or increasing $z(B(P))$ by at least $1/5$ at step i , and notice that according to the above random process, each Z_i is lower bounded by an independent Bernoulli variable with probability $1/5$. Thus if we define $Z = \sum_{i \in [1000 \cdot \ln(1/\varepsilon)]} Z_i$, we have $\mathbf{E}[Z] \geq 200 \cdot \ln(1/\varepsilon)$, and thus by standard coupling arguments and Chernoff bounds, we have that

$$\Pr[Z \leq 100 \cdot \ln(1/\varepsilon)] \leq \Pr[Z \leq (1 - 1/2) \cdot \mathbf{E}[Z]] \leq e^{-(1/2) \cdot (1/2)^2 \cdot 200 \cdot \ln(1/\varepsilon)} \leq \varepsilon^{10}.$$

But if the path does not terminate within $1000 \cdot \ln(1/\varepsilon)$ steps and $Z \geq 100 \cdot \ln(1/\varepsilon)$, then $z(B(P)) \geq 20 \cdot \ln(1/\varepsilon)$. \square

We now prove Theorem 3.22.

Proof of Theorem 3.22. By Lemma 3.16, the probability that H_τ has a primary path P with $z(B(P)) \geq 20 \cdot \ln(1/\varepsilon)$ starting at v is at most ε^{10} . Thus, for a bad vertex u , the probability that H_τ has some primary path P rooted at u with $|P| \geq 2000 \cdot \ln(1/\varepsilon)$ and $z(B(P)) \leq 20 \cdot \ln(1/\varepsilon)$ is at least $\varepsilon^6 - \varepsilon^{10} \geq \varepsilon^6/2$.

Let $k = 20 \cdot \ln(1/\varepsilon)$ and $\ell = 2000 \cdot \ln(1/\varepsilon)$. Let \mathcal{P}_u be the set of all primary paths P rooted at u such that $z(B(P)) \leq k$ and $|P| = \ell$ starting at u . Since all such primary paths with length more than ℓ are extensions of those with length exactly ℓ , we have $\sum_{P \in \mathcal{P}_u} \Pr[P \text{ is in } H_\tau] \geq \varepsilon^6/2$. For each such path $P \in \mathcal{P}_u$, consider the two vertices w_ℓ^P and $w_{\ell-1}^P$ at distances ℓ and $\ell - 1$ respectively from u . For each such vertex w_j^P ($j \in \{\ell - 1, \ell\}$), charge $(2/\varepsilon^6) \cdot \Pr[P \text{ is in } H_\tau] \cdot y_{w_j^P}$. Then the sum of these charges is

$$\sum_{P \in \mathcal{P}_u} (2/\varepsilon^6) \cdot \Pr[P \text{ is in } H_\tau] \cdot \underbrace{(y_{w_\ell^P} + y_{w_{\ell-1}^P})}_{\geq 1} \geq (2/\varepsilon^6) \cdot \sum_{P \in \mathcal{P}_u} \Pr[P \text{ is in } H_\tau] \geq 1.$$

Notice that the fact $(y_{w_\ell^P} + y_{w_{\ell-1}^P}) \geq 1$ follows because y_w 's form a feasible dual solution (to the vertex cover problem).

On the other hand, consider how many times each vertex is charged. For this, for every vertex w , let \mathcal{Q}_w be the set of primary predecessor paths Q of w such that $|Q| = \ell - 1$ and $z(B(P)) \leq k$. As $|Q| = \ell - 1 \geq 1000 \cdot \ln(1/\varepsilon)$ for all $Q \in \mathcal{Q}_w$, by Lemma 3.23, $\sum_{Q \in \mathcal{Q}_w} \Pr[Q \text{ is in } H_\tau] \leq \varepsilon^{10}$. For a primary predecessor path $Q \in \mathcal{Q}_w$ (or one of its extensions), the vertex w can be charged at most twice according to the above charging mechanism. Since any predecessor path of w with length more than $\ell - 1$ must be an extension of one with length exactly $\ell - 1$, we have that the amount w is charged is at most

$$\sum_{Q \in \mathcal{Q}_w} 2 \cdot 2 \cdot \Pr[Q \text{ is in } H_\tau] \cdot y_w / \varepsilon^6 \leq 4 \cdot (\varepsilon^{10} / \varepsilon^6) \cdot y_w \leq 4 \cdot \varepsilon^4 \cdot y_w.$$

Summing over all $w \in V$ and using Lemma 3.3, the total charge is at most

$$\sum_{w \in v} 4 \cdot \varepsilon^4 \cdot y_w \leq 4 \cdot \varepsilon^4 \cdot \beta \cdot \sum_{e \in E} x_e \leq \varepsilon^3 \sum_{e \in E} x_e. \quad \square$$

3.4.7 Calculating the Competitive Ratio of Algorithm 3

We now show that the competitive ratio of Algorithm 3 is indeed $(1/2 + \alpha)$ competitive for some sufficiently small absolute constant $\alpha > 0$, thus proving Theorem 1.2. This essentially combines the facts that for good vertices, the matching probability is very close to the fractional values of incident edges, and that the number of bad vertices is very small compared to the total value of the fractional algorithm (over the entire graph).

Proof of Theorem 1.2. Let OPT denote the size of the maximum cardinality matching in the input graph G . Then, by Lemma 3.5 and our choice of $f = f_{1+2\varepsilon}$ and $\beta = 2 - \varepsilon \geq \beta^*(f_{1+2\varepsilon})$, we have that $\sum_e x_e \geq (1/\beta) \cdot \text{OPT} \geq (1/2 + \varepsilon/4) \cdot \text{OPT}$, where the x_e 's are the fractional values we compute in Algorithm 3.

Now let M be the matching output by Algorithm 3. We have

$$\begin{aligned} \mathbf{E}[|M|] &= \sum_{e \in E} \Pr[e \text{ is matched}] \\ &\geq \sum_{\text{good } v \in V} (1 - \varepsilon^2) \cdot \sum_{u \in N_v(v)} x_{uv} && \text{(By Theorem 3.18)} \\ &\geq (1 - \varepsilon^2) \cdot \left(\sum_{e \in E} x_e - \sum_{\text{bad } v \in V} \sum_{u \in N_v(v)} x_{uv} \right) \\ &\geq (1 - \varepsilon^2) \cdot \left(\sum_{e \in E} x_e - \sum_{\text{bad } v \in V} 1 \right) && \left(\sum_{u \in N_v(v)} x_{uv} \leq 1 \right) \\ &\geq (1 - \varepsilon^2) \cdot \left(\sum_{e \in E} x_e - \varepsilon^3 \sum_{e \in E} x_e \right) && \text{(By Theorem 3.22)} \\ &\geq (1 - 2\varepsilon^2) \cdot \sum_{e \in E} x_e \\ &\geq (1 - 2\varepsilon^2) \cdot (1/2 + \varepsilon/4) \cdot \text{OPT} \\ &\geq (1/2 + \varepsilon/5) \cdot \text{OPT}, \end{aligned}$$

where the last line holds for a sufficiently small constant $\varepsilon > 0$. □

Appendix

A Deferred Proofs of Section 3.3

Here we prove that a change of the realized arc choices of any vertex does not change the matched status of more than two vertices (at any point in time). This is Lemma 3.9, restated below.

Lemma 3.9. *Let G_τ and $G_{\tau'}$ be two realizations of the random digraph where all the vertices in the two graphs make the same choices except for one vertex v . Then the number of vertices that have different matched status (free/matched) in the matchings computed in H_τ and $H_{\tau'}$ at any point of time is at most two.*

Proof. We consider the evolution, following each vertex arrival, of the matchings M_τ and $M_{\tau'}$ computed in H_τ and $H_{\tau'}$, respectively, as well as the set of vertices with different matched status in these matchings, denoted by $D := (M_\tau \setminus M_{\tau'}) \cup (M_{\tau'} \setminus M_\tau)$. The set D is empty before the first arrival and remains empty until the arrival of v , as all earlier vertices than v have the same primary and secondary arcs and have the same set of free neighbors in H_τ and $H_{\tau'}$ (as $D = \emptyset$, by induction). Now, if immediately after v arrives it remains free in both M_τ and $M_{\tau'}$, or it is matched to the same neighbor in both matchings, then clearly D remains empty. Otherwise, either v is matched to different neighbors in M_τ and $M_{\tau'}$, or v is matched in one of these matchings but not in the other. Both these cases result in $|D| = 2$. We now show by induction that the cardinality of D does not increase following subsequent arrivals, implying the lemma.

Let u be some vertex which arrives after v . If when u arrives u is matched to the same neighbor w in M_τ and $M_{\tau'}$ or if u remains free in both matchings, then D is unchanged. If u is matched to some w on arrival in M_τ , but not in $M_{\tau'}$, then since the arcs of u are the same in G_τ and $G_{\tau'}$, this implies that w must have been free in M_τ but not in $M_{\tau'}$, and so $D \ni w$. Therefore, after u arrives, we have $D \leftarrow (D \setminus \{w\}) \cup \{u\}$, and so D 's cardinality is unchanged. Finally, if u is matched to two distinct neighbors, denoted by w and w' , respectively, then one of (u, w) and (u, w') must be the primary arc of u in both G_τ and $G_{\tau'}$. Without loss of generality, say (u, w) is this primary arc. Since u is matched to w in M_τ but not in $M_{\tau'}$, then w must be free in M_τ when u arrives, but not in $M_{\tau'}$, and so $D \ni w$. Consequently, we have that after u arrives we have $D \leftarrow S$ for some set $S \subseteq (D \setminus \{w\}) \cup \{w'\}$, and so D 's cardinality does not increase. \square

B Deferred Proofs of Section 3.1

Here we prove the bound on the fractional degree x_u in terms of its dual value, restated below.

Lemma 3.3. *For any vertex $u, v \in V$, let y_u be the potential of u prior to arrival of v . Then the fractional degree just before v arrives, $x_u := \sum_{w \in N_v(u)} x_{uw}$, is bounded as follows:*

$$\frac{y_u}{\beta} \leq x_u \leq \frac{y_u + f(1 - y_u)}{\beta}.$$

Proof. Let y_0 be u 's potential after u 's arrival. For the lower bound, note that it suffices to prove that every increase in the fractional degree is bounded below by the increase in the potential divided by β . When vertex u first arrived, we consider two cases.

1. $y_0 > 0$ (thus $y_0 = 1 - \theta > 0$, and so $\theta < 1$), then the increase in u 's fractional degree was:

$$\sum_{v \in N_u(u)} \frac{(\theta - y_v)^+}{\beta} \left(1 + \frac{1 - \theta}{f(\theta)} \right) = \frac{f(\theta) + 1 - \theta}{\beta} = \frac{f(1 - y_0) + y_0}{\beta} \geq \frac{y_0}{\beta}.$$

2. $y_0 = 0$ (thus $\theta = 1$), then the increase in u 's fractional degree was:

$$\sum_{v \in N_u(u)} \frac{(\theta - y_v)^+}{\beta} \left(1 + \frac{1 - \theta}{f(\theta)}\right) = \sum_{v \in N_u(u)} \frac{(\theta - y_v)^+}{\beta} \geq 0 = \frac{y_0}{\beta}.$$

For every subsequent increase of the fractional degree due to a newly-arrived vertex we have that:

$$\frac{(\theta - y_u^{old})^+}{\beta} \left(1 + \frac{1 - \theta}{f(\theta)}\right) \geq \frac{(\theta - y_u^{old})^+}{\beta},$$

Which concludes the proof for the lower bound.

For the upper bound, by [25, Invariant 1], we have that

$$\beta \cdot x_u \leq y_c + f(1 - y_0) + \int_{y_0}^{y_c} \frac{1 - x}{f(x)} dx. \quad (30)$$

This upper bound can be simplified by using Equation (5), as follows. Taking (30), adding and subtracting $1 + f(1 - y_u)$ and writing the integral $\int_{y_0}^{y_u} \frac{1-x}{f(x)} dx$ as the difference of two integrals $\int_{y_0}^1 \frac{1-x}{f(x)} dx - \int_{y_u}^1 \frac{1-x}{f(x)} dx$, and relying on Equation (5), we find that

$$\begin{aligned} \beta \cdot x_u &\leq y_c + f(1 - y_0) + \int_{y_0}^{y_c} \frac{1 - x}{f(x)} dx \\ &= \left(1 + f(1 - y_0) + \int_{y_0}^1 \frac{1 - x}{f(x)} dx\right) - 1 + y_c + \int_1^{y_c} \frac{1 - x}{f(x)} dx \\ &= \beta^*(f) + y_c - \left(1 + f(1 - y_c) + \int_{y_c}^1 \frac{1 - x}{f(x)} dx\right) + f(1 - y_c) \\ &= \beta^*(f) + y_c - \beta^*(f) + f(1 - y_c) \\ &= y_c + f(1 - y_c), \end{aligned}$$

from which the lemma follows. □

C Deferred Proofs of Section 3.4.2

In this section we present the proofs deferred from Section 3.4.2. We start by presenting a more manageable form for the function $f = f_{1+2\varepsilon}$ which we use.

A function in the WW family is determined by a parameter $k \geq 1$ and takes the following form

$$f_\kappa(\theta) = \left(\frac{1 + \kappa}{2} - \theta\right)^{\frac{1+\kappa}{2\kappa}} \left(\theta + \frac{\kappa - 1}{2}\right)^{\frac{\kappa-1}{2\kappa}}.$$

Letting $\kappa = 1 + 2\varepsilon$, we get that $f := f_\kappa$ is of the form

$$\begin{aligned} f(\theta) &= (1 + \varepsilon - \theta)^{\frac{1+\varepsilon}{1+2\varepsilon}} \cdot (\theta + \varepsilon)^{\frac{\varepsilon}{1+2\varepsilon}} \\ &= (1 + \varepsilon - \theta) \cdot \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta}\right)^{\frac{\varepsilon}{1+2\varepsilon}}. \end{aligned}$$

Clearly this is water filling when $\varepsilon = 0$ and otherwise we have that the first term is like water filling and then the second term is less than 1 for $z \leq 1/2$ and greater than 1 if $z > 1/2$.

By Taylor expansion, we obtain the following more manageable form for f .

Lemma C.1. *There exists $\varepsilon_0 \in (0, 1)$ such that for every $\varepsilon \in (0, \varepsilon_0)$ and every $\theta \in [0, 1]$, we have*

$$f(\theta) \leq (1 - \theta) \left(1 + \varepsilon \ln \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta} \right) \right) + 1.01\varepsilon.$$

Proof. Taking the Taylor expansion of e^x , we find that

$$\begin{aligned} f(\theta) &= (1 + \varepsilon - \theta) \cdot \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta} \right)^{\frac{\varepsilon}{1+2\varepsilon}} \\ &= (1 + \varepsilon - \theta) \cdot \sum_{i=0}^{\infty} \frac{\left(\ln \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta} \right) \cdot \frac{\varepsilon}{1+2\varepsilon} \right)^i}{i!} \\ &= (1 + \varepsilon - \theta) \left(1 + \ln \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta} \right) \cdot \frac{\varepsilon}{1+2\varepsilon} \right) + o(\varepsilon) \\ &= (1 + \varepsilon - \theta) + (1 - \theta) \ln \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta} \right) \cdot \frac{\varepsilon}{1+2\varepsilon} + o(\varepsilon) \\ &= (1 + \varepsilon - \theta) + (1 - \theta)\varepsilon \ln \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta} \right) + o(\varepsilon) \\ &= (1 - \theta) \left(1 + \varepsilon \ln \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta} \right) \right) + \varepsilon + o(\varepsilon). \end{aligned}$$

To be precise, for $\theta \in [0, 1]$ and $0 < \varepsilon \leq \varepsilon_0 \leq 1$ (implying for example $\frac{\theta + \varepsilon}{1 + \varepsilon - \theta} \leq \frac{2}{\varepsilon}$), we will show that terms dropped in the third, fourth and fifth lines are all at most some $O((\ln(\frac{1}{\varepsilon}) \cdot \varepsilon)^2) = o(\varepsilon)$, from which the lemma follows as the sum of these terms is at most 0.01ε for $\varepsilon \leq \varepsilon_0$ and ε_0 sufficiently small.

Indeed, in the third line, we dropped

$$(1 + \varepsilon - \theta) \cdot \sum_{i=2}^{\infty} \frac{\left(\ln \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta} \right) \cdot \frac{\varepsilon}{1+2\varepsilon} \right)^i}{i!} \leq 2 \cdot \sum_{i=2}^{\infty} \frac{(\ln(\frac{2}{\varepsilon}) \cdot \varepsilon)^i}{i!} \leq \sum_{i=2}^{\infty} \frac{(\ln(\frac{2}{\varepsilon}) \cdot \varepsilon)^i}{i^2} = O((\ln(1/\varepsilon) \cdot \varepsilon)^2),$$

where the last step used that $\ln(1/\varepsilon) \cdot \varepsilon \leq 1$ holds for all $\varepsilon \geq 0$. In the fourth line, we dropped

$$\varepsilon \cdot \ln \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta} \right) \cdot \frac{\varepsilon}{1+2\varepsilon} \leq \varepsilon^2 \cdot \ln(2/\varepsilon) = O((\ln(1/\varepsilon) \cdot \varepsilon)^2).$$

Finally, in the fifth line, we dropped

$$(1 - \theta) \cdot \left(\varepsilon - \frac{\varepsilon}{1+2\varepsilon} \right) \cdot \ln \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta} \right) \leq 1 \cdot (\varepsilon^2/(1+2\varepsilon)) \cdot \ln(2/\varepsilon) = O((\ln(1/\varepsilon) \cdot \varepsilon)^2). \quad \square$$

Given this more manageable form for f , we can now turn to prove Lemma 3.10, restated below.

Lemma 3.10. *(Basic bounds on conditional probabilities z_u) There exist absolute constants $c \in (0, 1)$ and $C > 1/c > 1$ and $\varepsilon_0 \in (0, 1)$ such that for every $\varepsilon \in (0, \varepsilon_0)$ the following holds: for every vertex $v \in V$, if y_u is the dual variable of a neighbor $u \in N_v(v)$ before v 's arrival and θ is the value chosen by Algorithm 1 on v 's arrival, then for z_u as defined in Algorithm 3, we have:*

- (1) If $\theta \notin (c, 1 - c)$, then $\sum_{u \in N_v(v)} z_u \leq 1$,
- (2) If $\theta \in [0, 1]$, then $\sum_{u \in N_v(v)} z_u \leq 1 + C\varepsilon$,

- (3) If $\sum_{u \in N_v(v)} z_u > 1$, then $z_u \leq C\sqrt{\varepsilon}$ for every $u \in N_v(v)$,
- (4) If $\sum_{u \in N_v(v)} z_u > 1$, then for every $u \in N_v(v)$ such that $z_u > 0$, one has $y_u \in [c/2, 1 - c/2]$, and
- (5) For all $u \in N_v(v)$, one has $z_u \leq 1/2 + O(\sqrt{\varepsilon})$.

Proof. We begin by getting a generic upper bound for z_u . We note that each edge e is matched by Algorithm 3 with probability at most x_e by Line 12. Therefore, u is matched before v arrives with probability at most $x_u := \sum_{w \in N_v(u) \setminus \{v\}} x_{wu}$, the fractional degree of u before v arrives. Therefore, by Lemma 3.3, the probability that u is free is at least

$$\Pr[u \text{ free when } v \text{ arrives}] \geq 1 - x_u \geq 1 - \frac{y_u + f(1 - y_u)}{\beta}, \quad (31)$$

from which, together with the definition of $x_{uv} = \frac{1}{\beta}(\theta - y_u)^+ \left(1 + \frac{1-\theta}{f(\theta)}\right)$, we obtain the following upper bound on z_u :

$$z_u = \frac{x_{uv}}{\Pr[u \text{ is free when } v \text{ arrive}]} \leq \frac{\frac{1}{\beta}(\theta - y_u)^+ \left(1 + \frac{1-\theta}{f(\theta)}\right)}{1 - \frac{y_u + f(1 - y_u)}{\beta}} = \frac{(\theta - y_u) \left(1 + \frac{1-\theta}{f(\theta)}\right)}{\beta - (y_u + f(1 - y_u))}. \quad (32)$$

We start by upper bounding $\sum_{u \in N_v(v)} z_u$, giving a bound which will prove useful in the proofs of both (1) and (2). Recall that θ is defined as the largest $\theta \leq 1$ such that

$$\sum_{u \in N_v(v)} (\theta - y_u)^+ \leq f(\theta). \quad (33)$$

Summing (32) over all $u \in N_v(v)$, we find that

$$\begin{aligned} \sum_{u \in N_v(v)} z_u &\leq \sum_{u \in N_v(v)} \frac{(\theta - y_u)^+ \cdot \left(1 + \frac{1-\theta}{f(\theta)}\right)}{\beta - (\theta + f(1 - \theta))} \quad (f(\cdot) \text{ is non-increasing, by Observation 3.4}) \\ &\leq \frac{f(\theta) + 1 - \theta}{\beta - (\theta + f(1 - \theta))} \quad (\text{by (33) and } \beta \geq \beta^*(f) = 1 + f(0) \geq \theta + f(1 - \theta)) \end{aligned}$$

We therefore wish to upper bound $\frac{f(\theta) + 1 - \theta}{\beta - \theta - f(1 - \theta)}$. To this end let $\gamma(\theta, \varepsilon) := \varepsilon \ln \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta} \right)$. Before proceeding to the proof, it would be useful to summarize some properties of the function $\gamma(\theta, \varepsilon)$.

1. $\gamma(\theta, \varepsilon) = -\gamma(1 - \theta, \varepsilon)$ for all $\theta \in [0, 1]$.
2. For c, ε_0 sufficiently small we have for all $\theta \in [0, c]$ that $\gamma(\theta, \varepsilon) \leq \varepsilon \ln \left(\frac{c + \varepsilon}{1 + \varepsilon - c} \right) \leq -20 \cdot \varepsilon$, and for all $\theta \in (1 - c, 1]$ that $\gamma(\theta, \varepsilon) \geq \varepsilon \ln \left(\frac{1 - c + \varepsilon}{1 + \varepsilon - (1 - c)} \right) \geq 20 \cdot \varepsilon$.
3. $\gamma(\theta, \varepsilon) \cdot (1 - 2\theta) \leq 0$ for $\theta \in [0, 1]$, since $\gamma(\theta, \varepsilon) \leq 0$ for $\theta \leq 1/2$ and $\gamma(\theta, \varepsilon) \geq 0$ for $\theta \geq 1/2$.
4. $\theta \cdot \gamma(\theta, \varepsilon) \geq -\varepsilon$ for all $\theta \in [0, 1]$.

The last property follows from $\ln \left(\frac{1 + \varepsilon - \theta}{\theta + \varepsilon} \right) \leq \ln \left(\frac{1 + \varepsilon + \theta}{\theta + \varepsilon} \right) \leq \ln \left(1 + \frac{1}{\theta + \varepsilon} \right) \leq \frac{1}{\theta + \varepsilon} \leq \frac{1}{\theta}$, which implies in particular that $\theta \cdot \gamma(\theta, \varepsilon) = \theta \cdot \varepsilon \cdot \left(-\ln \left(\frac{1 + \varepsilon - \theta}{\theta + \varepsilon} \right) \right) \geq -\varepsilon$.

We will use γ as shorthand for $\gamma(\theta, \varepsilon)$. Recalling that $\beta = 2 - \varepsilon$ and using Lemma C.1, we have:

$$\begin{aligned}
\frac{f(\theta) + 1 - \theta}{\beta - (\theta + f(1 - \theta))} &\leq \frac{(1 - \theta) \left(1 + \varepsilon \ln \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta}\right)\right) - \theta + 1 + 1.01\varepsilon}{2 - \varepsilon - \theta - \theta \left(1 + \varepsilon \ln \left(\frac{1 - \theta + \varepsilon}{\theta + \varepsilon}\right)\right) - 1.01\varepsilon} \\
&\leq \frac{(1 - \theta)(2 + \gamma) + 2\varepsilon}{2 - 2\theta + \theta\gamma - 3\varepsilon} \\
&= 1 + \frac{\gamma(1 - 2\theta) + 5\varepsilon}{2 - 2\theta + \theta\gamma - 3\varepsilon}.
\end{aligned} \tag{34}$$

We will continue by proving that the second term is negative. First we prove that the denominator is positive. To this end, first consider the case when $\theta \in [0, c)$. In this case for ε_0, c sufficiently small one has that: $2 - 2\theta + \theta\gamma - 2\varepsilon > 2 - 2\theta - \varepsilon - 2\varepsilon > 0$ from Item 4. Moreover, when $\theta \in (1 - c, 1]$ one has that $\theta > \frac{1}{2}$ (since c is small) and $\gamma \geq 20\varepsilon$ from Item 2. Thus $2 - 2\theta + \theta\gamma - 2\varepsilon \geq \theta\gamma - 2\varepsilon \geq \frac{1}{2} \cdot 20\varepsilon - 3\varepsilon = 7\varepsilon > 0$. Now, it remains to prove that the numerator is always negative. When $\theta \in [0, c)$ we have that $1 - 2\theta \geq 3/4$ (since c is small) and $\gamma \leq -20\varepsilon$ from Item 2, therefore $\gamma(1 - 2\theta) + 5\varepsilon \leq \gamma \cdot \frac{3}{4} + 5 \cdot (-\frac{\gamma}{20}) = \frac{\gamma}{2} < 0$. In the case where $\theta \in (1 - c, 1]$, we have that $1 - 2\theta < -3/4$, and $\theta > 1/2$ (since c is small), and $\gamma \geq 20\varepsilon$ from Item 2, thus $\gamma(1 - 2\theta) + 5\varepsilon \leq -\frac{3}{4} \cdot 20\varepsilon + 5\varepsilon = -10\varepsilon < 0$.

We now turn to (2). We assume that $\theta \in (c, 1 - c)$, since otherwise the claim is trivial, by (1). We have by (34) that $\frac{f(\theta) + 1 - \theta}{\beta - (\theta + f(1 - \theta))} \leq 1 + \frac{\gamma(1 - 2\theta) + 5\varepsilon}{2 - 2\theta + \theta\gamma - 3\varepsilon}$. We have that $\gamma(1 - 2\theta) + 5\varepsilon \leq 5\varepsilon$ from Item 3. Furthermore, using Item 4 we have that $2 - 2\theta + \theta\gamma - 3\varepsilon \geq 2c + -4\varepsilon > c$ for a sufficiently small ε_0 . Overall, the second term is bounded above by $\frac{5}{c} \cdot \varepsilon < C \cdot \varepsilon$, for $C > \frac{5}{c} > \frac{1}{c}$ as required.

We now prove (3). Note that by (1), $\sum_{u \in N_v(v)} z_u > 1$ implies that $\theta \in (c, 1 - c)$. Now, for every $u \in N_v(v)$, let $\alpha_u := \frac{(\theta - y_u)^+}{f(\theta)}$, so that $y_u = \theta - f(\theta) \cdot \alpha_u$ if $y_u \leq \theta$. We also note that by definition of α_u and our choice of θ , we have $\sum_{u \in N_v(v)} \alpha_u = \sum_{u \in N_v(v)} \frac{(\theta - y_u)^+}{f(\theta)} \leq 1$. In the proof of (3) and (4) we will assume for notational simplicity that all $u \in N_v(v)$ have $y_u \leq \theta$, implying $z_u \geq 0$. Summing up (32) over all $u \in N_v(v)$ and substituting in α_u , we thus find that

$$\begin{aligned}
\sum_{u \in N_v(v)} z_u &\leq \sum_{u \in N_v(v)} \frac{(\theta - y_u)^+ (1 + \frac{1 - \theta}{f(\theta)})}{\beta - (y_u + f(1 - y_u))} \\
&= \sum_{u \in N_v(v)} \alpha_u \cdot \frac{f(\theta) + 1 - \theta}{\beta - (y_u + f(1 - y_u))} \\
&\leq \sum_{u \in N_v(v)} \alpha_u \cdot \frac{f(\theta) + 1 - \theta}{2 - y_u - f(1 - y_u) - 2.01\varepsilon} \quad (\text{by Lemma C.1 and } \beta = 2 - \varepsilon) \\
&\leq \sum_{u \in N_v(v)} \alpha_u \cdot \frac{f(\theta) + 1 - \theta}{2 - 4\varepsilon - 2y_u},
\end{aligned}$$

In the last transition we used again (as in Item 4) that $y_u \cdot \varepsilon \ln \left(\frac{1 - y_u + \varepsilon}{y_u + \varepsilon}\right) \leq \varepsilon$, which implies $f(\theta) \leq 1 - \theta + \varepsilon$ for all $\theta \in [0, 1]$. Substituting $y_u = \theta - f(\theta) \cdot \alpha_u$ into the above upper bound on

$\sum_{u \in N_v(v)} z_u$, we get

$$\begin{aligned} \sum_{u \in N_v(v)} z_u &\leq \sum_{u \in N_v(v)} \alpha_u \cdot \frac{f(\theta) + 1 - \theta}{2 - 4\varepsilon - 2\theta + 2f(\theta) \cdot \alpha_u} \\ &= \sum_{u \in N_v(v)} \alpha_u \cdot \frac{f(\theta) + 1 - \theta}{2 - 4\varepsilon - 2\theta} - \sum_{u \in N_v(v)} \frac{(f(\theta) + 1 - \theta) \cdot 2f(\theta) \cdot \alpha_u^2}{(2 - 4\varepsilon - 2\theta) \cdot (2 - 4\varepsilon - 2\theta + 2f(\theta) \cdot \alpha_u)}, \end{aligned} \quad (35)$$

using the elementary identity $\frac{1}{a+b} = \frac{1}{a} - \frac{b}{a(a+b)}$ for appropriate a and b . Now, both terms in the last line of (35) can be significantly simplified, as follows. For the former term, again using that $f(\theta) \leq 1 - \theta + \varepsilon$, together with $\sum_{u \in N_v(v)} \alpha_u \leq 1$ noted above, we find that

$$\sum_{u \in N_v(v)} \alpha_u \cdot \frac{f(\theta) + 1 - \theta}{2 - 4\varepsilon - 2\theta} \leq \sum_{u \in N_v(v)} \alpha_u \cdot \frac{2 + \varepsilon - 2\theta}{2 - 4\varepsilon - 2\theta} = \sum_{u \in N_v(v)} \alpha_u \cdot \left(1 + \frac{5\varepsilon}{2 - 4\varepsilon - 2\theta}\right) \leq 1 + O(\varepsilon), \quad (36)$$

where in the last step we used that $\theta \leq 1 - c$ and c is some fixed constant. For the second term in the last line of (35), we note that

$$\sum_{u \in N_v(v)} \frac{(f(\theta) + 1 - \theta) \cdot 2f(\theta) \cdot \alpha_u^2}{(2 - 4\varepsilon - 2\theta) \cdot (2 - 4\varepsilon - 2\theta + 2f(\theta) \cdot \alpha_u)} = \Omega(1) \cdot \left(\sum_{u \in N_v(v)} \alpha_u^2 \right). \quad (37)$$

To see this, first note that for $\theta \in (c, 1 - c)$, the numerator of each summand of the LHS is at least $2f(c)^2 \cdot \alpha_u^2 \geq \Omega(\alpha_u^2)$, since f is decreasing by Observation 3.4 and $f(c) \geq \frac{1}{2} \cdot (1 + \varepsilon - c) \geq \Omega(1)$ for c and ε sufficiently small. To verify the first inequality of this lower bound for $f(c)$, recall that $f(c) = (1 + \varepsilon - c) \cdot \left(\frac{c + \varepsilon}{1 + \varepsilon - c}\right)^{\frac{\varepsilon}{1 + 2\varepsilon}}$. Now, for ε tending to zero and $c < 1/2$, the term $\left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta}\right)^{\frac{\varepsilon}{1 + 2\varepsilon}}$ tends to one as ε tends to zero. Therefore for ε sufficiently small we have $f(c) \geq \frac{1}{2} \cdot (1 + \varepsilon - c)$ for all $c < 1/2$. We now turn to upper bounding the denominator of each summand in the LHS of Equation (37). Indeed, substituting $y_u = \theta - f(\theta) \cdot \alpha_u$, we find that each such denominator is at most $(2 - 4\varepsilon - 2\theta) \cdot (2 - 4\varepsilon - 2\theta + 2f(\theta) \cdot \alpha_u) \leq (1/2) \cdot (2 - 4\varepsilon - 2y_u) \leq (1/2) \cdot (2 - 4\varepsilon - 2c) \leq O(1)$ for c and ε sufficiently small. Note that both numerator and denominator are positive for sufficiently small c and ε_0 . Substituting the bounds of (36) and (37) into (35), we obtain

$$\sum_{u \in N_v(v)} z_u \leq 1 + O(\varepsilon) - \Omega(1) \cdot \left(\sum_{u \in N_v(v)} \alpha_u^2 \right). \quad (38)$$

From Eq. (38) and $\sum_{u \in N_v(v)} z_u > 1$ by assumption of (3), we get that

$$\sum_{u \in N_v(v)} \alpha_u^2 \leq C' \varepsilon \quad (39)$$

for an absolute constant $C' > 1$, since otherwise $\sum_{u \in N_v(v)} z_u \leq 1$. Finally, it remains to note that

$$\begin{aligned}
\sum_{u \in N_v(v)} z_u^2 &= \sum_{u \in N_v(v)} \left(\frac{\alpha_u \cdot (f(\theta) + 1 - \theta)}{\beta - (y_u + f(1 - y_u))} \right)^2 \\
&\leq \left(\sum_{u \in N_v(v)} \alpha_u^2 \right) \cdot \left(\frac{f(\theta) + 1 - \theta}{\beta - (\theta + f(1 - \theta))} \right)^2 && \text{(by Observation 3.4 and } y_u \leq \theta) \\
&\leq \left(\sum_{u \in N_v(v)} \alpha_u^2 \right) \cdot \left(\frac{f(\theta) + 1 - \theta}{\beta - (1 - c + f(c))} \right)^2 && \text{(by Observation 3.4 and } \theta \leq 1 - c) \\
&\leq \left(\sum_{u \in N_v(v)} \alpha_u^2 \right) \cdot \left(\frac{1 - \theta + \varepsilon + 1 - \theta}{\beta - (1 - c + 1 - c + \varepsilon)} \right)^2 && (f(c) \leq 1 - c + \varepsilon) \\
&\leq \left(\sum_{u \in N_v(v)} \alpha_u^2 \right) \cdot \frac{2}{2c - 2\varepsilon} \\
&\leq C\varepsilon,
\end{aligned}$$

for some constant $C \geq \frac{2}{2c - 2\varepsilon}$. Thus $z_u^2 \leq \sum_{u \in N_v(v)} z_u \leq C\varepsilon$ and so $z_u \leq \sqrt{C \cdot \varepsilon} \leq C\sqrt{\varepsilon}$, as claimed.

We now prove (4). Since $\sum_{u \in N_v(v)} z_u > 1$ implies $\theta \in (c, 1 - c)$ by (1), using the definition of α_u 's from the proof of (3) together with the fact that $\alpha_u \leq C'\sqrt{\varepsilon}$ for every $u \in N_v(v)$ by (39) and the fact that $f(\theta) \leq 2$ for all $\theta \in [0, 1]$ (by Lemma C.1), we get that

$$y_u = \theta - f(\theta) \cdot \alpha_u \in [c - O(\sqrt{\varepsilon}), 1 - c] \subseteq [c/2, 1 - c/2],$$

for sufficiently small $\varepsilon_0 > 0$, as required.

As for (5), simplifying (32) and using the fact that $\theta - y_u \leq f(\theta)$, we get

$$z_u \leq \frac{\theta - y_u + 1 - \theta}{\beta - y_u - f(1 - y_u)} = \frac{1 - y_u}{\beta - y_u - f(1 - y_u)}.$$

Recall from Lemma C.1 that for all $\theta \in [0, 1]$, we have $f(\theta) \leq (1 - \theta) \left(1 + \varepsilon \ln \left(\frac{\theta + \varepsilon}{1 + \varepsilon - \theta} \right) \right) + 1.01\varepsilon$, which implies the following:

1. For all $\theta \in [0, 1]$, we have $f(\theta) \leq 1 - \theta + \sqrt{\varepsilon}$, and
2. For $\theta < e^{-10}$, we have $f(\theta) \leq (1 - \theta)(1 + \varepsilon(\ln((e^{-10} + \varepsilon)/(1 - e^{-10} + \varepsilon))) + 1.01\varepsilon) \leq 1 - \theta - 2\varepsilon$.

Suppose that $y_u \leq 1 - e^{-10}$. Then using Item 1, we have

$$\begin{aligned}
z_u &\leq \frac{1 - y_u}{\beta - y_u - f(1 - y_u)} \leq \frac{1 - y_u}{2 - \varepsilon - y_u - y_u - \sqrt{\varepsilon}} \\
&\leq \frac{1 - y_u}{2(1 - y_u) - 2\sqrt{\varepsilon}} \leq 1/2 + \frac{2\sqrt{\varepsilon}}{2e^{-10} - 2\sqrt{\varepsilon}} \leq 1/2 + O(\sqrt{\varepsilon}).
\end{aligned}$$

Now suppose that $y_u > 1 - e^{-10}$. Then $1 - y_u < e^{-10}$, and so by Item 2, $f(1 - y_u) \leq 1 - y_u - 2\varepsilon$. Thus we have

$$z_u \leq \frac{1 - y_u}{\beta - y_u - f(1 - y_u)} \leq \frac{1 - y_u}{2 - \varepsilon - y_u - (y_u - 2\varepsilon)} = \frac{1 - y_u}{2(1 - y_u) + \varepsilon} \leq 1/2,$$

completing the proof. \square

Finally, we rely on Lemma C.1 to prove that the fractional solution maintained by Line 4 is $1/\beta$ competitive, as implied by Lemma 3.5 and the following restated fact.

Fact 3.12. *For all sufficiently small $\varepsilon > 0$, we have that $2 - \varepsilon \geq \beta^*(f_{1+2\varepsilon})$.*

Proof. Let us denote as before $f = f_{1+2\varepsilon}$. Recall that $\beta^*(f) = 1 + f(0)$. By Lemma C.1, this is at most $1 + f(0) \leq 1 + \left(1 + \varepsilon \ln\left(\frac{\varepsilon}{1+\varepsilon}\right)\right) + 1.01\varepsilon$. But for small enough ε , we have that $\ln\left(\frac{\varepsilon}{1+\varepsilon}\right) \leq -2.01$, implying that $1 + f(0) \leq 2 - \varepsilon$, as claimed. \square

References

- [1] AGEEV, A. A. AND SVIRIDENKO, M. I. 2004. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization* 8, 3, 307–328.
- [2] BIRNBAUM, B. AND MATHIEU, C. 2008. On-line bipartite matching made simple. *ACM SIGACT News* 39, 1, 80–87.
- [3] BUCHBINDER, N., SEGEV, D., AND TKACH, Y. 2018. Online algorithms for maximum cardinality matching with edge arrivals. *Algorithmica*, 1–19.
- [4] CHIPLUNKAR, A., TIRODKAR, S., AND VISHWANATHAN, S. 2015. On randomized algorithms for matching in the online preemptive model. In *Proceedings of the 23rd Annual European Symposium on Algorithms (ESA)*. 325–336.
- [5] COHEN, I. R. AND WAJC, D. 2018. Randomized online matching in regular graphs. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 960–979.
- [6] DEVANUR, N. R., JAIN, K., AND KLEINBERG, R. D. 2013. Randomized primal-dual analysis of ranking for online bipartite matching. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 101–107.
- [7] EDEN, A., FELDMAN, M., FIAT, A., AND SEGAL, K. 2018. An economic-based analysis of ranking for online bipartite matching. *arXiv preprint arXiv:1804.06637*.
- [8] EDMONDS, J. 1965a. Maximum matching and a polyhedron with 0, 1-vertices. *Journal of research of the National Bureau of Standards B* 69, 125-130, 55–56.
- [9] EDMONDS, J. 1965b. Paths, trees, and flowers. *Canadian Journal of mathematics* 17, 3, 449–467.
- [10] EPSTEIN, L., LEVIN, A., SEGEV, D., AND WEIMANN, O. 2013. Improved bounds for online preemptive matching. In *Proceedings of the 30th International Symposium on Theoretical Aspects of Computer Science (STACS)*. 389.
- [11] FEIGE, U. 2018. Tighter bounds for online bipartite matching. *arXiv preprint arXiv:1812.11774*.
- [12] GANDHI, R., KHULLER, S., PARTHASARATHY, S., AND SRINIVASAN, A. 2006. Dependent rounding and its applications to approximation algorithms. *Journal of the ACM (JACM)* 53, 3, 324–360.

- [13] GOEL, G. AND MEHTA, A. 2008. Online budgeted matching in random input models with applications to adwords. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 982–991.
- [14] GURUGANESH, G. P. AND SINGLA, S. 2017. Online matroid intersection: Beating half for random arrival. In *Proceedings of the 19th Conference on Integer Programming and Combinatorial Optimization (IPCO)*. 241–253.
- [15] HUANG, Z., KANG, N., TANG, Z. G., WU, X., ZHANG, Y., AND ZHU, X. 2018a. How to match when all vertices arrive online. In *Proceedings of the 50th Annual ACM Symposium on Theory of Computing (STOC)*. 17–29.
- [16] HUANG, Z., PENG, B., TANG, Z. G., TAO, R., WU, X., AND ZHANG, Y. 2019. Tight competitive ratios of classic matching algorithms in the fully online model. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2875–2886.
- [17] HUANG, Z., TANG, Z. G., WU, X., AND ZHANG, Y. 2018b. Online vertex-weighted bipartite matching: Beating $1-1/e$ with random arrivals. In *Proceedings of the 45th International Colloquium on Automata, Languages and Programming (ICALP)*. 1070–1081.
- [18] KARP, R. M., VAZIRANI, U. V., AND VAZIRANI, V. V. 1990. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*. 352–358.
- [19] KUHN, H. W. 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly* 2, 1-2, 83–97.
- [20] LEE, E. AND SINGLA, S. 2017. Maximum matching in the online batch-arrival model. In *Proceedings of the 19th Conference on Integer Programming and Combinatorial Optimization (IPCO)*. 355–367.
- [21] LOVÁSZ, L. AND PLUMMER, M. D. 2009. *Matching theory*. Vol. 367. American Mathematical Society.
- [22] MEHTA, A. 2013. Online matching and ad allocation. *Foundations and Trends® in Theoretical Computer Science* 8, 4, 265–368.
- [23] SCHRIJVER, A. 2003. *Combinatorial optimization: polyhedra and efficiency*. Vol. 24. Springer Science & Business Media.
- [24] TIRODKAR, S. AND VISHWANATHAN, S. 2017. Maximum matching on trees in the online preemptive and the incremental dynamic graph models. In *Proceedings of the 23rd International Computing and Combinatorics Conference (COCOON)*. 504–515.
- [25] WANG, Y. AND WONG, S. C.-w. 2015. Two-sided online bipartite matching and vertex cover: Beating the greedy algorithm. In *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP)*. 1070–1081.