

The Greedy Algorithm is *not* Optimal for On-Line Edge Coloring

Amin Saberi and David Wajc

Stanford University
{saber, wajc}@stanford.edu

Abstract

Nearly three decades ago, Bar-Noy, Motwani and Naor showed that no online edge-coloring algorithm can edge color a graph optimally. Indeed, their work, titled “the greedy algorithm is optimal for on-line edge coloring”, shows that the competitive ratio of 2 of the naïve greedy algorithm is best possible online. However, their lower bound required bounded-degree graphs, of maximum degree $\Delta = O(\log n)$, which prompted them to conjecture that better bounds are possible for higher-degree graphs. While progress has been made towards resolving this conjecture for restricted inputs and arrivals or for random arrival orders, an answer for fully general *adversarial* arrivals remained elusive.

We resolve this thirty-year-old conjecture in the affirmative, presenting a $(1.9 + o(1))$ -competitive online edge coloring algorithm for general graphs of degree $\Delta = \omega(\log n)$ under vertex arrivals. At the core of our results, and of possible independent interest, is a new online algorithm which rounds a fractional bipartite matching x online under vertex arrivals, guaranteeing that each edge e is matched with probability $(1/2 + c) \cdot x_e$, for a constant $c > 0.027$.

1 Introduction

An edge coloring of a graph is a decomposition of its edge-set into few vertex-disjoint edge-sets (matchings), or *colors*. Edge coloring a graph of maximum degree Δ trivially requires at least Δ colors, and this is tight for bipartite graphs, by the century-old result of König [28]. For general graphs, Δ colors are not always sufficient (e.g., in odd-length cycles), yet $\Delta + 1$ colors are always sufficient, by Vizing’s Theorem [35].

Algorithmically matching, or approximating, the optimal $\Delta(+1)$ colors needed to edge color a graph has been the focus of much concentrated effort, for numerous computational models. These include offline, online, distributed, parallel, and dynamic algorithms (see, e.g., [7, 8, 9, 11, 13, 25, 31, 34, 36] and references therein). These different models’ specific challenges naturally impose limitations on the attainable approximations. For example, Holyer’s Theorem [20] rules out efficient offline algorithms for computing an optimal edge coloring in general graphs, unless $P=NP$.

For online algorithms, the challenge is in making immediate and irrevocable decisions concerning edges’ colors after only part of the input is revealed. For example, the input graph can either be revealed edge-by-edge (edge arrivals) or vertex-by-vertex (vertex arrivals), and an online algorithm must assign colors to edges after they are revealed, immediately and irrevocably. The measure of an online algorithm is its competitive ratio, which is the worst-case ratio of the number of colors used by the algorithm to those of the optimal offline algorithm, namely, Δ or $\Delta + 1$.

In both the edge-arrival and vertex-arrival settings, a simple greedy algorithm has competitive ratio 2. The natural question, then, is whether a better online algorithm exists. Some thirty years ago, Bar-Noy, Motwani and Naor [4] showed that this competitive ratio of 2 is best possible, and no online algorithm (randomized or deterministic) can do better, in either arrival model.

However, noting that their result only holds for bounded-degree n -node graphs, of maximum degree $\Delta = O(\log n)$, Bar-Noy et al. conjectured that better algorithms exist for graphs of sufficiently high maximum degree.

Conjecture 1.1 ([4]). *There exists a $(2 - \Omega(1))$ -competitive online edge coloring algorithm under vertex arrivals in n -node graphs of maximum degree $\Delta = \omega(\log n)$.*

Bar-Noy et al. conjectured that the same holds under the more challenging edge-arrival model, and that moreover a $(1 + o(1))$ -competitive algorithm exists. These conjectures remain out of reach, though progress has been made on them over the years. For edge arrivals, a positive resolution of the stronger conjecture was achieved under the assumption of *random order* arrivals, where the input is generated adversarially, but its arrival order is randomly permuted by nature [1, 3, 5]. For adversarial vertex arrivals, Cohen et al. [9] showed that for *bipartite* graphs under one-sided vertex arrivals (vertices of one side are given, and the other side's vertices arrive), the conjectured $(1 + o(1))$ -competitive ratio is achievable for $\Delta = \omega(\log n)$. Whether the competitive ratio of 2 of the greedy algorithm is optimal under *general* vertex arrivals, in *general* graphs, however, remained open.

We answer the above open question, resolving Conjecture 1.1 in the affirmative.

Theorem 1.2. *There exists an online edge coloring algorithm which is $(1.897 + o(1))$ -competitive w.h.p. on general n -node graphs with maximum degree $\Delta = \omega(\log n)$ under vertex arrivals.*

Remark 1. *For general Δ , the $o(1)$ term in the above theorem is of the form $\gamma\sqrt{\log n/\Delta}$, for some constant $\gamma > 0$. This implies a better than two approximation ratio for sufficiently large $\Delta = O(\log n)$. For simplicity of exposition, we do not elaborate on this point.*

1.1 Techniques

To obtain our results, we combine and extend several previous algorithmic ideas.

Our starting point is the following natural recursive approach, due to Karloff and Shmoys [25], which reduces edge coloring a general graph G to edge coloring random bipartite subgraphs. Their idea was to assign each vertex to either side of a random subgraph uniformly, resulting in a bipartite subgraph H of G with maximum degree $\Delta/2 + o(\Delta)$ for $\Delta = \omega(\log n)$, by standard tail bounds. Consequently, applying an α -approximate algorithm to the random bipartite graph and recursing on the remaining edges is easily shown to result in an edge coloring using $\alpha \cdot \Delta/2 + o(\Delta) + \alpha \cdot \Delta/4 + o(\Delta) \cdots = \alpha \cdot \Delta + o(\Delta)$ colors. Importantly for us, this approach, originally used in the context of NC algorithms by [25], is implementable online, by sampling the random bipartitions in advance. (See Appendix A.)

At this point, one might be tempted to use the online algorithm of Cohen et al. [9] for these random bipartite subgraphs. Unfortunately, the reduction of Karloff and Shmoys [25] applied to online edge coloring with general vertex arrivals requires an online algorithm for bipartite graphs with *interleaved* arrivals, and not one-sided arrivals, as handled by [9]. To instantiate the Karloff-Shmoys approach, we therefore present a $(2 - c)$ -competitive edge coloring algorithm for interleaved vertex arrivals in bipartite graphs, which, when combined with the approach of [25], then extends to general graphs.

To obtain an edge-coloring algorithm for bipartite graphs under interleaved vertex arrival, we extend the approach of Cohen et al. [9], who showed that an $(\alpha + o(1))$ -competitive edge coloring can be achieved by repeatedly applying a matching algorithm which matches each edge with

probability $(1/\alpha)/\Delta$. For each vertex of degree $\Delta(1 - o(1))$, such a matching results in v being matched with probability $(1/\alpha) \cdot (1 - o(1))$. Repeating the above a super-logarithmic number of times (making use of $\Delta = \omega(\log n)$) therefore decreases the maximum degree of the graph at a rate of roughly one per α colors used. Cohen et al. used this approach with $\alpha = 1 + o(1)$, using an online matching algorithm from [10], on bipartite graphs under one-sided arrivals. We observe that this approach extends to arbitrary α and any arrival model, including interleaved vertex arrivals in bipartite graphs. (See Appendix B.)

Motivated by the above discussion, we design an online matching algorithm for bipartite graphs under interleaved arrivals, which matches each edge with probability $(1/2 + c)/\Delta$, for some constant $c > 0$. More generally, and of possible independent interest, we design an online rounding algorithm for bipartite fractional matchings under interleaved vertex arrivals, with a multiplicative factor of $1/2 + c$. That is, we show how, given a bipartite graph G and a fractional matching x in G revealed vertex-by-vertex, one can output a randomized matching which matches each edge e in G with probability $(1/2 + c) \cdot x_e$. This extends a similar online rounding algorithm previously developed by the authors with Papadimitriou and Pollner [33] in the context of online stochastic optimization, but which only works under one-sided vertex arrivals, and is therefore insufficient for our needs. This new rounding algorithm is the technical meat of this paper, and is presented in Section 3.

Combining the above, we obtain Theorem 1.2, and the positive resolution of Conjecture 1.1.

1.2 Related Work

The first positive results for online edge coloring were under random order edge arrivals. In this setting, Aggarwal et al. [1] showed that a $(1 + o(1))$ -competitive ratio is achievable in dense multigraphs with maximum degree $\Delta = \omega(n^2)$. Bahmani et al. [3] then showed that the greedy algorithm is sub-optimal for any graph of maximum degree $\Delta = \omega(\log n)$. Achieving the best of both these results, Bhattacharya et al. [5] recently obtained a $(1 + o(1))$ -competitive algorithm for graphs of maximum degree $\Delta = \omega(\log n)$. As stated above, the only prior algorithm which outperforms the greedy algorithm under *adversarial* arrivals is the algorithm of Cohen et al. [9] for bipartite graphs under one-sided vertex arrivals. In this work, we remove the assumption of bipartiteness and one-sided arrivals, and show how to outperform greedy in general graphs under arbitrary vertex arrivals.

Our work also ties into the long line of work on online matching, initiated by Karp, Vazirani and Vazirani [26]. (See e.g., [2, 16, 18, 19, 21, 32] and references therein and [29] for a survey of earlier work.) Historically, most research on online matching considered bipartite graphs with one-sided arrivals, due to applications in Internet advertising [17, 30]. A recent line of work considers such problems subject to interleaved vertex arrivals (motivated by more dynamic two-sided markets), as well as vertex arrivals in general graphs [2, 19, 21, 22, 37]. Our rounding algorithm for bipartite graphs with interleaved arrivals adds to the list of tools for tackling problems in this space.

Few of the works in the online (bipartite) matching literature rely on randomized rounding. At first blush, this seems surprising, given the integrality of the bipartite fractional matching polytope, and the multitude of competitive fractional algorithms for problems in this area [6, 17, 21, 22, 24, 37]. However, as pointed out in [12] and elaborated upon in [10], lossless rounding of a fractional matching x is impossible in online settings. In particular, outputting a matching \mathcal{M} which matches each edge e in a bipartite graph with probability $\Pr[e \in \mathcal{M}] = x_e$ is impossible in online settings, though it is easy to do offline. A natural question, then, is what is the highest value of $\alpha < 1$ for which one can guarantee $\Pr[e \in \mathcal{M}] \geq \alpha \cdot x_e$ when rounding bipartite fractional matchings online. The batched OCRS of Ezra et al. [15] gives $\alpha = 1/2$, unfortunately too low for our purposes. In prior work [33], motivated by a variation of the online Bayesian selection problem, we improve this bound to $\alpha = 0.51$, though only for one-sided arrivals, which is insufficient for our needs here. In this work we generalize this result, achieving a slightly higher $\alpha = 0.527$, subject to *interleaved* vertex arrivals.

2 Preliminaries

The underlying (a priori unknown) input to our problem is an n -node graph $G = (V, E)$ of maximum degree Δ (with n and Δ both known). The vertices of G are revealed over time. For notational convenience, we associate the $n := |V|$ vertices with the numbers in $[n]$ by order of appearance, and denote by $u < v$ the fact that u arrives before v . When a vertex v arrives (at *time* v), all its edges (u, v) to its previously-arrived neighbors $u < v$ are revealed. After v arrives, and before arrival of vertex $v + 1$, an online edge coloring algorithm must decide, irrevocably, which color to assign to all edges (u, v) with $u < v$. The objective is to minimize the number of distinct colors used.

As outlined in the introduction, we will rely on the ability to edge color general graphs by recursively coloring random bipartite subgraphs, as first proposed by Karloff and Shmoys [25], in the context of NC algorithms. The extension and proof for online settings is essentially the same, and is provided, for completeness, in Appendix A.

Lemma 2.1. *(Implied by [25]) Given an online edge coloring algorithm which is α -competitive w.h.p. on bipartite graphs of maximum degree $\Delta = \omega(\log n)$ under interleaved vertex arrivals, there exists an online edge coloring algorithm which is $(\alpha + o(1))$ -competitive w.h.p. on general graphs of maximum degree $\Delta = \omega(\log n)$ under vertex arrivals.*

The following lemma, implied by the recent work of Cohen et al. [9], reduces α -competitive edge coloring to online matching algorithms which match each edge with probability $(1/\alpha)/\Delta$. The proof is provided, for completeness, in Appendix B.

Lemma 2.2. *(Implied by [9]) Let \mathcal{A} be an online matching algorithm which on any (bipartite) graph of maximum degree $\Delta \leq \Delta'$ under vertex arrivals, matches each edge with probability at least $1/(\alpha\Delta')$. Then, there exists an online edge coloring algorithm \mathcal{A}' which is $(\alpha + o(1))$ -competitive w.h.p. for (bipartite) graphs of maximum degree $\Delta = \omega(\log n)$ under vertex arrivals.*

Motivated by Lemma 2.2, we show how to (approximately) round fractional matchings online. These are assignments of nonnegative $x_e \geq 0$ to edges $e \in E$, satisfying the fractional matching constraint, $\sum_{e \ni v} x_e \leq 1$ for all $v \in V$. This is a fractional relaxation of the matching constraint, which stipulates that the degree of any vertex in a matching be at most one. Fittingly, we refer to $\sum_{w < v} x_{u,w}$ as the *fractional degree* of u before arrival of v (or at its arrival time, if $u = v$). We shall show how to round fractional matchings up to a multiplicative error of $\alpha < 2$. This rounding subroutine applied to the fractional matching assigning value $1/\Delta$ to each edge of the graph thus matches each edge with probability $1/(\alpha\Delta)$. Combined with lemmas 2.1 and 2.2, this yields our $(\alpha + o(1))\Delta$ coloring algorithm.

2.1 Negative Association

In our work we will need to bound positive correlations between variables. At the core of these proofs will be a use of *negatively associated* random variables. This section introduces this notion of negative dependence and its properties which we use.

Definition 2.3 ([23, 27]). *Random variables X_1, \dots, X_n are negatively associated (NA) if every two monotone nondecreasing functions f and g defined on disjoint subsets of the variables in \vec{X} are negatively correlated. That is,*

$$\mathbb{E}[f \cdot g] \leq \mathbb{E}[f] \cdot \mathbb{E}[g]. \quad (1)$$

The following simple example of NA variables will prove useful for us.

Proposition 2.4 (0-1 Principle [14]). *Let $X_1, \dots, X_n \in \{0, 1\}$ be binary random variables satisfying $\sum_i X_i \leq 1$ always. Then, the variables X_1, \dots, X_n are NA.*

Negative association is closed under several operations, allowing to construct more elaborate NA distributions from simpler NA distributions as above (see [14, 23, 27]).

Proposition 2.5 (Independent Union). *Let X_1, \dots, X_n be NA and Y_1, \dots, Y_m be NA, with $\{X_i\}_i$ independent of $\{Y_j\}_j$. Then, the variables $X_1, \dots, X_n, Y_1, \dots, Y_m$ are all NA.*

Proposition 2.6 (Function Composition). *Let X_1, \dots, X_n be NA variables, and let f_1, \dots, f_k be monotone nondecreasing functions defined on disjoint subsets of the variables in \vec{X} . Then the variables $f_1(\vec{X}), \dots, f_k(\vec{X})$ are NA.*

An immediate corollary of negative association, obtained by considering the functions $f(\vec{X}) = X_i$ and $g(\vec{X}) = X_j$ for $i \neq j$, is pairwise negative correlation.

Proposition 2.7 (NA implies Negative Correlation). *Let X_1, \dots, X_n be NA variables. Then, for all $i \neq j$, we have that $\text{Cov}(X_i, X_j) \leq 0$.*

3 Rounding Bipartite Fractional Matchings Online

In this section we present an online algorithm which (approximately) rounds a bipartite fractional matching under interleaved vertex arrivals. In what follows, we let $c \geq 0.027$ be the largest value below 0.03 satisfying

$$(1/2 - c)(1 - 4c)(1/2 - c - 6c/(1/2 - c)) - 2c \geq 0. \quad (2)$$

We note that this choice of $c \leq 0.03$ also satisfies the following.¹

$$\min\{1/2 - c, 1 - 4c, 1 - 6c/(1/2 - c)^2\} \geq 0. \quad (3)$$

We show the following.

Theorem 3.1. *There exists an online algorithm which, given an (unknown) bipartite graph G under interleaved vertex arrivals, together with a fractional matching x in G , outputs a random matching \mathcal{M} matching each edge $e \in E$ with probability*

$$\Pr[e \in \mathcal{M}] = (1/2 + c) \cdot x_e \geq 0.527 \cdot x_e. \quad (4)$$

We now turn to describing the algorithm claimed by the above theorem.

3.1 Intuition and Algorithm

Before presenting our algorithm, we describe the approach used to obtain Theorem 3.1 under one-sided arrivals [33], and then discuss the new ideas needed to extend this result to interleaved arrivals.

Naturally, an edge (u, v) with $u < v$ (i.e., v arriving later than u) can only be matched if u is not already matched before the arrival of v . We denote by $F_{u,v}$ the event that u is free (i.e., is not matched in \mathcal{M}) prior to the arrival of v . The guarantee of Theorem 3.1 implies the following closed form for the probability of this event.

$$\Pr[F_{u,v}] = g(u, v) := 1 - \sum_{w < v} (1/2 + c) \cdot x_{u,w}. \quad (5)$$

¹We encourage the reader to think of $c \rightarrow 0$, and note that inequalities (2) and (3) hold for sufficiently small constant $c > 0$. Our choice of $c \approx 0.027$ is simply the largest satisfying all these constraints.

To achieve marginal probabilities of $\Pr[(u, v) \in \mathcal{M}] = (1/2 + c) \cdot x_{u,v}$, our first step is to have every arriving vertex v pick a random neighbor $u < v$ with probability $x_{u,v}$, and then, if u is free, we match (u, v) with probability $q_{u,v} := \min(1, (1/2 + c)/g(u, v))$. For neighbors u of low fractional degree upon arrival of v , i.e., $\sum_{w < v} x_{u,w} \leq \frac{1/2-c}{1/2+c}$, this last probability is precisely $q_{u,v} = (1/2+c)/\Pr[F_{u,v}]$. Consequently, we match each such edge (u, v) with probability $\Pr[(u, v) \in \mathcal{M}] = x_{u,v} \cdot \Pr[F_{u,v}] \cdot (1/2 + c)/\Pr[F_{u,v}] = (1/2 + c) \cdot x_{u,v}$, as desired. For edges (u, v) for which u has *high* fractional degree, on the other hand, this only gives us $\Pr[(u, v) \in \mathcal{M}] \geq (1/2 - c) \cdot x_{u,v}$, and this can be tight.

To increase the probability of an edge (u, v) to be matched to the desired $(1/2 + c) \cdot x_{u,v}$, we repeat this process a second time, making a second pick, if v is not matched after its first pick. Here, we must argue that the variables $\{F_{u,v} \mid u < v\}$ do not have strong positive correlation. Indeed, if, as an extreme case, we had $F_{u,v} = F_{w,v}$ always for all $u, w < v$, and v had only high-degree neighbors (for which $q_{u,v} = 1$), then if v is not matched to its first pick, then all its neighbors must be matched, and v is therefore never matched as a second pick. This implies that a second pick does not increase $\Pr[(u, v) \in \mathcal{M}]$ in this case. As shown in [33], under one-sided arrivals, this problematic scenario does not occur, since the matched status of neighbors of v is rather weak. For interleaved arrivals, however, the underlying argument does not carry through, as we now explain.

3.1.1 Extension to Interleaved Arrivals

The key difference between one-sided and interleaved arrivals is that now we require small positive correlation between the matched statuses of every two nodes on the same side of the bipartition, rather than just nodes on the “offline side”. For one-sided arrivals, the weak positive correlation between offline vertices was due to two factors. (1) low-degree offline vertices are matched only due to semi-adaptive matching choices, where precisely one neighbor of an arriving online vertex is picked, and at most one is matched. (That is, they are only matched as a first pick.) Therefore, by the 0-1 Principle (Proposition 2.4) and closure properties of NA distributions (propositions 2.5 and 2.6), the indicators for a vertex to be matched when it has low fractional degree are NA, and hence are negatively correlated. (2) On the other hand, the probability of a node to be matched when it has high degree is low, since each edge is matched with probability $(1/2 + c) \cdot x_{u,v}$, and the residual fractional degree when v has high degree is $1 - \frac{1/2-c}{1/2+c} = \frac{2c}{1/2+c} \leq 4c$. Putting (1) and (2) together, we find that the matched statuses of any two offline vertices have small correlation.

Unfortunately, under interleaved arrivals, the above is no longer true. In particular, if a vertex v has low fractional degree upon arrival, it may still be matched as a second pick upon arrival (due to its high-degree neighbors). Consequently, the indicators for vertices on the same side of the bipartition being matched when they have low fractional degree are no longer negatively associated, thus undoing the entire argument used to bound $\text{Cov}(F_{u,v}, F_{w,v})$ for vertices $u, w < v$ on the same side of the bipartition.

To overcome this problem, we have each arriving vertex v with low fractional degree upon arrival only pick once, and rely on its low fractional degree to pick each neighbor with higher probability. In particular, when such a vertex v arrives, we pick at most one neighbor with probability $x_{u,v} \cdot \frac{1/2+c}{1/2-c}$. (Since v has low fractional degree on arrival, $\sum_{u < v} x_{u,v} \leq \frac{1/2-c}{1/2+c}$, this is well-defined.) Then, if this picked vertex u is free, we match (u, v) with probability $\frac{1/2-c}{\Pr[F_{u,v}]} = \frac{1/2-c}{g(u,v)} (\leq 1)$, resulting in the edge (u, v) being matched with probability $x_{u,v} \cdot (1/2 + c)$. Crucially for our analysis, this now allows us to show that the indicators for vertices (in the same side of the graph) to be matched when they have low fractional degree is again negatively associated. This then results in the matched status of vertices again being decomposable into two variables, with the first being negatively correlated, and the second having low probability, from which we

obtain that vertices on the same side of the bipartition have low correlation.²

This discussion gives rise to Algorithm 1, which we prove in this section provides the guarantees of Theorem 3.1.

Algorithm 1

```

1: Init:  $\mathcal{M} \leftarrow \emptyset$ 
2: for all vertices  $v$ , on arrival do
3:   read  $\{x_{u,v} \mid u < v\}$ 
4:   if  $\sum_{u < v} x_{u,v} \leq \frac{1/2-c}{1/2+c}$  then
5:     pick at most one  $u < v$  with probability  $x_{u,v} \cdot \frac{1/2+c}{1/2-c}$ 
6:     if  $u \neq \text{nil}$  and  $u$  is unmatched in  $\mathcal{M}$  then
7:       with probability  $\frac{1/2-c}{g(u,v)}$  do
8:          $\mathcal{M} \leftarrow \mathcal{M} \cup \{(u, v)\}$ 
9:   else
10:    pick at most one  $u < v$  with probability  $x_{u,v}$ 
11:    if  $u \neq \text{nil}$  and  $u$  is unmatched in  $\mathcal{M}$  then
12:      with probability  $\min\left(1, \frac{1/2+c}{g(u,v)}\right)$  do
13:         $\mathcal{M} \leftarrow \mathcal{M} \cup \{(u, v)\}$ 
14:    if  $v$  is still unmatched in  $\mathcal{M}$  then
15:      pick at most one  $u < v$  with probability  $x_{u,v}$ 
16:      if  $u \neq \text{nil}$  and  $u$  is unmatched in  $\mathcal{M}$  then
17:        with probability  $p_{u,v}$  guaranteeing  $\Pr[(u, v) \in \mathcal{M}] = (1/2 + c) \cdot x_{u,v}$  do
18:           $\mathcal{M} \leftarrow \mathcal{M} \cup \{(u, v)\}$ 
19: Output  $\mathcal{M}$ 

```

3.2 High-Level Analysis

For our analysis and proof of Theorem 3.1, we will assume, by way of an inductive proof, that Equation (4) holds for all edges (u, w) with $u, w < v$ and therefore that for each $u < v$ we have $\Pr[F_{u,v}] = g(u, v)$, as stated in Equation (5).

Given the inductive hypothesis, it is easy to verify that Algorithm 1 guarantees marginal probabilities of each edge to be matched to be precisely $(1/2+c) \cdot x_e$. Indeed, for an arriving vertex v with low fractional degree, $\sum_{u < v} x_{u,v} \leq \frac{1/2-c}{1/2+c}$ (lines 4-8), since by the inductive hypothesis u is free at time v with probability $\Pr[F_{u,v}] = g(u, v)$, we have that

$$\Pr[(u, v) \in \mathcal{M}] = x_{u,v} \cdot \frac{1/2+c}{1/2-c} \cdot g(u, v) \cdot \frac{1/2-c}{g(u, v)} = (1/2+c) \cdot x_{u,v}.$$

In the alternative case of lines 9-18, we trivially have that each edge (u, v) with $u < v$ is matched with probability precisely $\Pr[(u, v) \in \mathcal{M}] = (1/2+c) \cdot x_e$, due to lines 17-18. The crux of the analysis, then, is in proving that this algorithm is well-defined, and in particular that there exists some probabilities $p_{u,v}$ as stated in Line 17.

We note that all probabilistic lines in the algorithm except for Line 17 are trivially well-defined. First, if v has low fractional degree before time v , i.e., $\sum_{u < v} x_{u,v} \leq \frac{1/2-c}{1/2+c}$, then the probability of any neighbor to be picked in Line 5 is at most $\sum_{u < v} x_{u,v} \cdot \frac{1/2+c}{1/2-c} \leq 1$, and so this

²We note that Gamlath et al. [19] followed a superficially similar rounding approach, using two choices. As they only required bounds on the (unweighted) matching's size, their analysis relied on showing that *globally* positive correlation is low. As we desire high matching probability on an edge-by-edge (or at least vertex-by-vertex) basis, we must follow a more delicate approach.

line is well-defined. Next, by the fractional matching constraint, we have that $\sum_{u < v} x_{u,v} \leq 1$, and consequently lines 10 and 15 are well-defined. Finally, by the fractional matching constraint, we have that $\sum_{w < v} (1/2 + c) \cdot x_{u,w} \leq 1/2 + c$, and therefore

$$\Pr[F_{u,v}] = g(u, v) \geq 1/2 - c. \quad (6)$$

Consequently, the term $\frac{1/2-c}{g(u,v)}$ in Line 7 is indeed a probability, by our choice of $c = 0.027 \leq 1/2$. We now turn to proving that probabilities $p_{u,v}$ as stated in Line 17 do indeed exist.

First, to show that $p_{u,v} \geq 0$, we must show that the probability of edge (u, v) to be matched as a first pick in Line 13 does not on its own exceed $(1/2 + c) \cdot x_{u,v}$.

Observation 3.2. *The probability of an edge (u, v) to be matched in Line 13 is at most*

$$\Pr[(u, v) \text{ added to } \mathcal{M} \text{ in Line 13}] \leq (1/2 + c) \cdot x_{u,v}.$$

Proof. By the inductive hypothesis, we have that $\Pr[F_{u,v}] = g(u, v)$. Consequently,

$$\Pr[(u, v) \text{ added to } \mathcal{M} \text{ in Line 13}] = x_{u,v} \cdot \min\left(1, \frac{1/2 + c}{g(u, v)}\right) \cdot g(u, v) \leq (1/2 + c) \cdot x_{u,v}. \quad \square$$

Corollary 3.3. *The parameter $p_{u,v}$ in Line 17 satisfies $p_{u,v} \geq 0$.*

The core of the analysis will then be in proving that $p_{u,v} \leq 1$. For this, we will need to argue that a second pick in lines 14-18 is likely to result in (u, v) being matched, provided we set $p_{u,v} \leq 1$ high enough. We prove as much in the next section.

3.3 Core of the Analysis

In this section we prove that the second pick is likely to result in a match. To this end, we prove that the matched statuses of neighbors of an arriving vertex v have low positive correlation (if any). More formally, if $G = (V_1, V_2, E)$ is our bipartite graph, we will prove the following.

Lemma 3.4. *For any $i = 1, 2$, vertex v and vertices $u, w < v$ with $u, w \in V_i$,*

$$\text{Cov}(F_{u,v}, F_{w,v}) \leq 6c.$$

Since the covariance of two binary variables A and B is equal to that of their complements, $\text{Cov}(A, B) = \text{Cov}(1 - A, 1 - B)$, we will concern ourselves with bounding $\text{Cov}(M_{u,v}, M_{w,v})$, where $M_{u,v} := 1 - F_{u,v}$ is an indicator for u being matched in \mathcal{M} before v arrives.

For this proof, we write $M_{u,v}$ as the sum of two Bernoulli variables, $M_{u,v} = M_{u,v}^L + M_{u,v}^H$. The indicators $M_{u,v}^L$ and $M_{u,v}^H$ correspond to u being matched to some neighbor w at a time z when u had low or high fractional degree, respectively. That is,

$$M_{u,v}^L := \mathbb{I}\left[(u, w) \in \mathcal{M} \text{ for some } w < v \text{ with } \sum_{z < \min\{u, w\}} x_{u,z} \leq \frac{1/2 - c}{1/2 + c}\right],$$

with $M_{u,v}^H = M_{u,v} - M_{u,v}^L$ defined analogously.

In what follows, we will show that for any vertex v and index $i = 1, 2$, the variables $\{M_{u,v}^L \mid u \in V_i\}$ are negatively correlated, while the variables $\{M_{u,v}^H \mid u \in V_i\}$ have low probability, which implies that they have low positive correlation with any other binary variable. These bounds will allow us to bound the correlation of the sums $M_{u,v} = M_{u,v}^L + M_{u,v}^H$.

We start by proving the negative correlation between $M_{u,v}^L$ variables, and indeed proving negative association of these variables.

Lemma 3.5. *For any $i = 1, 2$ and vertex v , the variables $\{M_{u,v}^L \mid u < v, u \in V_i\}$ are NA.*

By Proposition 2.7, this implies that the above variables are negatively correlated.

Corollary 3.6. *For any $i = 1, 2$, vertex v and earlier vertices $u, w < v$ with $u, w \in V_i$,*

$$\text{Cov}(M_{u,v}^L, M_{w,v}^L) \leq 0.$$

Proof of Lemma 3.5. Recall that $M_{u,v}^L$ is an indicator for u being matched before arrival of v before it has high fractional degree. By definition of Algorithm 1, this implies that a matching event accounted for by $M_{u,v}^L$ can only occur in lines 8 or 13. Such matches occur due to u picking a neighbor or being picked as a neighbor in line 5 or 10, and the probabilistic test in line 7 or 12 (respectively), passing, if the picked vertex was previously unmatched in \mathcal{M} . We imagine we perform the probabilistic tests in lines 7 and 12 *before* testing whether the picked vertex was unmatched in \mathcal{M} .

For vertices $w < z$, let $A_{w,z}$ be an indicator for z picking w in line 5 or 10, and the probabilistic test in line 7 or 12 (respectively) passing. Then, by the 0-1 Principle (Proposition 2.4), we have that for any vertex z , the variables $\{A_{w,z} \mid w < z\}$ are NA. Moreover, the families of variables $\{A_{w,z} \mid w < z\}$ for distinct z are NA. Therefore, by closure of NA under independent union (Proposition 2.5), the variables $\{A_{w,z} \mid z, w < z\}$ are NA. For notational simplicity, letting $A_{z,w} := A_{w,z}$ for $z > w$ (recall that we only defined $A_{w,z}$ for $w < z$), we find that if z' is the smaller of $v - 1$ and the first time z that u has high fractional degree, the variables $M_{u,v}^L$ are precisely equal to

$$M_{u,v}^L := \bigvee_{w \leq z'} A_{w,u}.$$

Indeed, this is due to u being matched while it has low fractional degree upon the first time that it is picked by a neighbor (or it picks a neighbor) in line 5 or 10, and the corresponding probabilistic test in line 7 or 12 passes. Therefore, by closure of NA under monotone function composition (Proposition 2.6), the variables $\{M_{u,v}^L \mid u \in V_i\}$, which are monotone nondecreasing functions of disjoint subsets of the variables $A_{w,u}$ by bipartiteness, are NA.³ \square

We now turn to upper bounding the probability of the event $M_{u,v}^H$.

Lemma 3.7. *For any edge (u, v) with $u < v$, we have that $\Pr[M_{u,v}^H] \leq 2c$.*

Proof. Recall that by the inductive hypothesis, $\Pr[(u, w) \in \mathcal{M}] = (1/2 + c) \cdot x_{u,w}$. On the other hand, by the fractional matching constraint, we have that $\sum_{w < v} x_{u,w} \leq 1$, and therefore $\Pr[M_{u,v}] \leq 1/2 + c$. On the other hand, if we denote by z_u the first time u has high fractional degree, then either $z_u \geq v$, in which case $\Pr[M_{u,v}^H] = 0$, or

$$\Pr[M_{u,v}^L] \geq \sum_{w < z_u} x_{u,w} \cdot (1/2 + c) \geq \frac{1/2 - c}{1/2 + c} \cdot (1/2 + c) = 1/2 - c,$$

in which case we have

$$\Pr[M_{u,v}^H] = \Pr[M_{u,v}] - \Pr[M_{u,v}^L] \leq 2c. \quad \square$$

We are now ready to prove Lemma 3.4, whereby vertices u, w on the same side of the bipartition have weakly correlated matched statuses, namely $\text{Cov}(F_{u,v}, F_{w,v}) \leq 6c$.

Proof. By definition of covariance, the binary variables $F_{u,v}$ and $F_{w,v}$ satisfy $\text{Cov}(F_{u,v}, F_{w,v}) = \text{Cov}(1 - F_{u,v}, 1 - F_{w,v}) = \text{Cov}(M_{u,v}, M_{w,v})$ (see Proposition B.3). We therefore turn to upper bounding the covariance of the variables $M_{u,v}$ and $M_{w,v}$.

³This is the only place in our analysis where we use bipartiteness.

By the additive law of covariance, the covariance of the variables $M_{u,v} = M_{u,v}^L + M_{u,v}^H$ and $M_{w,v} = M_{w,v}^L + M_{w,v}^H$, denoted by $(\star) = \text{Cov}(M_{u,v}, M_{w,v})$, satisfies

$$\begin{aligned}
(\star) &= \text{Cov}(M_{u,v}^L + M_{u,v}^H, M_{w,v}^L + M_{w,v}^H) \\
&= \text{Cov}(M_{u,v}^L, M_{w,v}^L) + \text{Cov}(M_{u,v}^L, M_{w,v}^H) + \text{Cov}(M_{u,v}^H, M_{w,v}^L) + \text{Cov}(M_{u,v}^H, M_{w,v}^H) \\
&\leq 0 + \Pr[M_{u,v}^L, M_{w,v}^H] + \Pr[M_{u,v}^H, M_{w,v}^L] + \Pr[M_{u,v}^H, M_{w,v}^H] \\
&\leq 0 + \Pr[M_{w,v}^H] + \Pr[M_{u,v}^H] + \Pr[M_{u,v}^H] \\
&\leq 6c.
\end{aligned}$$

Here, the first inequality follows from Corollary 3.6, the second inequality follows from the trivial bound on covariance of Bernoulli variables A and B given by $\text{Cov}(A, B) = \Pr[A, B] - \Pr[A] \cdot \Pr[B] \leq \Pr[A, B] \leq \Pr[A]$, and the final inequality follows from Lemma 3.7. \square

Lemma 3.4 now allows us to argue that if u has high degree upon arrival of v , then $F_{u,v}$ is nearly independent of the event R_v , whereby v is rejected (not matched) after its first pick of u_1 (possibly $u_1 = \text{nil}$). In particular, we have the following.

Lemma 3.8. *Let $u < v$ be a vertex of high fractional degree, $\sum_{w < v} x_{u,w}$, upon arrival of v . Then, for all $w \neq u$ (including possibly $w = \text{nil}$), we have*

$$\Pr[F_{u,v}, R_v, u_1 = w] \geq \Pr[F_{u,v}] \cdot \Pr[R_v, u_1 = w] \cdot \left(1 - \frac{6c}{(1/2 - c)^2}\right).$$

Proof. For $w = \text{nil}$ the claim follows from the event $u_1 = \text{nil}$ implying R_v , and being independent of $F_{u,v}$.

$$\Pr[F_{u,v}, R_v, u_1 = \text{nil}] = \Pr[F_{u,v}] \cdot \Pr[R_v, u_1 = \text{nil}].$$

Next, let $w < v$ be some neighbor of v . If we denote by $q_{w,v} := \min\left(1, \frac{1/2+c}{g(w,v)}\right)$ the probability that w does not reject v if it is picked first and is free, then the probability that $u_1 = w$ and v gets rejected in its first pick is

$$\Pr[R_v, u_1 = w] = x_{w,v} \cdot (1 - q_{w,v} \cdot \Pr[F_{w,v}]) \geq x_{w,v} \cdot (1/2 - c), \quad (7)$$

where the inequality follows from $\Pr[F_{w,v}] = g(w, v)$ by Equation (6), which implies that $q_{w,v} \cdot \Pr[F_{w,v}] \leq 1/2 + c$. Similarly, the probability u is free, $u_1 = w$ and v gets rejected in its first pick is

$$\begin{aligned}
\Pr[F_{u,v}, R_v, u_1 = w] &= x_{w,v} \cdot (\Pr[F_{u,v}] - q_{w,v} \cdot \Pr[F_{w,v}, F_{u,v}]) \\
&\geq x_{w,v} \cdot (\Pr[F_{u,v}] - q_{w,v} \cdot (\Pr[F_{w,v}] \cdot \Pr[F_{u,v}] + 6c)), \\
&\geq x_{w,v} \cdot (\Pr[F_{u,v}] - q_{w,v} \cdot (\Pr[F_{w,v}] \cdot \Pr[F_{u,v}]) - 6c) \\
&\geq x_{w,v} \cdot \Pr[F_{u,v}] \cdot \left(1 - q_{w,v} \cdot \Pr[F_{w,v}] - \frac{6c}{1/2 - c}\right) \\
&\geq \Pr[F_{u,v}] \cdot \Pr[R_v, u_1 = w] \cdot \left(1 - \frac{6c}{(1/2 - c)^2}\right),
\end{aligned}$$

where the first inequality follows from Lemma 3.4, the second inequality follows from the trivial bound $q_{wv} \leq 1$, the third inequality follows from $\Pr[F_{u,v}] = g(u, v) \geq 1/2 - c$ by Equation (5), and the final inequality follows from Equation (7). \square

In what follows we denote by $x_{\text{nil},v} := 1 - \sum_{w < v} x_{w,v}$ the probability with which $u_1 = \text{nil}$. From Lemma 3.8 and Equation (7), as well as $\Pr[R_v, u_1 = \text{nil}] = \Pr[u_1 = \text{nil}] = x_{\text{nil},v}$, we obtain the following lower bound on $\Pr[F_{u,v}, R_v, u_1 = w]$ in terms of $x_{w,v}$.

Corollary 3.9. *For any vertex v and w (possibly $w = \text{nil}$), we have that*

$$\Pr[F_{u,v}, R_v, u_1 = w] \geq \Pr[F_{u,v}] \cdot x_{w,v} \cdot \left(1/2 - c - \frac{6c}{1/2 - c}\right).$$

Finally, we are ready to prove that $p_{u,v}$ is a probability, and in particular $p_{u,v} \leq 1$.

Lemma 3.10. *The parameter $p_{u,v}$ in Line 17 satisfies $p_{u,v} \in [0, 1]$.*

Proof. Non-negativity of $p_{u,v}$ was proven in Corollary 3.3. We turn to proving that $p_{u,v} \leq 1$ suffices to guarantee $\Pr[(u, v) \in \mathcal{M}] \geq (1/2 + c) \cdot x_{u,v}$, from which we obtain that there exists some $p_{u,v} \in [0, 1]$ which results in $\Pr[(u, v) \in \mathcal{M}] = (1/2 + c) \cdot x_{u,v}$.

By Equation (6) we have that $\Pr[F_{u,v}] = g(u, v) \geq 1/2 - c$, and therefore

$$\Pr[(u, v) \in \mathcal{M} \text{ in Line 13}] = x_{u,v} \cdot \min\left(1, \frac{1/2 + c}{g(u, v)}\right) \cdot g(u, v) \geq (1/2 - c) \cdot x_{u,v}. \quad (8)$$

We therefore wish to prove that the probability of (u, v) being matched in Line 18 is at least $2c \cdot x_{u,v}$, for some choice of $p_{u,v} \leq 1$. And indeed,

$$\begin{aligned} \Pr[(u, v) \text{ added to } \mathcal{M} \text{ in Line 18}] &= x_{u,v} \cdot \sum_{w \neq u} \Pr[F_{u,v}, R_v, u_1 = w] \cdot p_{u,v} \\ &\geq x_{u,v} \cdot \Pr[F_{u,v}] \cdot \sum_{w \neq u} x_{w,v} \cdot \left(1/2 - c - \frac{6c}{1/2 - c}\right) \cdot p_{u,v} \\ &\geq x_{u,v} \cdot (1/2 - c) \cdot (1 - 4c) \cdot \left(1/2 - c - \frac{6c}{1/2 - c}\right) \cdot p_{u,v} \\ &\geq 2c \cdot x_{u,v}, \end{aligned}$$

where the first inequality follows from Corollary 3.9 and Equation (3). The second inequality holds due to Equation (5) implying $\Pr[F_{u,v}] \geq 1/2 - c$ and due to vertex u having high degree at time v , and therefore by the fractional matching constraint $x_{u,v} \leq 1 - \frac{1/2 - c}{1/2 + c} = \frac{2c}{1/2 + c} \leq 4c$, and hence $\sum_{w \neq u} x_{w,v} \geq 1 - 4c \geq 0$ (again using Equation (3)). The final inequality holds for $p_{u,v} = 1$ and for our choice of c , by Equation (2).

Consequently, combining the above with Equation (8), we find that setting $p_{u,v} = 1$ results in (u, v) being matched in either Line 13 or Line 18 with probability at least

$$\Pr[(u, v) \in \mathcal{M}] \geq (1/2 + c) \cdot x_{u,v}. \quad (9)$$

As the probability of (u, v) being added to \mathcal{M} in Line 18 is monotone increasing in $p_{u,v}$, we conclude that there exists some $p_{u,v} \in [0, 1]$ for which Equation (9) holds with equality. \square

Conclusion of Algorithm 1's analysis. To conclude, Algorithm 1 is well-defined, and this algorithm outputs a random matching \mathcal{M} which matches each edge e with probability precisely $\Pr[e \in \mathcal{M}] = (1/2 + c) \cdot x_e$. Theorem 3.1 follows.

Remark 2. Computational Aspects: *As described, the only way we are aware of to implement Line 17 exactly (and in particular, computing all $p_{u,v}$ exactly) is using an exponential-time algorithm maintaining the joint distributions as they evolve. However, a simple modification of the algorithm, resulting in a polynomial-time algorithm with a $(1 + o(1))$ additional multiplicative loss in each edge's matching probability, can be readily obtained by approximately estimating the above $p_{u,v}$ up to $(1 \pm o(1))$ multiplicative errors, by standard monte carlo methods. As this results in rather cumbersome descriptions and subsequent calculations, and since running time is not our focus, we do not expand on this.*

4 Putting it all Together

In this section we prove our main result, Theorem 1.2, restated below for ease of reference.

Theorem 1.2. *There exists an online edge coloring algorithm which is $(1.897 + o(1))$ -competitive w.h.p. on general n -node graphs with maximum degree $\Delta = \omega(\log n)$ under vertex arrivals.*

Proof. For a graph of maximum degree at most Δ , assigning x -value $1/\Delta$ to each edge yields a fractional matching. Applying Algorithm 1 to this fractional matching in a bipartite graph under vertex arrivals results in each edge being matched with probability $0.527/\Delta$, by Theorem 3.1. Therefore, by Lemma 2.2, there exists an online edge coloring algorithm whose competitive ratio is $(1/0.527 + o(1)) \approx 1.897 + o(1)$ w.h.p. on bipartite graphs of maximum degree $\Delta = \omega(\log n)$ under (interleaved) vertex arrivals. Finally, Lemma 2.1 together with union bound implies that the same competitive ratio (up to $o(1)$ terms) carries over to general graphs under vertex arrivals. \square

Remark 3. *Our analysis extends to prove the slightly tighter result, whereby there exist constants $c_1, c_2 > 0$ and a $(2 - c_1)$ -competitive online algorithm for n -node graphs of maximum degree at least $c_2 \cdot \log n$ under vertex arrivals. (See Remark 1.) For brevity's sake, we omit the details.*

5 Conclusion

In this work we resolve the longstanding conjecture of Bar-Noy, Motwani and Naor, namely Conjecture 1.1. That is, we show that, while for bounded-degree graphs the greedy algorithm's competitive ratio of 2 is optimal among online algorithms, for high-degree graphs this is not the case.

Some natural questions remain. What is the best achievable competitive ratio? Is a ratio of $1 + o(1)$ possible, as for one-sided arrivals in bipartite graphs and random-order edge arrivals [5, 9]? Can the same be achieved under adversarial *edge* arrivals? Bar-Noy et al. [4] suggested a candidate algorithm for this latter model, but its analysis seems challenging. Finally, does the online rounding Algorithm 1 have more applications beyond edge coloring?

Acknowledgements We thank Janardhan Kulkarni for drawing our attention to [25]. This research is partially supported by NSF award 1812919, ONR award N000141912550, and a gift from Cisco Research.

Appendix

A The Karloff-Shmoys Approach: Online

Here we substantiate our earlier assertion that α -competitive online edge coloring on high-degree graphs is equivalent (up to $o(1)$ terms) to the same task on high-degree *bipartite* graphs. That is, we outline the proof of Lemma 2.1, restated below for ease of reference.

Lemma 2.1. *(Implied by [25]) Given an online edge coloring algorithm which is α -competitive w.h.p. on bipartite graphs of maximum degree $\Delta = \omega(\log n)$ under interleaved vertex arrivals, there exists an online edge coloring algorithm which is $(\alpha + o(1))$ -competitive w.h.p. on general graphs of maximum degree $\Delta = \omega(\log n)$ under vertex arrivals.*

Proof. The general graph edge coloring algorithm relies on the following subroutine for sampling balanced random subgraphs in subgraphs of maximum degree $\Delta' \geq 18 \cdot \sqrt{\Delta \log n}$. (Note that $\Delta \geq 18\sqrt{\Delta \log n}$, by the hypothesis, whereby $\Delta = \omega(\log n)$.) Assign each vertex to a set $V_i \subseteq V$ with $i = 1, 2$ chosen uniformly at random. For any vertex $v \in V$, let $d(v)$ denotes the degree of v in G , and D_v denotes the (random) degree of v in the random bipartite subgraph $H =$

$H(V_1, V_2, E \cap (V_1 \times V_2))$. Then, we have that $\mathbb{E}[D_v] = d(v)/2 \leq \Delta'/2$. By Chernoff's Bound (Proposition B.1), since D_v is the sum of independent Bernoulli(1/2) variables, we have that, for $\epsilon = \sqrt[4]{\log n/\Delta} = o(1)$,

$$\Pr[D_v \geq (\Delta'/2) \cdot (1 + \epsilon)] \leq \exp\left(\frac{-(\Delta'/2) \cdot \epsilon^2}{3}\right) \leq \frac{1}{n^3}, \quad (10)$$

using $\Delta' \geq 18 \cdot \sqrt{\Delta \cdot \log n}$, and consequently $\Delta \cdot \epsilon^2 \geq 18 \log n$. The same high-probability bound holds for $d(v) - D_v$, which is identically distributed to D_v .

To achieve an online edge coloring algorithm for G from the above, we apply the α -competitive edge coloring algorithm to the random bipartite H , and recursively apply the same approach to the random subgraph induced by the edges outside of H , namely $G \setminus H = G[E \setminus (V_1 \times V_2)]$, until H is guaranteed to have degree at most $18 \cdot \sqrt{\Delta \cdot \log n}$ w.h.p. We note that this approach can be applied online, by assigning to each vertex v on arrival a side of each of the recursive random bipartitions. Moreover, the colors of each recursive level number ℓ can be associated with a contiguous set of integers of cardinality $\alpha \cdot \Delta \cdot ((1 + \epsilon)/2)^\ell$, which is the high probability upper bound on the number of colors used in this recursive call. Repeating the above recursively for $t := \log_{2/(1+\epsilon)}(18\sqrt{\Delta \cdot \log n}) \leq \log n$ levels results in a random uncolored subgraph of maximum degree at most $18\sqrt{\Delta \cdot \log n} = o(\Delta)$ w.h.p., which we color greedily.

Taking union bound over the $O(n^2)$ bad events (some vertex degree D_v exceeding $\Delta' \cdot ((1 + \epsilon)/2)$ in a random bipartite subgraph or its complement in a subgraph whose maximum degree is $\Delta' \geq 18\sqrt{\Delta \cdot \log n}$, or any of the bipartite edge coloring algorithms failing to be α competitive on the subgraph it is applied to), we have that w.h.p., the number of colors C used is, as desired, at most

$$\begin{aligned} C &\leq \alpha \cdot \Delta \cdot \frac{1 + \epsilon}{2} + \alpha \cdot \Delta \cdot \left(\frac{1 + \epsilon}{2}\right)^2 + \cdots + \alpha \cdot \Delta \cdot \left(\frac{1 + \epsilon}{2}\right)^t + 36 \cdot \sqrt{\Delta \cdot \log n} \\ &\leq \alpha \cdot \sum_{i \geq 1} \Delta \cdot \left(\frac{1 + \epsilon}{2}\right)^i + 36 \cdot \sqrt{\Delta \cdot \log n} \\ &= \alpha \cdot \Delta \cdot \frac{1 + \epsilon}{1 - \epsilon} + o(\Delta) \\ &= (\alpha + o(1)) \cdot \Delta. \quad \square \end{aligned}$$

Remark 4. *As stated in the introduction, we note that the above reduction from general to bipartite graphs results in bipartite graphs with interleaved vertex arrivals.*

B Edge Coloring from Random Matchings

In this section, we show how to reduce edge coloring in (bipartite) graphs under vertex arrivals to computing a random matching which matches each edge with probability $\Omega(1/\Delta)$.

Lemma 2.2. *(Implied by [9]) Let \mathcal{A} be an online matching algorithm which on any (bipartite) graph of maximum degree $\Delta \leq \Delta'$ under vertex arrivals, matches each edge with probability at least $1/(\alpha\Delta')$. Then, there exists an online edge coloring algorithm \mathcal{A}' which is $(\alpha + o(1))$ -competitive w.h.p. for (bipartite) graphs of maximum degree $\Delta = \omega(\log n)$ under vertex arrivals.*

Proof. If $\alpha > 2$, then the claim follows trivially from the greedy algorithm's 2-competitiveness. We therefore assume $\alpha \leq 2$. We give a subroutine which decreases the uncolored degree of a subgraph of maximum degree $\Delta' \geq 48 \cdot \sqrt[4]{\Delta^3 \log n}$ at a rate of one per $\alpha + o(1)$ colors w.h.p. (Note that $\Delta \geq 48 \sqrt[4]{\Delta^3 \log n}$, by the hypothesis, whereby $\Delta = \omega(\log n)$.)

Our subroutine is as follows. Let $L := 12\sqrt{\Delta \log n}$ and $\epsilon := \sqrt[4]{(\log n)/\Delta} (= o(1) \leq 1/2)$. We note that by our choice of L and ϵ and our lower bound on Δ' , we have that

$$4L/\Delta' \leq 48\sqrt{\Delta \log n}/48\sqrt[4]{\Delta^3 \log n} = \sqrt[4]{(\log n)/\Delta} = \epsilon. \quad (11)$$

For $i = 1, \dots, \lceil \alpha \cdot L \rceil$, we run Algorithm \mathcal{A} , which matches each edge with probability at least $(1/\alpha)/\Delta'$, and color all previously-uncolored matched edges in this run of \mathcal{A} using a new (common) color. Fix a vertex v whose degree in the subgraph is at least $d(v) \geq \Delta' - \lceil \alpha \cdot L \rceil$ and let X_1, \dots, X_L be indicators of v having an edge colored during application $i = 1, \dots, \lceil \alpha \cdot L \rceil$ of Algorithm \mathcal{A} . Since vertex v can have at most $\lceil \alpha \cdot L \rceil \leq 2 \cdot L$ edges colored during these L applications of Algorithm \mathcal{A} , we find that the number of uncolored edges of v at any point during this subroutine is at least $\Delta' - 2\lceil \alpha \cdot L \rceil \geq \Delta' - 4L$, independently of previous random choices. On the other hand, since each uncolored edge is matched (and hence colored) with probability at least $(1/\alpha)/\Delta'$, we have that for any history \mathcal{H} of random choices in applications $1, 2, \dots, i-1$ of \mathcal{A} , application i of \mathcal{A} results in one of the (at least) $\Delta' - 4L$ uncolored edges of v being colored with probability at least

$$\Pr[X_i \mid \mathcal{H}] \geq (1/\alpha) \cdot (\Delta' - 4L)/\Delta' = (1/\alpha) \cdot (1 - 4L/\Delta') \geq (1/\alpha) \cdot (1 - \epsilon), \quad (12)$$

where the last inequality relied on Equation (11). Combining Equation (12) with standard coupling arguments (Proposition B.2) together with a Chernoff Bound (Proposition B.1), we find that the number of colored edges of v , denoted by $X := \sum_i X_i$ satisfies

$$\Pr[X \leq L \cdot (1 - \epsilon)^2] \leq \exp\left(\frac{-L \cdot (1 - \epsilon) \cdot \epsilon^2}{2}\right) \leq \exp\left(\frac{-L \cdot \epsilon^2}{4}\right) = \frac{1}{n^3},$$

where the second inequality follows from $\epsilon \leq 1/2$ and the equality follows from choice of L and ϵ . Union bounding over the n vertices, we obtain the following high probability bound on the maximum degree of the uncolored subgraph H after the $\lceil \alpha \cdot L \rceil$ applications of \mathcal{A} :

$$\Pr[\Delta(H) \geq \Delta' - L \cdot (1 - \epsilon)^2] \leq \frac{1}{n^2}. \quad (13)$$

We now describe how to make use of this subroutine. For $r = 1, \dots, \Delta/L$ phases, let $\Delta_i := \Delta - (i-1) \cdot L \cdot (1 - \epsilon)^2$. If $\Delta_i < 48 \sqrt[4]{\Delta^3 \log n}$, apply the greedy coloring. Otherwise, apply the above subroutine with $\Delta' = \Delta_i$. A simple inductive argument together with union bound, relying on Equation (13), shows that for $i = 1, 2, \dots, \Delta/L (\leq n)$, the uncolored subgraph after the first $i-1$ phases has maximum degree at most $\Delta' \leq \Delta_i$ w.h.p., or alternatively it has maximum degree at most $\Delta' \leq 48 \cdot \sqrt[4]{\Delta^3 \log n} = o(\Delta)$. Moreover, each of these Δ/L phases requires at most $\lceil \alpha \cdot L \rceil \leq \alpha \cdot L + 1$ colors, by definition, and therefore these Δ/L phases require at most $\alpha \cdot \Delta + \Delta/L = (\alpha + o(1)) \cdot \Delta$ colors in total. Finally, after these phases we are guaranteed that the maximum degree of the uncolored subgraph is at most $\min\{48 \cdot \sqrt[4]{\Delta^3 \log n}, \Delta - (\Delta/L) \cdot L \cdot (1 - \epsilon)^2\} = o(\Delta)$. Applying the greedy algorithm to this uncolored subgraph after the Δ/L phases thus requires a further $2 \cdot o(\Delta) = o(\Delta)$ colors. This results in a proper edge coloring using $(\alpha + o(1)) \cdot \Delta$ colors w.h.p.

Finally, we note that the above algorithm can be implemented online under vertex arrivals, since \mathcal{A} works under vertex arrivals. In particular, when a vertex arrives, we perform the next steps of the different copies of Algorithm \mathcal{A} (with the different settings of Δ_i) on the uncolored subgraphs obtained from each phase, simulating the arrival of a vertex in each such uncolored subgraph. Combined with the above, this yields the desired result: an edge coloring algorithm which is $(\alpha + o(1))$ -competitive on general n -node graphs of maximum degree $\Delta = \omega(\log n)$ under vertex arrivals. \square

Remark 5. Lemma 2.2 naturally extends to edge arrivals. Unfortunately, no algorithm matching each edge with probability $(1/\alpha)/\Delta$ subject to edge arrivals is currently known for any constant $\alpha < 2$.

Remark 6. The approach of Lemma 2.2 only requires matching algorithms which match each edge with probability $(1/\alpha)/\Delta$ for subgraphs of the input graph. Consequently, improved matching algorithms, with smaller $\alpha \geq 1$, for any downward-closed family of graphs \mathcal{F} imply a similar improved $(\alpha + o(1))$ -competitive edge coloring algorithm for the same family.

B.1 Probability Basics

Here we include, for completeness, a number of basic probabilistic results used in this paper.

Proposition B.1 (Chernoff Bound). *Let $X = \sum_i X_i$ be the sum of independent Bernoulli random variables $X_i \sim \text{Bernoulli}(p_i)$, with expectation $\mu := \mathbb{E}[X] = \sum_i p_i$. Then, for any $\epsilon \in (0, 1)$, and $\kappa \geq \mu$,*

$$\Pr[X \geq \kappa \cdot (1 + \epsilon)] \leq \exp\left(\frac{-\kappa \cdot \epsilon^2}{3}\right).$$

$$\Pr[X \leq \mu \cdot (1 - \epsilon)] \leq \exp\left(\frac{-\mu \cdot \epsilon^2}{2}\right).$$

Proposition B.2 (Coupling). *Let X_1, \dots, X_m be binary random variables such that for all i and $\vec{x} \in \{0, 1\}^{i-1}$,*

$$\Pr\left[X_i = 1 \mid \bigwedge_{\ell \in [i-1]} (X_\ell = x_\ell)\right] \geq p_i.$$

If $\{Y_i \sim \text{Bernoulli}(p_i)\}_i$ are independent random variables, then for any $k \in \mathbb{R}$,

$$\Pr\left[\sum_i X_i \leq k\right] \leq \Pr\left[\sum_i Y_i \leq k\right].$$

Proposition B.3. *Let A and B be Bernoulli random variables. Then*

$$\text{Cov}(A, B) = \text{Cov}(1 - A, 1 - B).$$

References

- [1] Gagan Aggarwal, Rajeev Motwani, Devavrat Shah, and An Zhu. Switch scheduling via randomized edge coloring. In *Proceedings of the 44th Symposium on Foundations of Computer Science (FOCS)*, pages 502–512, 2003.
- [2] Itai Ashlagi, Maximilien Burq, Chinmoy Dutta, Patrick Jaillet, Amin Saberi, and Chris Sholley. Edge weighted online windowed matching. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, pages 729–742, 2019.
- [3] Bahman Bahmani, Aranyak Mehta, and Rajeev Motwani. Online graph edge-coloring in the random-order arrival model. *Theory of Computing*, 8(1):567–595, 2012.
- [4] Amotz Bar-Noy, Rajeev Motwani, and Joseph Naor. The greedy algorithm is optimal for on-line edge coloring. *Information Processing Letters (IPL)*, 44(5):251–253, 1992.
- [5] Sayan Bhattacharya, Fabrizio Grandoni, and David Wajc. Online edge coloring algorithms via the nibble method. In *Proceedings of the 32nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2830–2842, 2021.
- [6] Niv Buchbinder, Kamal Jain, and Joseph (Seffi) Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. In *Proceedings of the 15th Annual European Symposium on Algorithms (ESA)*, pages 253–264. 2007.
- [7] Yi-Jun Chang, Qizheng He, Wenzheng Li, Seth Pettie, and Jara Uitto. The complexity of distributed edge coloring with small palettes. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2633–2652, 2018.

- [8] Moses Charikar and Paul Liu. Improved algorithms for edge colouring in the w-streaming model. In *Proceedings of the 4th Symposium on Simplicity in Algorithms (SOSA)*, pages 181–183, 2021.
- [9] Ilan Reuven Cohen, Binghui Peng, and David Wajc. Tight bounds for online edge coloring. In *Proceedings of the 60th Symposium on Foundations of Computer Science (FOCS)*, pages 1–25, 2019.
- [10] Ilan Reuven Cohen and David Wajc. Randomized online matching in regular graphs. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 960–979, 2018.
- [11] Richard Cole, Kirstin Ost, and Stefan Schirra. Edge-coloring bipartite multigraphs in $O(E \log D)$ time. *Combinatorica*, 21(1):5–12, 2001.
- [12] Nikhil R Devanur, Kamal Jain, and Robert D Kleinberg. Randomized primal-dual analysis of ranking for online bipartite matching. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 101–107, 2013.
- [13] Ran Duan, Haoqing He, and Tianyi Zhang. Dynamic edge coloring with improved approximation. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1937–1945, 2019.
- [14] Devdatt Dubhashi and Desh Ranjan. Balls and bins: A study in negative dependence. *BRICS Report Series*, 3(25), 1996.
- [15] Tomer Ezra, Michal Feldman, Nick Gravin, and Zhihao Gavin Tang. Online stochastic maximum weight matching: prophet inequality for vertex and edge arrival models. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 769–787, 2020.
- [16] Matthew Fahrbach, Zhiyi Huang, Runzhou Tao, and Morteza Zadimoghaddam. Edge-weighted online bipartite matching. In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*, pages 412–423, 2020.
- [17] Jon Feldman, Nitish Korula, Vahab Mirrokni, S Muthukrishnan, and Martin Pál. Online ad assignment with free disposal. In *Proceedings of the 5th Conference on Web and Internet Economics (WINE)*, pages 374–385, 2009.
- [18] Yiding Feng, Rad Niazadeh, and Amin Saberi. Two-stage stochastic matching with application to ride hailing. In *Proceedings of the 32nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2862–2877, 2021.
- [19] Buddhima Gamlath, Michael Kapralov, Andreas Maggiori, Ola Svensson, and David Wajc. Online matching with general arrivals. In *Proceedings of the 60th Symposium on Foundations of Computer Science (FOCS)*, pages 26–37, 2019.
- [20] Ian Holyer. The np-completeness of edge-coloring. *SIAM Journal on Computing (SICOMP)*, 10(4):718–720, 1981.
- [21] Zhiyi Huang, Ning Kang, Zhihao Gavin Tang, Xiaowei Wu, Yuhao Zhang, and Xue Zhu. Fully online matching. *Journal of the ACM (JACM)*, 67(3):1–25, 2020.
- [22] Zhiyi Huang, Zhihao Gavin Tang, Xiaowei Wu, and Yuhao Zhang. Fully online matching ii: Beating ranking and water-filling. In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*, pages 1380–1391, 2020.
- [23] Kumar Joag-Dev and Frank Proschan. Negative association of random variables with applications. *The Annals of Statistics*, pages 286–295, 1983.

- [24] Bala Kalyanasundaram and Kirk R Pruhs. An optimal deterministic algorithm for online b -matching. *Theoretical Computer Science (TCS)*, 233(1):319–325, 2000.
- [25] Howard J. Karloff and David B. Shmoys. Efficient parallel algorithms for edge coloring problems. *J. Algorithms*, 8(1):39–52, 1987.
- [26] Richard M Karp, Umesh V Vazirani, and Vijay V Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 352–358, 1990.
- [27] Alam Khursheed and KM Lai Saxena. Positive dependence in multivariate distributions. *Communications in Statistics - Theory and Methods*, 10(12):1183–1196, 1981.
- [28] Dénes König. Über graphen und ihre anwendung auf determinantentheorie und mengenlehre. *Mathematische Annalen*, 77(4):453–465, 1916.
- [29] Aranyak Mehta. Online matching and ad allocation. *Foundations and Trends® in Theoretical Computer Science*, 8(4):265–368, 2013.
- [30] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *Journal of the ACM (JACM)*, 54(5):22, 2007.
- [31] Rajeev Motwani, Joseph (Seffi) Naor, and Moni Naor. The probabilistic method yields deterministic parallel algorithms. *Journal of Computer and System Sciences*, 49(3):478–516, 1994.
- [32] Joseph Seffi Naor and David Wajc. Near-optimum online ad allocation for targeted advertising. *ACM Transactions on Economics and Computation (TEAC)*, 6(3-4):16:1–16:20, 2018.
- [33] Chrisots Papadimitriou, Tristan Pollner, Amin Saberi, and David Wajc. Online stochastic max-weight bipartite matching: Beyond prophet inequalities. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 2021. To appear.
- [34] Hsin-Hao Su and Hoa T. Vu. Towards the locality of vizing’s theorem. In *Proceedings of the 51st Annual ACM Symposium on Theory of Computing (STOC)*, pages 355–364, 2019.
- [35] Vadim G Vizing. On an estimate of the chromatic class of a p -graph. *Diskret analiz*, 3:25–30, 1964.
- [36] David Wajc. Rounding dynamic matchings against an adaptive adversary. In *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 194–207, 2020.
- [37] Yajun Wang and Sam Chiu-wai Wong. 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)*, pages 1070–1081, 2015.