COMBINATORICA
Bolyai Society – Springer-Verlag

Combinatorica 1–32

PERFECT MATCHINGS IN $\tilde{O}(n^{1.5})$ TIME IN REGULAR BIPARTITE GRAPHS

ASHISH GOEL*, MICHAEL KAPRALOV†, SANJEEV KHANNA‡

Received 23 July, 2009

We consider the well-studied problem of finding a perfect matching in d-regular bipartite graphs with $2n$ vertices and $m=nd$ edges. While the best-known algorithm for general bipartite graphs (due to Hopcroft and Karp) takes $O(m\sqrt{n})$ time, in regular bipartite graphs, a perfect matching is known to be computable in $O(m)$ time. Very recently, the $O(m)$ bound was improved to $O(\min\{m, \frac{n^{2.5}\ln n}{d}\})$ expected time, an expression that is bounded by $\tilde{O}(n^{1.75})$. In this paper, we further improve this result by giving an $O(\min\{m, \frac{n^2\ln^3 n}{d}\})$ expected time algorithm for finding a perfect matching in regular bipartite graphs; as a function of n alone, the algorithm takes expected time $O((n\ln n)^{1.5})$.

To obtain this result, we design and analyze a two-stage sampling scheme that reduces the problem of finding a perfect matching in a regular bipartite graph to the same problem on a subsampled bipartite graph with $O(n\ln n)$ edges. The first-stage is a sub-linear time uniform sampling that reduces the size of the input graph while maintaining certain structural properties of the original graph. The second-stage is a non-uniform sampling that takes linear-time (on the reduced graph) and outputs a graph with $O(n\ln n)$ edges, while preserving a matching with high probability. This matching is then recovered using the Hopcroft-Karp algorithm. While the standard analysis of Hopcroft-Karp also gives us an $\tilde{O}(n^{1.5})$ running time, we present a tighter analysis for our special case that results in the stronger $\tilde{O}(\min\{m, \frac{n^2}{d}\})$ time mentioned earlier.

Our proof of correctness of this sampling scheme uses a new correspondence theorem between cuts and Hall's theorem "witnesses" for a perfect matching in a bipartite graph that we prove. We believe this theorem may be of independent interest; as another example application, we show that a perfect matching in the support of an $n\times n$ doubly stochastic matrix with m non-zero entries can be found in expected time $\tilde{O}(m+n^{1.5})$.

Mathematics Subject Classification (2010): .... Fill in, please

* Research supported by NSF ITR grant 0428868, NSF CAREER award 0339262, and a grant from the Stanford-KAUST alliance for academic excellence.

† Research supported by a Stanford Graduate Fellowship.

‡ Supported in part by a Guggenheim Fellowship, an IBM Faculty Award, and by NSF Award CCF-0635084.

1. Introduction

A bipartite graph $G=(P,Q,E)$ with vertex set $P\cup Q$ and edge set $E\subseteq P\times Q$ is said to be regular if every vertex has the same degree d. We use $m=nd$ to denote the number of edges in G and n to represent the number of vertices in P (as a consequence of regularity, P and Q have the same size). Regular bipartite graphs are a fundamental combinatorial object, and arise, among other things, in expander constructions, scheduling, routing in switch fabrics, and task-assignment [15,1,6].

A regular bipartite graph of degree d can be decomposed into exactly d perfect matchings, a fact that is an easy consequence of Hall's theorem [4], and is closely related to the Birkhoff-von Neumann decomposition of a doubly stochastic matrix [3,17]. Finding a matching in a regular bipartite graph is a well-studied problem, starting with the algorithm of König in 1916 [13], which is now known to run in time $O(mn)$. The well-known bipartite matching algorithm of Hopcroft and Karp [9] can be used to obtain a running time of $O(m\sqrt{n})$. In graphs where d is a power of 2, the following elegant idea, due to Gabow and Kariv [7], leads to an algorithm with $O(m)$ running time. First, compute an Euler tour of the graph (in time $O(m)$) and then follow this tour in an arbitrary direction. Exactly half the edges will go from left to right; these form a regular bipartite graph of degree $d/2$. The total running time $T(m)$ thus follows the recurrence $T(m)=O(m)+T(m/2)$ which yields $T(m)=O(m)$. Extending this idea to the general case proved quite hard, and after a series of improvements (eg. by Cole and Hopcroft [5], and then by Schrijver [16] to $O(md)$), Cole, Ost, and Schirra [6] gave an $O(m)$ algorithm for the case of general d. Their main interest was in edge coloring of general bipartite graphs, where finding perfect matchings in regular bipartite graphs is an important subroutine. Very recently, Goel, Kapralov, and Khanna [8], gave a sampling-based algorithm that computes a perfect matching in d-regular bipartite graphs in $O(\min\{m,\frac{n^{2.5}\ln n}{d}\})$ expected time, an expression that is bounded by $\tilde{O}(n^{1.75})$. The algorithm of [8] uses uniform sampling to reduce the number of edges in the input graph while preserving a perfect matching, and then runs the Hopcroft-Karp algorithm on the sampled graph.

Our Results and Techniques: We present a significantly faster algorithm for finding perfect matchings in regular bipartite graphs.

Theorem 1.1. *There is an* $O\left(\min\{m,\frac{n^2\ln^3 n}{d}\}\right)$ *expected time algorithm to find a perfect matching in a* d*-regular bipartite graph* G.

As a function of n alone, the running time stated above is $O((n\ln n)^{1.5})$. Since the $O(m)$ running time is guaranteed by the algorithm of Cole, Ost, and Schirra, we are only concerned with the case where d is $\Omega(\sqrt{n}\ln n)$. For this regime, our algorithm reduces the perfect matching problem on a regular bipartite graph G to the same problem on a (not necessarily regular) sparse bipartite graph H with $O(n\ln n)$ edges. This reduction takes time $O(\frac{n^2\ln^3 n}{d})$. We then use the Hopcroft-Karp algorithm on H to recover a perfect matching. A black-box use of the analysis of the Hopcroft-Karp algorithm would suggest a running time of $O(\frac{n^2\ln^3 n}{d}+n^{1.5}\ln n)$. However, we show that the final sampled graph has some special structure that guarantees that the Hopcroft-Karp algorithm would complete in time $O(\frac{n^2\ln^2 n}{d})$ whp.

For every pair $A \subseteq P, B \subseteq Q$, we define a *witness set* $W(A,B)$ to be the set of all edges going from A to $Q \setminus B$. Of particular interest are what we call *Hall witness sets*, which correspond to $|A| > |B|$; the well-known Hall's theorem [4] says that a bipartite graph $H(P,Q,E_H)$ contains a perfect matching iff E_H includes an edge from each Hall witness set. Thus any approach that reduces the size of the input bipartite graph by sampling must ensure that some edge from every Hall witness set is included in the sampled graph; otherwise the sampled graph no longer contains a perfect matching. Goel, Kapralov, and Khanna [8] showed that no *uniform sampling* scheme on a d-regular bipartite graph can reduce the number of edges to $o(\frac{n^2}{d\ln n})$ while preserving a perfect matching, and hence their $\tilde{O}(n^{1.75})$-time algorithm is the best possible running time achievable via uniform sampling followed by a black-box invocation of the Hopcroft-Karp analysis.

In order to get past this barrier, we use here a two-stage sampling process. The first stage is a uniform sampling (along the lines of [8]) which generates a reduced-size graph $G'=(P,Q,E')$ that preserves not only a perfect matching but also a key relationship between the sizes of "relevant" witness sets and cuts in the graph G. The second stage is to run the non-uniform Benczúr-Karger sampling scheme [2] on G' to generate a graph G'' with $\tilde{O}(n)$ edges while preserving a perfect matching whp. Since this step requires $\tilde{\Omega}(|E'|)$ time, we crucially rely on the fact that G' does not contain too many edges.

While our algorithm is easy to state and understand, the proof of correctness is quite involved. The Benczúr-Karger sampling was developed to generate, for any graph, a weighted subgraph with $\tilde{O}(n)$ edges that approximately preserves the size of all cuts in the original graph. The central idea underlying our result is to show that there exists a collection of *core* witness sets that can be identified in an almost one-one manner with cuts in the graph such that the probability mass of edges in each witness set is comparable to the probability mass of the edges in the cut identified with it. Further, every witness set in the graph has a "representative" in this collection of core witness sets. Informally, this allows us to employ cut-preserving sampling schemes such as Benczúr-Karger as "witness-preserving" schemes.

We note here that the natural mapping which assigns the witness set of a pair (A, B) to the cut edges associated with this pair can map arbitrarily many witness sets to the same cut and is not useful for our purposes. One of our contributions is an uncrossing theorem for witness sets, that we refer to as the *proportionate uncrossing theorem.* Informally speaking, it says that given any collection of witness sets $\mathcal{R}$ such that the probability mass of each witness set is comparable to that of its associated cut, there exists another collection $\mathcal{T}$ of witness sets such that (i) the natural mapping to cuts as defined above is *half-injective* for $\mathcal{T}$, that is, at most two witness sets in $\mathcal{T}$ map to any given cut, (ii) the probability mass of each witness set is comparable to the probability mass of its associated cut, and (iii) any subset of edges that hits every witness set in $\mathcal{T}$ also hits every witness set in $\mathcal{R}$. The collection $\mathcal{T}$ is referred to as a proportional uncrossing of $\mathcal{R}$. As shown in Figure 1(a), we can not achieve an injective mapping, and hence the half-injectivity is unavoidable.

We believe the half-injective correspondence between witness sets and cuts, as facilitated by the proportionate uncrossing theorem, is of independent interest, and will perhaps have other applications in this space of problems. We also emphasize here that the uncrossing theorem holds for all bipartite graphs, and not only regular bipartite graphs. Indeed, the graph G' on which we invoke this theorem does not inherit the regularity property of the original graph G. As another illustrative example, consider the celebrated Birkhoff-von Neumann theorem [4,17] which says that every doubly stochastic matrix can be expressed as a convex combination of permutation matrices (i.e., perfect matchings). In some applications, it is of interest to do an iterative decomposition whereby a single matching is recovered in each iteration. The best-known bound for this problem, to our knowledge, is an $O(mb)$ time algorithm that follows from the work of Gabow and Kariv [7]; here b denotes the maximum number of bits needed to express any entry in M. The following theorem is an easy consequence of our proportionate uncrossing result.

Theorem 1.2. *Given an $n \times n$ doubly-stochastic matrix M with m non-zero entries, one can find a perfect matching in the support of M in $\tilde{O}(m + n^{1.5})$ expected time.*

The proof of this theorem and a discussion of known results about this problem are given in section 6. Though this result itself represents only a modest improvement over the earlier $O(mb)$ running time, it is an instructive illustration of the utility of the proportionate uncrossing theorem.

It is worth noting that while the analysis of Goel, Kapralov, and Khanna was along broadly similar lines (sample edges from the original graph, followed by running the Hopcroft-Karp algorithm), the proportionate uncrossing theorem developed in this paper requires significant new ideas and is crucial to incorporating the non-uniform sampling stage into our algorithm. Further, the running time of the Hopcroft-Karp algorithm is easily seen to be $\Omega(m\sqrt{n})$ even for the 2-regular graph consisting of $\Theta(\sqrt{n})$ disjoint cycles of lengths $2, 4, \ldots, \sqrt{n}$ respectively; the stronger analysis for our special case requires both our uncrossing theorem as well as a stronger decomposition[1]. As a step in this analysis, we prove the independently interesting fact that after sampling edges from a d-regular bipartite graph with rate $\frac{c\ln n}{d}$, for some suitable constant c, we obtain a graph that has a matching of size $n - O(n/d)$ whp and such a matching can be found in $O(n/d)$ augmenting phases of the Hopcroft-Karp algorithm whp.

Organization: Section 2 reviews and presents some useful corollaries of relevant earlier work. In section 3, we establish the proportionate uncrossing theorem. In section 4, we present and analyze our two-stage sampling scheme, and section 5 outlines the stronger analysis of the Hopcroft-Karp algorithm for our special case. Section 6 contains the proof of Theorem 1.2 and a discussion of known results on finding perfect matchings in the support of double stochastic matrices.

2. Preliminaries

In this section, we adapt and present recent results of Goel, Kapralov, and Khanna [8] as well as the Benczúr-Karger sampling theorem [2] for our purposes, and also prove a simple technical lemma for later use.

2.1. Bipartite decompositions and relevant witness pairs

Let $G = (P, Q, E)$ be a regular bipartite graph, with vertex set $P \cup Q$ and edge set $E \subseteq P \times Q$. Consider any partition of P into k sets $P_1, P_2, \ldots, P_k$, and a partition of Q into $Q_1, Q_2, \ldots, Q_k$. Let G_i denote the (not necessarily regular) bipartite graph (P_i, Q_i, E_i) where $E_i = E \cap (P_i \times Q_i)$. We will call this a "decomposition" of G.

Given $A \subseteq P$ and $B \subseteq Q$, define the witness set corresponding to the pair (A, B), denoted $W(A, B)$, as the set of all edges between A and $Q \setminus B$, and define the cut $C(A, B)$ as the set of all edges between $A \cup B$ and $(P \setminus A) \cup (Q \setminus B)$. The rest of the definitions in this section are with respect to some arbitrary but fixed decomposition of G.

[1] It is known that the Hopcroft-Karp algorithm terminates quickly on bipartite expanders [14], but those techniques don't help in our setting since we start with an arbitrary regular bipartite graph.

Definition 2.1. An edge $(u,v)\in E$ is relevant if $(u,v)\in E_i$ for some i.

Definition 2.2. Let E_R be the set of all relevant edges. A pair (A,B) is said to be relevant if

1. $A\subseteq P_i$ and $B\subseteq Q_i$ for some i,
2. $|A|>|B|$, and
3. There does not exist another $A'\in P_i$, $B'\in Q_i$, such that $A'\subset A$, $|A'|>|B'|$, and $W(A',B')\cap E_R\subseteq W(A,B)\cap E_R$.

Informally, a relevant pair is one which is contained completely within a single piece in the decomposition, and is "minimal" with respect to that piece. The following lemma is implicit in [8] and is proved in appendix A for completeness.

Lemma 2.3. *Let $\mathcal{R}$ denote all relevant pairs (A,B) with respect to a decomposition of $G(P,Q,E)$, and let E_R denote all relevant edges. Consider any graph $G^*=(P,Q,E^*)$. If for all $(A,B)\in\mathcal{R}$, we have $W(A,B)\cap E^*\cap E_R\neq\emptyset$, then G^* has a perfect matching.*

2.2. A Corollary of Benczúr-Karger sampling scheme

The Benczúr-Karger sampling theorem [2] shows that for any graph, a relatively small *non-uniform* edge sampling rate suffices to ensure that every cut in the graph is hit by the sampled edges (i.e. it has a non-empty intersection) with high probability. The sampling rate used for each edge e inversely depends on its strength, as defined below.

Definition 2.4. [2] A k-strong component of a graph H is a maximal vertex-induced subgraph of H with edge-connectivity k. The strength of an edge e in a graph H is the maximum value of k such that a k-strong component contains e.

Definition 2.5. Given a graph $H=(V,E)$, let $H_{[j]}=(V,E_{[j]})$ denote the subgraph of H restricted to edges of strength j or higher, where j is some integer in $\{1,2,\dots,|V|\}$.

It is easy to see that whenever a cut in a graph $H(V,E)$ contains an edge of strength k, then the cut must contain at least k edges. Furthermore, for any $1<j\leq|V|$, each connected component of graph $H_{[j]}$ is contained inside some connected component of $H_{[j-1]}$. The Benczúr-Karger theorem utilizes these properties to show that it suffices to sample each edge e with probability $\Theta(\min\{1,\ln n/s_e\})$.

We now extend this sampling result to any collection of edge-sets for which there exists an injection (one-one mapping) to cuts of comparable inverse strengths. The statement of our Theorem 2.6 closely mirrors the Benczúr-Karger sampling theorem, and the proof is also along the same general lines. However, the proof does not follow from the Benczúr-Karger sampling theorem in a black-box fashion, so a proof is provided in appendix B.

Theorem 2.6. *Let $H(V,E)$ be any graph on n vertices, and let $\mathcal{C}$ denote the set of all possible edge cuts in H, and $\gamma \in (0,1]$ be a constant. Let H' be a subgraph of H obtained by sampling each edge e in H with probability $p_e = \min\left\{1, \frac{c \ln n}{\gamma s_e}\right\}$, where s_e denotes the strength of edge e, and c is a suitably large constant. Further, let $\mathcal{X}$ be a collection of subsets of edges, and let f be a one-one (not necessarily onto) mapping from $\mathcal{X}$ to $\mathcal{C}$ satisfying $\sum_{e\in X} 1/s_e > \gamma \sum_{e\in f(X)} 1/s_e$ for all $X \in \mathcal{X}$. Then*

$$\sum_{X\in\mathcal{X}} \Pr[\text{No edge in } X \text{ is chosen in } H'] \le \frac{1}{n^2}.$$

The result below from [2] bounds the number of edges chosen by the sampling in Theorem 2.6.

Theorem 2.7. *Let $H(V,E)$ be any graph on n vertices, and let H' be a subgraph of H obtained by sampling each edge e in H with probability $p_e = \min\left\{1, \frac{c \ln n}{s_e}\right\}$, where s_e denotes the strength of edge e, and c is any constant. Then with probability at least $1 - \frac{1}{n^2}$, the graph H' contains at most $c'n \ln n$ edges, where c' is another suitably large constant.*

We conclude with a simple property of integer multisets that we will use later. A similar statement was used in [11] (lemma 4.5). A proof is provided in appendix C for completeness.

Lemma 2.8. *Let S_1 and S_2 be two arbitrary multisets of positive integers such that $|S_1| > \gamma |S_2|$ for some $\gamma > 0$. Then there exists an integer j such that*

$$\sum_{i \ge j \text{ and } i \in S_1} \frac{1}{i} > \gamma \left(\sum_{i \ge j \text{ and } i \in S_2} \frac{1}{i} \right).$$

3. Proportionate uncrossing of witness sets

Consider a bipartite graph $G = (P, Q, E)$, with a non-negative weight function t defined on the edges. Assume further that we are given a set of "relevant edges" $E_R \subseteq E$. We can extend the definition of t to sets of edges, so that $t(S) = \sum_{e\in S} t(e)$, where $S \subseteq E$.

Definition 3.1. For any $\gamma > 0$ and $A \subseteq P, B \subseteq Q$, the pair (A, B) is said to be γ-thick with respect to (G, t, E_R) if $t(W(A, B) \cap E_R) > \gamma t(C(A, B))$, *i.e.*, the total weight of the relevant edges in $W(A, B)$ is *strictly* more than γ times the total weight of $C(A, B)$. A set of pairs $\mathcal{R} = \{(A_1, B_1), (A_2, B_2), \ldots, (A_K, B_K)\}$ where each $A_i \subseteq P$ and each $B_i \subseteq Q$ is said to be a γ-thick collection with respect to (G, t, E_R) if every pair $(A_i, B_i) \in \mathcal{R}$ is γ-thick.

The quantities G, t, and E_R will be fixed for this section, and for brevity, we will omit the phrase "with respect to (G, t, E_R)" in the rest of this section.

Before defining proportionate uncrossings of witness sets, we will informally point out the motivation for doing so. If a pair (A, B) is γ-thick for some constant γ, and if we know that a sampling process where edge e is chosen with probability t chooses some edge from $C(A, B)$ with high probability, then increasing the sampling probability by a factor of $1/\gamma$ should result in some relevant edge from $W(A, B)$ being chosen with high probability as well, a fact that would be very useful in the rest of this paper. The sampling sub-routines that we employ in the rest of this paper are analyzed by using union-bound over all cuts, and in order to apply the same union bound, it would be useful if each witness set were to correspond to a unique cut. However, in figure 1(a), we show two pairs (A, B) and (X, Y) which are both $(1/3)$-thick with respect to the uniform weight function $t \equiv 1$ but correspond to the same cut; we call this a "crossing" of the pairs (A, B) and (X, Y), drawing intuition from the figure. In general, we can have many witness sets that map to the same cut. We would like to "uncross" these witness sets by finding subsets of each witness set that map to unique cuts, but there is no way to uncross figure 1(a) in this fashion. Fortunately, and somewhat surprisingly, this is the worst case: any collection of γ-thick pairs can be uncrossed into another collection such that all the pairs in the new collection are also γ-thick (hence the term proportionate uncrossing), every original witness set has a representative in this new collection, and no more than two new pairs have the same cut. Figure 1(b) shows two $\frac{1}{3}$-thick pairs that can be uncrossed using a single $\frac{1}{3}$-thick representative, $(A \cap X, B \cap Y)$. We will spend the rest of this section formalizing the notion of proportionate uncrossings and proving their existence. The uncrossing process is algorithmically inefficient, but we only need to demonstrate existence for the purpose of this paper. The arguments in this section represent the primary technical contribution of this paper; these arguments apply to bipartite graphs in general (not necessarily regular), and may be independently interesting.

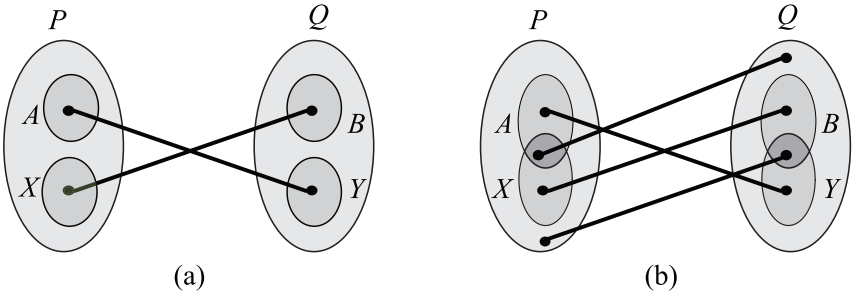

Figure 1. Both (a) and (b) depict two $\frac{1}{3}$-thick pairs (A,B) and (X,Y) that have different witness sets but the same cut (i.e. $W(A,B)\neq W(X,Y)$ but $C(A,B)=C(X,Y)$). The pairs in (a) can not be uncrossed, whereas the pairs in (b) can be uncrossed by choosing the single pair $(A\cap X, B\cap Y)$ as a representative.

3.1. Proportionate uncrossings: definitions and properties

Definition 3.2. A γ-uncrossing of a γ-thick collection $\mathcal{R}$ is another γ-thick collection of pairs $\mathcal{T}$ that satisfies the three properties below:

P1: For every pair $(A,B)\in\mathcal{R}$ there exists a pair $(A',B')\in\mathcal{T}$ such that $C(A',B')\subseteq C(A,B)$, and $W(A',B')\subseteq W(A,B)$. We will refer to (A',B') as a representative of (A,B).

P2: For every $(A',B')\in\mathcal{T}$, there exists $(A,B)\in\mathcal{R}$ such that $C(A',B')\subseteq C(A,B)$.

P3: *(Half-injectivity):* There can not be three distinct pairs $(A,B),(A',B')$, and (A'',B'') in $\mathcal{T}$ such that $C(A,B)=C(A',B')=C(A'',B'')$.

Since $\mathcal{T}$ has the same (or larger) thickness as the thickness guarantee that we had for $\mathcal{R}$, it seems appropriate to refer to $\mathcal{T}$ as a proportionate uncrossing of $\mathcal{R}$.

Definition 3.3. A γ-partial-uncrossing of a γ-thick collection $\mathcal{R}$ is another γ-thick collection of pairs $\mathcal{T}$ which satisfies properties P1, P2 above but not necessarily P3.

The following three lemmas follow immediately from the two definitions above, and it will be useful to state them explicitly. Informally, the first says that every collection is its own partial uncrossing, the second says that uncrossings can be composed, and the third says that the union of the partial uncrossings of two collections is a partial uncrossing of the union of the collections.

Lemma 3.4. *If $\mathcal{R}$ is a γ-thick collection, then $\mathcal{R}$ is a γ-partial uncrossing of itself.*

Lemma 3.5. *If $\mathcal{S}$ is a γ-partial uncrossing of a γ-thick collection $\mathcal{R}$, and $\mathcal{T}$ is a γ-uncrossing of $\mathcal{S}$, then $\mathcal{T}$ is also a γ-uncrossing of $\mathcal{R}$.*

Lemma 3.6. *If $\mathcal{R}_1$ and $\mathcal{R}_2$ are two γ-thick collections, $\mathcal{T}_1$ is a γ-partial-uncrossing of $\mathcal{R}_1$, and $\mathcal{T}_2$ is a γ-partial-uncrossing of $\mathcal{R}_2$, then $\mathcal{T}_1 \cup \mathcal{T}_2$ is a γ-partial-uncrossing of $\mathcal{R}_1 \cup \mathcal{R}_2$.*

3.2. Proportionate uncrossings: an existence theorem

The main technical result of this section is the following:

Theorem 3.7. *For every γ-thick collection $\mathcal{R}$, there exists a γ-uncrossing of $\mathcal{R}$.*

The proof is via induction over the "largest cut" corresponding to any pair in the collection $\mathcal{R}$; each inductive step "uncrosses" the witness sets which corresponds to this largest cut. Before proving this theorem, we need to provide several useful definitions and also establish a key lemma.

Define some total ordering $\prec$ over all subsets of E which respects set cardinality, so that if $|E_1| < |E_2|$, then $E_1 \prec E_2$. Overload notation to use $C(\mathcal{R})$ to denote the set of cuts $\{C(A,B)\colon (A,B) \in \mathcal{R}\}$. Analogously, use $W(\mathcal{R})$ to denote the set of witness sets corresponding to pairs in $\mathcal{R}$. Since $C(A,B)$ may be equal to $C(A',B')$ for $(A,B) \neq (A',B')$, it is possible that $|C(\mathcal{R})|$ may be smaller than $|\mathcal{R}|$. In fact, if $\mathcal{R}$ and $|C(\mathcal{R})|$ are equal, then $\mathcal{R}$ is its own γ-uncrossing and the theorem is trivially true. Similarly, it is possible that $W(A,B)$ is equal to $W(A',B')$ for two different pairs (A,B) and (A',B') in $\mathcal{R}$. However, suppose $W(A,B) = W(A',B')$ and $C(A,B) = C(A',B')$ for two different pairs (A,B) and (A',B') in $\mathcal{R}$. In this case, we can remove one of the two pairs from the collection to obtain a new collection $\mathcal{R}'$; it is easy to see that a γ-uncrossing of $\mathcal{R}'$ is also a γ-uncrossing of $\mathcal{R}$. So we will assume without loss of generality that for any two pairs (A,B) and (A',B') in $\mathcal{R}$, either $W(A,B) \neq W(A',B')$ or $C(A,B) \neq C(A',B')$; we will call this the *non-redundancy* assumption.

We will now prove a key lemma which contains the meat of the uncrossing argument. When we use this lemma later in the proof of Theorem 3.7, we will only use the fact that there exists a γ-partial-uncrossing of $\mathcal{R}$, where $\mathcal{R}$ satisfies the preconditions of the lemma. However, the stronger claim of existence of a γ-uncrossing does not require much additional work and appears to be an interesting graph theoretic argument in its own right, so we prove this stronger claim.

Lemma 3.8. *If $\mathcal{R}$ is a γ-thick collection such that $|\mathcal{R}| > 2$, $\mathcal{R}$ satisfies the non-redundancy assumption, and $C(\mathcal{R})$ contains a single set S, then there* *exists a γ-uncrossing $\mathcal{T}$ of $\mathcal{R}$. Further, for every pair $(A,B)\in\mathcal{T}$, we have $C(A,B)\subset S$.*

Proof. Let $\mathcal{R}=\{(A_1,B_1),(A_2,B_2),\ldots,(A_J,B_J)\}$. Since $C(A_i,B_i)=S$ for all i, we know by the non-redundancy assumption that $W(A_i,B_i)\neq W(A_{i'},B_{i'})$ for $i\neq i'$. We break the proof down into multiple stages.

1. *Definition of Venn witnesses and Venn cuts.* For any J-dimensional bit-vector $b\in\{0,1\}^J$, define

$$A_{(b)}=\left(P\cap\left(\bigcap_{b_i=1}A_i\right)\right)\setminus\left(\bigcup_{b_i=0}A_i\right),$$

and similarly,

$$B_{(b)}=\left(Q\cap\left(\bigcap_{b_i=1}B_i\right)\right)\setminus\left(\bigcup_{b_i=0}B_i\right).$$

We overload notation and use $W_{(b)}$ to denote the witness set $W(A_{(b)},B_{(b)})$ and $C_{(b)}$ to denote the cut set $C(A_{(b)},B_{(b)})$. A node u belongs to $A_{(b)}$ if it is in every set A_i such that $b_i=1$ and not in any of the sets A_i for which $b_i=0$. Thus, each $A_{(b)}$ corresponds to one of the regions in the Venn diagram of the sets $A_1,A_2,\ldots,A_J$, and the analogous statement holds for each $B_{(b)}$. Hence, we will refer to the sets $W_{(b)}$ and $C_{(b)}$ as the Venn-witness and the Venn-cut for b, respectively, and refer to the pair $(A_{(b)},B_{(b)})$ as a Venn pair. Also, we will use $\bar{b}$ to refer to a vector which differs from b in every bit.

2. *The special structure of Venn witnesses and Venn cuts.* Consider an edge (u,v) that goes out of $A_{(b)}$. Suppose that edge goes to $B_{(d)}$ where $d\neq b$ and $d\neq\bar{b}$. Then there must exist $1\leq i,i'\leq J$ such that $b_i=d_i$ and $b_{i'}\neq d_{i'}$. Since $b_i=d_i$, either $u\in A_i,v\in B_i$ (if $b_i=d_i=1$) or $u\notin A_i,v\notin B_i$ (if $b_i=d_i=0$). In either case the edge (u,v) does not belong to the cut $C(A_i,B_i)$, and since all pairs in $\mathcal{R}$ have the same cut S, we conclude that $(u,v)\notin S$. On the other hand, since $b_{i'}\neq d_{i'}$, either $u\in A_{i'},v\notin B_{i'}$ (if $b_{i'}=1,d_{i'}=0$) or $u\notin A_{i'},v\in B_{i'}$ (if $b_{i'}=0,d_{i'}=1$). In either case the edge (u,v) belongs to the cut $C(A_{i'},B_{i'})$ and hence to S, which is a contradiction. Thus, *any edge from $A_{(b)}$ goes to either $B_{(b)}$ or $B_{(\bar{b})}$.*

   If the edge (u,v) goes to $B_{(b)}$, then it does not belong to any witness set in $W(\mathcal{R})$, any Venn witness set, any Venn cut, or S. If (u,v) goes to $B_{(\bar{b})}$, then it belongs to S, to the Venn witness set $W_{(b)}$, to the Venn cuts $C_{(b)}$ and $C_{(\bar{b})}$, and to no other Venn witness set or Venn cut. This edge also belongs to $W(A_i,B_i)$ for all i such that $b_i=1$. These observations, and the definitions of Venn witnesses, cuts, and pairs easily lead to the following consequences:

$$W_{(b)} \cap W_{(d)} = \emptyset \text{ if } b \neq d, \tag{1}$$

$$W(A_i,B_i) = \bigcup_{b\in\{0,1\}^J :\, b_i=1} W_{(b)}, \tag{2}$$

$$C_{(b)} = C_{(\bar{b})}, \tag{3}$$

$$C_{(b)} \cap C_{(d)} = \emptyset \text{ if } b \neq d \text{ and } b \neq \overline{d}, \tag{4}$$

$$(\forall i, 1 \leq i \leq J) \colon S = \bigcup_{b\in\{0,1\}^J :\, b_i=1} C_{(b)}, \tag{5}$$

and finally,

$$W_{(b)} \cup W_{(\bar{b})} = C_{(b)}. \tag{6}$$

3. *The collection $\mathcal{T}$.* Define $\mathcal{T}$ to consist of all γ-thick Venn pairs $(A_{(b)},B_{(b)})$ where b is not the all zero vector.
4. *Proving that $\mathcal{T}$ is a γ-uncrossing of $\mathcal{R}$.* **(P1):** Fix some $i, 1 \leq i \leq J$. Since $\mathcal{R}$ is a γ-thick collection, it follows from the definition that (A_i,B_i) must be a γ-thick pair. From equations 2 and 1, we know that $t(W(A_i,B_i) \cap E_R) = \sum_{b\in\{0,1\}^J :\, b_i=1} t(W_{(b)} \cap E_R)$. We also know, from equations 4 and 5, that $t(S) = \sum_{b\in\{0,1\}^J :\, b_i=1} t(C_{(b)})$. Hence, there must be some $b \in \{0,1\}^J$ such that $b_i=1$ and $(A_{(b)},B_{(b)})$ is γ-thick, which in turn implies that $(A_{(b)},B_{(b)})$ is in $\mathcal{T}$. This is the representative of (A_i,B_i) and hence $\mathcal{T}$ satisfies P1. **(P2):** This follows trivially from equation 5. **(P3):** From equation 4 we know that there are only two possible Venn pairs (specifically, $(A_{(b)},B_{(b)})$ and $(A_{(\bar{b})},B_{(\bar{b})})$) that have the same non-empty cut $C_{(b)}$. Observe that our definition of γ-thickness involves "strict inequality", and hence Venn pairs where the Venn witness set and the Venn cut are both empty can't be γ-thick and can't be in $\mathcal{T}$.
5. *Proving that $C(X,Y) \subset S$ for all pairs $(X,Y) \in \mathcal{T}$.* Any cut $C(A,B) \in C(\mathcal{T})$ is of the form $C_{(b)}$ for some J-dimensional bit vector b. Each $C_{(b)} \subseteq S$, from equation 5. We will now show that this containment is strict. Suppose not, *i.e.*, there exists some $C_{(b)} = S$. By equation 3, $C_{(\bar{b})} = S$ as well. Since $J > 2$, either b or $\overline{b}$ must have two bits that are set to 1; without loss of generality, assume that $b_1 = b_2 = 1$. From equations 1 and 6, we know that $C_{(b)}$ (and hence S) is the disjoint union of $W_{(b)}$ and $W_{(\bar{b})}$. Any edge in $W_{(b)}$ must belong to both $W(A_1,B_1)$ and $W(A_2,B_2)$, whereas any edge in $W_{(\bar{b})}$ can not belong to either $W(A_1,B_1)$ or $W(A_2,B_2)$. Hence,

$W(A_1,B_1) = W(A_2,B_2) = W_{(b)}$ which contradicts the non-redundancy assumption on $\mathcal{R}$. Therefore, we must have $C_{(b)} \subset S$. ▮

Proof of Theorem 3.7. The proof will be by induction over the largest set in $C(\mathcal{R})$ according to the ordering $\prec$. Let $M(\mathcal{R})$ denote this largest set.

For the base case, suppose $M(\mathcal{R})$ is the smallest set S under the ordering $\prec$. Then S must be singleton, $C(\mathcal{R})$ must have just a single set S, and $W(\mathcal{R})$ must also have a single witness set, which must be the same as S since $\mathcal{R}$ is γ-thick. By the non-redundancy assumption, $\mathcal{R}$ must have at most one pair, and is its own γ-uncrossing.

For the inductive step, consider any possible cut S and assume that the theorem is true when $M(\mathcal{R}) \prec S$. We will show that the theorem is also true when $M(\mathcal{R}) = S$, which will complete the inductive proof.

Suppose there is a unique $(A,B) \in \mathcal{R}$ such that $C(A,B) = S$. Intuitively, one would expect this to be the easy case, since there is no "uncrossing" to be done for S, and indeed, this case is quite straightforward. Define $\mathcal{R}' = \mathcal{R} - (A,B)$. Let $\mathcal{T}'$ denote a γ-uncrossing of $\mathcal{R}$, which is guaranteed to exist by the inductive hypothesis. Since $\mathcal{T}'$ is γ-thick, so is $\mathcal{T} = \mathcal{T}' \cup \{(A,B)\}$. The pair (A,B) clearly has a representative in $\mathcal{T}$ (itself), and any $(A',B') \in \mathcal{R} - (A,B)$ has a representative in $\mathcal{T}'$ and hence also in $\mathcal{T}$. Thus, $\mathcal{T}$ satisfies property P1 for being a γ-uncrossing of $\mathcal{R}$. Every set in $C(\mathcal{T}')$ is a subset of some cut in $C(\mathcal{R}')$ (by property P2) and $C(A,B)$ is also in $C(\mathcal{R})$, and hence $\mathcal{T}$ satisfies property P2 for being a γ-uncrossing of $\mathcal{R}$. Every set in $\mathcal{T}'$ is smaller than $C(A,B)$ according to $\prec$ and $\mathcal{T}'$ satisfies property P3. Hence, $\mathcal{T}$ also satisfies property P3. Thus, $\mathcal{T}$ is a γ-uncrossing of $\mathcal{R}$. If there are exactly two distinct pairs (A,B) and (A',B') in $\mathcal{R}$ such that $C(A,B) = C(A',B') = S$, then the same argument works again, except that $\mathcal{R}' = \mathcal{R} \setminus \{(A,B),(A',B')\}$ and $\mathcal{T} = \mathcal{T}' \cup \{(A,B),(A',B')\}$.

We now need to tackle the most interesting case of the inductive step, where there are more than two pairs in $\mathcal{R}$ that correspond to the same cut S. Write $\mathcal{R} = \mathcal{R}_1 \cup \mathcal{R}_2$ where $C(A,B) \prec S$ for all $(A,B) \in \mathcal{R}_1$ and $C(A,B) = S$ for all $(A,B) \in \mathcal{R}_2$. Recall that for two different pairs (A,B) and (A',B') in $\mathcal{R}_2$, we must have $W(A,B) \neq W(A',B')$ by the non-redundancy assumption. From Lemma 3.8, there exists a γ-partial-uncrossing, say $\mathcal{S}_2$, of $\mathcal{R}_2$ with the property that for every set $S' \in C(\mathcal{S}_2)$, we have $S' \subset S$, and hence $S' \prec S$. By Lemma 3.4, we know that $\mathcal{R}_1$ is its own γ-partial-uncrossing. Further, by definition of $\mathcal{R}_1$, every set $S' \in C(\mathcal{R}_1)$ must satisfy $S' \prec S$. Define $\mathcal{S} = \mathcal{R}_1 \cup \mathcal{S}_2$. By Lemma 3.6, $\mathcal{S}$ is a γ-partial-uncrossing of $\mathcal{R}_1 \cup \mathcal{R}_2$, *i.e.*, of $\mathcal{R}$. Further, for every cut $S' \in C(\mathcal{S})$, we have $S' \prec S$. Hence, by our inductive hypothesis, there exists a γ-uncrossing of $\mathcal{S}$; let $\mathcal{T}$ be a γ-uncrossing of $\mathcal{S}$. By Lemma 3.5, $\mathcal{T}$ is also a γ-uncrossing of $\mathcal{R}$, which completes the inductive proof. ▮

Remark 3.1. An alternate approach to relating cuts and witness sets is to suitably modify the proof of the Benczúr-Karger sampling theorem, circumventing the need for the proportionate uncrossing theorem. The idea is based on the observation that Karger's sampling theorem also holds for vertex cuts in graphs. Since Benczúr-Karger sampling theorem is proved using multiple invocations of Karger's sampling theorem, it is possible to set up a correspondence between cuts and witness sets using a vertex-cut version of the Benczúr-Karger sampling theorem. However, we prefer to use here the approach based on the proportionate uncrossing theorem as it is an interesting combinatorial statement in its own right.

4. An $\tilde{O}(n^{1.5})$ time algorithm for finding a perfect matching

We present here an $\tilde{O}(n^{1.5})$ time randomized algorithm to find a perfect matching in a given d-regular bipartite graph $G(P,Q,E)$ on $2n$ vertices. Throughout this section, we follow the convention that for any pair (A,B), the sets $C(A,B)$ and $W(A,B)$ are defined with respect to the graph G. Our starting point is the following theorem, established by Goel, Kapralov, and Khanna [8].[2]

Theorem 4.1. *Let $G(P,Q,E)$ be a d-regular bipartite graph, ϵ any number in $(0,\frac{1}{2})$, and c a suitably large constant that depends on ϵ. There exists a decomposition of G into $k=O(n/d)$ vertex-disjoint bipartite graphs, say $G_1=(P_1,Q_1,E_1), G_2=(P_2,Q_2,E_2),\dots,G_k=(P_k,Q_k,E_k)$, such that*

1. *Each G_i contains at least $d/2$ perfect matchings, and the minimum cut in each G_i is $\Omega(d^2/n)$.*
2. *Let $\mathcal{R}$ denote the set of relevant pairs with respect to this decomposition, and E_R denote the set of relevant edges. Then for each (A,B) in $\mathcal{R}$, we have $|W(A,B)\cap E_R|\geq\frac{1}{2}|C(A,B)|$.*
3. *Let $G'(P,Q,E')$ be a random graph generated by sampling the edges of G uniformly at random with probability $p=\frac{cn\ln n}{d^2}$. Then with probability at least $1-1/n$, for every pair $(A,B)\in\mathcal{R}$,*

$$|W(A,B)\cap E'\cap E_R| > (1-\epsilon)p|W(A,B)\cap E_R| > \left(\frac{1-\epsilon}{2(1+\epsilon)}\right)|C(A,B)\cap E'|.$$

The last condition above says that in addition to all cuts, all relevant witness edge sets are also preserved to within $(1\pm\epsilon)$ of their expected value

[2] Part 1 of Theorem 4.1 corresponds to theorem 2.3 in [8], part 2 is proved as part of the proof of theorem 2.1 in [8], and part 3 combines remark 2.5 in [8] with Karger's sampling theorem [10].

in G', with high probability. We emphasize here that the decomposition highlighted in Theorem 4.1 will be used only in the analysis of our algorithm; the algorithm itself is oblivious to this decomposition.

Our algorithm consists of the following three steps.

(S1) Generate a random graph $G' = (P, Q, E')$ by sampling edges of G *uniformly* at random with probability $p = \min\{1, \frac{c_1 n \ln n}{d^2}\}$ where c_1 is a constant as in Theorem 4.1.[3] We choose ϵ to be any fixed constant not larger than 0.2.

(S2) The graph G' contains $O(\frac{n^2 \ln n}{d})$ edges whp. We now run the Benczúr-Karger sampling algorithm [2] that takes $O(|E'| \ln^2 n)$ time to compute the strength s_e of every edge e and samples each edge e with probability p_e;[4] here p_e is as given by Theorem 2.6 with $\gamma = 1/3$. We show below that whp the graph $G'' = (P, Q, E'')$ obtained from this sampling contains a perfect matching. Also, let G''' be obtained by including each edge of G independently with probability $\min\{1, (c_2 \ln n)/d\}$ for a sufficiently large constant c_2.

(S3) Finally, we run the Hopcroft-Karp algorithm to obtain a maximum cardinality matching in $G'' \cup G'''$ in $O(n^{1.5} \ln n)$ time since G'' contains $O(n \ln n)$ edges whp by Theorem 2.7, and G''' has $O(n \ln n)$ edges whp by a simple application of Chernoff bounds. The rationale behind running the Hopcroft-Karp algorithm on $G'' \cup G'''$ is that G'' contains a perfect matching whp and the structure of G''' ensures that this matching can be found fast (see section 5 for the improved analysis of the Hopcroft-Karp algorithm that makes use of the structure of G''').

Running time: With high probability, the running time of this algorithm is bounded by $O(\frac{n^2}{d} \ln^3 n + n^{1.5} \ln n)$. Since we can always use the algorithm of Cole, Ost, and Schirra [6] instead, the final running time is $O(\min\{m, \frac{n^2}{d} \ln^3 n + n^{1.5} \ln n\})$ using standard bounds on the runtime of the Hopcroft-Karp algorithm. This reduces to $O(m)$ if $d \leq \sqrt{n} \ln n$; to $O(n^{1.5} \ln n)$ when $d \geq \sqrt{n} \ln^2 n$; and to $O((n \ln n)^{1.5})$ in the narrow range $\sqrt{n} \ln n < d < \sqrt{n} \ln^2 n$.

Correctness: To prove correctness, we need to show that G'' contains a perfect matching whp.

[3] The time required for this sampling is proportional to the number of edges chosen, assuming the graph is presented in an adjacency list representation with each list stored in an array.

[4] In fact, this sampling algorithm computes an upper bound on s_e, but this only affects the running time and the number of edges sampled by a constant factor.

Theorem 4.2. *The graph G'' contains a perfect matching with probability $1-O(1/n)$.*

Proof. Consider the decomposition defined in Theorem 4.1. Let $\mathcal{R}$ denote the set of relevant pairs with respect to this decomposition, and let E_R denote the set of all relevant edges with respect to this decomposition. We will now focus on proving that, with high probability, for every $(A,B)\in\mathcal{R}$, $W(A,B)\cap E_R\cap E''\neq\emptyset$; by Lemma 2.3, this is sufficient to prove the theorem.

For convenience, define $W'(A,B)=W(A,B)\cap E'$ and $C'(A,B)=C(A,B)\cap E'$. Assume for now that the low-probability event in Theorem 4.1 does not occur. Thus, by choosing $\epsilon\leq 0.2$, we know that for $\gamma=1/3$, every relevant pair $(A,B)\in\mathcal{R}$ satisfies $|W'(A,B)\cap E_R|>\gamma|C'(A,B)|$. Let s'_e denote the strength of e in G'. Recall that $G'_{[j]}=(V,E'_{[j]})$ is the graph with the same vertex set as G' but consisting of only those edges in E' which have strength at least j. Define $W'_{[j]}(A,B)$ to be the set of all edges in $W'(A,B)\cap E'_{[j]}$; define $C'_{[j]}(A,B)$ analogously. Finally, let $t(e)=1/s'_e$. Since $|W'(A,B)\cap E_R|>\gamma|C'(A,B)|$, by Lemma 2.8, there must exist a j such that

$$\sum_{e\in(W'(A,B)\cap E_R),s'_e\geq j}\frac{1}{s'_e}>\gamma\sum_{e\in C'(A,B),s'_e\geq j}\frac{1}{s'_e}>0,$$

which implies that (A,B) is γ-thick with respect to $(G'_{[j]},t,E_R)$, as defined in Definition 3.1. Partition $\mathcal{R}$ into $\mathcal{R}_{[1]},\mathcal{R}_{[2]},\ldots,\mathcal{R}_{[n]}$, such that if $(A,B)\in\mathcal{R}_{[j]}$, then (A,B) is γ-thick with respect to $(G'_{[j]},t,E_R)$, breaking ties arbitrarily if (A,B) can belong to multiple $\mathcal{R}_{[j]}$. Consider an arbitrary non-empty $\mathcal{R}_{[j]}$. Let $\mathcal{T}$ represent a γ-uncrossing of $\mathcal{R}_{[j]}$, as guaranteed by Theorem 3.7. By property P3 in Definition 3.1, no three pairs in a γ-uncrossing can have the same cut; partition $\mathcal{T}$ into $\mathcal{T}_1$ and $\mathcal{T}_2$ such that every pair $(A,B)\in\mathcal{T}_1$ has a unique cut $C'_{[j]}(A,B)$ and the same holds for $\mathcal{T}_2$. We focus on $\mathcal{T}_1$ for now. For any $(A,B)\in\mathcal{T}_1$, define $Y(A,B)=W'_{[j]}(A,B)\cap E_R$. Define $\mathcal{X}=\{Y(A,B)\colon (A,B)\in\mathcal{T}_1\}$. For any $X\in\mathcal{X}$, define $f(X)=C'_{[j]}(A,B)$ for some arbitrary $(A,B)\in\mathcal{T}_1$ such that $X=Y(A,B)$. The function f is one-one by construction, and since (A,B) is γ-thick, we know that $\sum_{e\in X}1/s'_e>\gamma\sum_{e\in f(X)}1/s'_e$. Thus, $\mathcal{X}$ satisfies the preconditions of Theorem 2.6. Further, the sampling probability p_e in step **(S2)** of the algorithm is chosen to correspond to $\gamma=1/3$. Thus, with probability at least $1-1/n^2$, $X\cap E''$ is non-empty for all $X\in\mathcal{X}$, *i.e.*, $W'_{[j]}(A,B)\cap E_R\cap E''\neq\emptyset$ for all $(A,B)\in\mathcal{T}_1$. Since $G'_{[j]}$ is a subgraph of G', we can conclude that $W'(A,B)\cap E_R\cap E''\neq\emptyset$ for all $(A,B)\in\mathcal{T}_1$ with probability at least $1-1/n^2$.

Since the analogous argument holds for $\mathcal{T}_2$, we obtain $W'(A,B)\cap E_R\cap E''\neq\emptyset$ for all $(A,B)\in\mathcal{T}$ with probability at least $1-2/n^2$. As $\mathcal{T}$ is a γ-uncrossing of $\mathcal{R}_{[j]}$, we use property P1 to conclude that $W'(A,B)\cap E_R\cap E''\neq\emptyset$ for all $(A,B)\in\mathcal{R}_{[j]}$, again with probability at least $1-2/n^2$. Applying the union bound over all j, we further conclude that $W'(A,B)\cap E_R\cap E''\neq\emptyset$ for all $(A,B)\in\mathcal{R}$ with probability at least $1-2/n$. As mentioned before, this suffices to prove that G'' has a perfect matching with probability at least $1-2/n$, by Lemma 2.3. We assumed that condition 3 in Theorem 4.1 is satisfied; this is violated with probability at most $\frac{1}{n}$, which proves that G'' has a perfect matching with probability at least $1-\frac{3}{n}$. ∎

As presented above, the algorithm takes time $\min\{\tilde{O}(n^{1.5}),O(m)\}$ with high probability, and outputs a perfect matching with probability $1-O(1/n)$. We conclude with two simple observations. First, it is easy to convert this into a Monte Carlo algorithm with a worst case running-time of $\min\{\tilde{O}(n^{1.5}),O(m)\}$, or a Las Vegas algorithm with an expected running-time of $\min\{\tilde{O}(n^{1.5}),O(m)\}$. If either the sampling process in steps **(S1)** or **(S2)** returns too many edges, or step **(S3)** does not produce a perfect matching, then (a) abort the computation to get a Monte Carlo algorithm, or (b) run the $O(m)$ time algorithm of Cole, Ost, and Schirra [6] to get a Las Vegas algorithm. Second, by choosing larger constants during steps **(S1)** and **(S2)**, it is easy to amplify the success probability to be at least $1-O\left(\frac{1}{n^j}\right)$ for any fixed $j\geq 1$.

5. An improved $O\left(\min\{nd,(n^2\ln^3 n)/d\}\right)$ bound on the runtime

In this section we give an improved analysis of the runtime of the Hopcroft-Karp algorithm on the subsampled graph, leading to an overall bound of $O\left(\min\{nd,(n^2\ln^3 n)/d\}\right)$ for our algorithm. We first give intuition behind the analysis, and then proceed to the technical details.

The runtime bound, which is proved in subsection 5.3, relies on breaking the execution of the algorithm into two parts. The first part corresponds to augmentation phases of the Hopcroft-Karp algorithm during which the size of the matching is at most $n-O(n/d)$. The second part corresponds to the final augmentation phases, during which the size of the matching is increased from $n-O(n/d)$ to n. The runtime of the second part is easy to bound: a phase of the Hopcroft-Karp algorithm takes time linear in the number of edges in the graph, which is $O(n\ln n)$ for the graphs G'' and G''' that are used in step **S3**. There can be at most $O(n/d)$ augmentations in the second part, so the bound of $O((n^2\log n)/d)$ follows. Thus, most of the development in this section is directed towards bounding the runtime of the first part. To prove the stated bound, we show in subsection 5.3 that during the execution of the first part, the length of augmenting paths does not exceed $\tilde{O}(n/d)$ (see Lemma 5.7). This is achieved by relating the length of augmenting paths with respect to partial matchings of size at most $n-O(n/d)$ to an expansion property for the set of witness sets $W(A,B)$ that naturally arise in the Hopcroft-Karp algorithm (see subsection 5.3). More precisely, these are witness sets $W(A,B)$ that are sufficiently *unbalanced* in the sense that $|A|>|B|+2n/d$. Their properties and their behavior under aggressive sampling of a d-regular bipartite graph G (at rate $p=c\log n/d$ for a constant $c>0$, the sampling rate that is used to obtain the graph G''' in **S2**) are studied in subsections 5.1 and 5.2. The main result of subsections 5.1 and 5.2 is Corollary 5.6, which lower bounds the size of the witness set $W(A,B), |A|>|B|+2n/d$ in terms of the size of the witness set $W(B,A)$ in such subsampled graphs (for $A\subseteq P, B\subseteq Q$, the set $W(B,A)$ is defined as $C(A,B)\setminus W(A,B)$). This corollary is later used in the main argument of subsection 5.3.

5.1. Combinatorial uncrossings

Theorem 5.2 below, which we state for general bipartite graphs, requires a variant of the uncrossing theorem that we formulate now. We introduce the definition of combinatorial uncrossings:

Definition 5.1. Let $\mathcal{R}$ be any collection of pairs $(A,B), A\subseteq P, B\subseteq Q$. A combinatorial uncrossing of $\mathcal{R}$ is a tuple $(\mathcal{T},\mathcal{I})$, where $\mathcal{T}$ is another collection and $\mathcal{I}$ is a mapping from $\mathcal{R}$ to subsets of $\mathcal{T}$, such that the following properties are satisfied:

Q1: For all $(A,B)\in\mathcal{R}$

1. $\{W(A',B')\}_{(A',B')\in\mathcal{I}(A,B)}$ are disjoint;
2. $\{C(A',B')\}_{(A',B')\in\mathcal{I}(A,B)}$ are disjoint;
3. $\{A'\cup B'\}_{(A',B')\in\mathcal{I}(A,B)}$ are disjoint;
4. $A'\subseteq A, B'\subseteq B$ for all $(A',B')\in\mathcal{I}(A,B)$;
5.

$$W(A,B) = \bigcup_{(A',B')\in\mathcal{I}(A,B)} W(A',B')$$
$$C(A,B) = \bigcup_{(A',B')\in\mathcal{I}(A,B)} C(A',B').$$

Q2: (Half-injectivity) There cannot be three distinct pairs (A,B), (A',B'), (A'',B'') in $\mathcal{T}$ such that $C(A,B)=C(A',B')=C(A'',B'')$.

The proof of existence of combinatorial uncrossings is along the lines of the proof of existence of γ-thick uncrossings (see Theorem 3.7), so we omit it here.

For a graph H we denote $W_H(A,B)=W(A,B)\cap E(H)$ and $C_H(A,B)=C(A,B)\cap E(H)$, and omit the subscript when the underlying graph is fixed.

Theorem 5.2. *There exists a constant $c>0$ such that for all $\epsilon>0$ such that for all bipartite graphs $G=(P,Q,E)$, $|P|=|Q|=n$, with a minimum cut of size at least κ the following holds. If a graph G^* is obtained by sampling edges of G uniformly at random with probability $p>\frac{c\ln n}{\epsilon^2\kappa}$, then whp for all $A\subseteq P$, and $B\subseteq Q$, we have*

$$p|W_G(A,B)| - \epsilon p|C_G(A,B)| \leq |W_{G^*}(A,B)| \leq p|W_G(A,B)| + \epsilon p|C_G(A,B)|.$$

Proof. Define $\mathcal{R}$ as the set of pairs $(A,B), A\subseteq P\cap V(G), B\subseteq Q\cap V(G)$. Denote a combinatorial uncrossing of $\mathcal{R}$ by $(\mathcal{T},\mathcal{I})$. We first prove the statement for pairs from $\mathcal{T}$, and then extend it to pairs from $\mathcal{R}$ to obtain the desired result.

Consider a pair $(A,B)\in\mathcal{T}$. Denote $\Delta_G(A,B) = |W_{G^*}(A,B)| - p|W_G(A,B)|$. We shall write $W(A,B)$ and $C(A,B)$ instead of $W_G(A,B)$ and $C_G(A,B)$ in what follows for brevity. We have by Chernoff bounds that for a given pair $(A,B)\in\mathcal{T}$

$$\begin{aligned}\mathbf{Pr}\left[|\Delta_G(A,B)| > \epsilon p|C(A,B)|\right] &< 2\exp\left[-\left(\frac{\epsilon|C(A,B)|}{|W(A,B)|}\right)^2 \frac{p|W(A,B)|}{3}\right]\\ &\leq 2\exp\left[-\epsilon^2\left(\frac{p|C(A,B)|}{3}\right)\right]\end{aligned}$$

since $|C(A,B)|\geq|W(A,B)|$. Since $\mathcal{T}$ satisfies Q2, we get that

$$\begin{aligned}&\mathbf{Pr}\left[\exists(A,B)\in\mathcal{T}\colon\ |\Delta_G(A,B)| > \epsilon p|C(A,B)|\right]\\ &< 2\sum_{W(A,B)\in W(\mathcal{T})} \exp\left[-\epsilon^2 p|C(A,B)|/3\right]\\ &\leq 4\sum_{C(A,B)\in C(\mathcal{T})} \exp\left[-\epsilon^2 p|C(A,B)|/3\right] = O(n^{-r})\end{aligned}$$

for $c=3(r+2)$ by Corollary 2.4 in [10]. This implies that for $c\geq 3(r+2)$ we have with probability $1-O(n^{-r})$ for all $(A,B)\in\mathcal{T}$

$$|\Delta_G(A,B)| \leq \epsilon p|C(A,B)|. \tag{7}$$

Now consider any pair $(A,B)\in\mathcal{R}$. Summing (7) over all $(A',B')\in\mathcal{I}(A,B)$ and using properties Q1.1-5, we get

$$|\Delta_G(A,B)| \leq \sum_{(A',B')\in\mathcal{I}(A,B)} \epsilon p|C(A',B')| = \epsilon p|C(A,B)|,$$

for all $(A,B)\in\mathcal{R}$ as required. ∎

5.2. Decomposition of the graph G

Corollary 5.6, which relates the size of sufficiently unbalanced witness sets in the sampled graph to the size of the corresponding cuts is the main result of this subsection. It follows from Theorem 5.2 and a stronger (than [8]) decomposition of bipartite d-regular graphs that we outline now. For $A\subseteq V(G)$ we denote the set of edges in the cut $(A,V(G)\setminus A)$ in G by $\delta(A)$.

Theorem 5.3. *Any d-regular graph G with $2n$ vertices can be decomposed into vertex-disjoint induced subgraphs $G_1=(P_1,Q_1,E_1)$, $G_2=(P_2,Q_2,E_2)$, $\dots$, $G_k=(P_k,Q_k,E_k)$, where $k\le 4n/d+1$, that satisfy the following properties:*

1. *The minimum cut in each G_i is at least $d/8$.*
2. $\sum_{i=1}^{k}|\delta_G(V(G_i))|\le 2n$.

To prove Theorem 5.3, we give a procedure that decomposes the graph G into vertex-disjoint induced subgraphs $G_1(P_1,Q_1,E_1)$, $G_2(P_2,Q_2,E_2)$, $\dots$, $G_k(P_k,Q_k,E_k)$, $k\le 4n/d+1$ such that the min-cut in G_j is at least $d/8$ and at most n edges run between pieces of the decomposition.

The procedure is as follows. Initialize $H_1:=G$, and set $i:=1$.

1. Find a smallest proper subset $X_i\subset V(H_i)$ such that $|\delta_{H_i}(X_i)|<d/4$. If no such set exists, define G_i to be the graph H_i and terminate.
2. Define G_i to be the subgraph of H_i induced by vertices in X_i, i.e. $X_i=P_i\cup Q_i=V(G_i)$. Also, define H_{i+1} to be the graph H_i with vertices from X_i removed.
3. Increment i and go to step 1.

We now prove that the output of the decomposition procedure satisfies the properties claimed above.

Lemma 5.4. *The min-cut in G_i is greater than $d/8$.*

Proof. If G_i contains a single vertex the min-cut is infinite by definition, so we assume wlog that G_i contains at least two vertices. The proof is essentially the same as the proof of property **P1** of the decomposition procedure in [8] (see Theorem 2.4).

Suppose that there exists a cut (V,V^c) in G_i where $V\subset V(G_i)$ and $V^c=V(G_i)\backslash V$, such that $|\delta_{G_i}(V)|\le d/8$ (note that it is possible that $V\cap P_i\neq\emptyset$ and $V\cap Q_i\neq\emptyset$). We have $|\delta_{H_i}(V)\backslash\delta_{G_i}(V)|+|\delta_{H_i}(V^c)\backslash\delta_{G_i}(V^c)|<d/4$ by the choice of X_i in (1). Suppose without loss of generality that $|\delta_{H_i}(V)\setminus\delta_{G_i}(V)|<d/8$. Then $|\delta_{H_i}(V)|<d/4$ and $V\subset X_i$, which contradicts the choice of X_i as the smallest cut of value at most $d/4$ in step (1) of the procedure. ∎

Lemma 5.5. *The number of steps in the decomposition procedure is $k \le 4n/d$, and at most n edges are removed in the process.*

Proof. We call a vertex $v \in V(G_i)$ *bad* if its degree in G_i is smaller than $d/2$. Note that for each $1 \le i \le k$ either G_i contains a bad vertex or $|V(G_i)| \ge d$.

Note that since strictly fewer than $d/4$ edges are removed in each iteration, the number of bad vertices created in the first j iterations is strictly less than $j(d/4)/(d/2) = j/2$. Hence, during at least half of the j iterations at least d vertices were removed from the graph, i.e.

$$\sum_{i=1}^{j} |V(G_i)| \ge (j/2) \cdot d = jd/2.$$

This implies that the process terminates in at most $4n/d$ steps, and the number of edges removed is at most $(4n/d) \cdot d/4 = n$. ▮

Proof of Theorem 5.3. The proof follows by putting together Lemmas 5.4 and 5.5. ▮

We overload notation here by denoting $W(B,A) = W(P \setminus A, Q \setminus B) = C(A,B) \setminus W(A,B)$ for $A \subseteq P, B \subseteq Q$. The main result of this subsection is the following.

Corollary 5.6. *Let $G^* = (P,Q,E^*)$ be a graph obtained by sampling the edges of a d-regular bipartite graph $G = (P,Q,E)$ on $2n$ vertices independently with probability p. There exists a constant $c > 0$ such that if $p > \frac{c \ln n}{\epsilon^2 d}$, then whp for all pairs $(A,B), A \subseteq P, B \subseteq Q$, $|A| \ge |B| + 2n/d$ one has that $|W(A,B) \cap E^*| > \frac{1-3\epsilon}{1+3\epsilon} |W(B,A) \cap E^*|$ for all $\epsilon < 1/10$. In particular, G^* contains a matching of size at least $n - 2n/d$ whp.*

Proof. Set $A_i = A \cap P_i, B_i = B \cap Q_i$, where $G_i = (P_i, Q_i, E_i)$ are the pieces of the decomposition obtained in Theorem 5.3. For each (A_i, B_i) such that G_i is not an isolated vertex we have by Lemma 5.4 and Theorem 5.2

$$||W_{G_i}(A_i,B_i) \cap E^*| - p|W_{G_i}(A_i,B_i)|| < \epsilon p |C_{G_i}(A_i,B_i)|.$$

If G_i is an isolated vertex, we have $|W_{G_i}(A_i,B_i) \cap E^*| = p|W_{G_i}(A_i,B_i)| = 0$. Since the latter estimate is stronger than the former, we shall not consider the isolated vertices separately in what follows.

Adding these inequalities over all i we get

$$\sum_{i=1}^{k} |W_{G_i}(A_i,B_i) \cap E^*| \ge p \sum_{i=1}^{k} |W_{G_i}(A_i,B_i)| - \epsilon p \sum_{i=1}^{k} |C_{G_i}(A_i,B_i)|. \quad (8)$$

Denote the set of edges removed during the decomposition process by E_r. Denote $E_1 = E_r \cap W(A,B)$ and $E_2 = E_r \cap W(B,A)$. Since $|W(A,B) \cap E^*| = \sum_{i=1}^k |W_{G_i}(A_i,B_i) \cap E^*| + |E_1 \cap E^*|$ and $\sum_{i=1}^k |W_{G_i}(A_i,B_i)| = |W(A,B)| - |E_1|$, this implies

$$|W(A,B) \cap E^*| \geq p|W(A,B)| - \epsilon p|C(A,B)| - p|E_1|. \tag{9}$$

Likewise, since $W(B,A) = W(P \setminus A, Q \setminus B)$, we have

$$|W(B,A) \cap E^*| \leq p|W(B,A)| + \epsilon p|C(A,B)| + p|E_2|.$$

Since $|A| \geq |B| + 2n/d$, we have $|W(A,B)| \geq |W(B,A)| + 2n$, so

$$\begin{aligned}|W(A,B) \cap E^*| &\geq p|W(A,B)| - \epsilon p|C(A,B)| - p|E_1| \\ &\geq p(|W(B,A)| + 2n) - \epsilon p|C(A,B)| - p|E_1| - p|E_2| \\ &\geq |W(B,A) \cap E^*| - 2\epsilon p|C(A,B)| + p(2n - |E_r|) \\ &\geq |W(B,A) \cap E^*| - 2\epsilon p|C(A,B)| + pn.\end{aligned}$$

Adding (9) for the pairs (A,B) and (B,A), we get $|C(A,B) \cap E^*| \geq (1-2\epsilon)p(|C(A,B)| - n)$, i.e. $p|C(A,B)| \leq \frac{1}{1-2\epsilon}|C(A,B) \cap E^*| + pn$. Hence, we have

$$\begin{aligned}&|W(A,B) \cap E^*| \geq |W(B,A) \cap E^*| - 2\epsilon p|C(A,B)| + pn \\ &\geq |W(B,A) \cap E^*| - \frac{2\epsilon}{1-2\epsilon}|C(A,B) \cap E^*| + (1-2\epsilon)pn \\ &\geq |W(B,A) \cap E^*| - \frac{2\epsilon}{1-2\epsilon}\left(|W(A,B) \cap E^*| + |W(B,A) \cap E^*|\right) + (1-2\epsilon)pn,\end{aligned}$$

which implies

$$|W(A,B) \cap E^*| > \frac{1-3\epsilon}{1+3\epsilon}|W(B,A) \cap E^*|$$

for $\epsilon < 1/10$. This completes the proof. ∎

Remark 5.1. The result in Corollary 5.6 is tight up to an $O(\ln d)$ factor for $d = \Omega(\sqrt{n})$.

Proof. The following construction gives a lower bound of $n - \Omega\left(\frac{n}{d \ln d}\right)$. Denote by $G_{n,d}$ the graph from Theorem 4.1 in [8] and denote by $G^*_{n,d}$ a graph obtained by sampling edges of $G_{n,d}$ at the rate of $\frac{c \ln n}{d}$ for a constant $c > 0$. Define the graph G as d disjoint copies of $G_{2d \ln d, d}$, and denote the sampled graph by G^*. Note that by Theorem 4.1 the maximum matching in each copy of $G^*_{2d \ln d, d}$ has size at most $2d \ln d - 1$ whp, and since the number of vertices in G is $N = 2d^2 \ln d$, the maximum matching in G^* has size at most $N - \Omega\left(\frac{N}{d \ln d}\right)$ whp. ∎

5.3. Runtime analysis of the Hopcroft-Karp algorithm

In this section we derive a bound on the runtime of the Hopcroft-Karp algorithm on the subsampled graph obtained in step **S2** of our algorithm. The main object of our analysis is the alternating level graph, which we now define. Given a partial matching of a graph $G=(P,Q,E)$, the alternating level graph is defined inductively. Define sets A_j and B_j, $j=1,\dots,L$ as follows. Let A_0 be the set of unmatched vertices in P and let $B_0=\emptyset$. Then let $B_{j+1}=\Gamma(A_j)\setminus\left(\bigcup_{i<j}B_i\right)$, where $\Gamma(A)$ is the set of neighbours of vertices in $A\subseteq V(G)$, and let A_j be the set of vertices matched to vertices from B_j. The construction terminates when either B_{j+1} contains an unmatched vertex or when $B_{j+1}=\emptyset$, and then we set $L=j$. We use the notation $A^{(j)}=\bigcup_{k\le j}A_k, B^{(j)}=\bigcup_{k\le j}B_k$. We now give an outline of the Hopcroft-Karp algorithm for convenience of the reader. Given a non-maximum matching, the algorithm starts by constructing the alternating level graph described above and stops when an unmatched vertex is found. Then the algorithm finds a maximal set of vertex-disjoint augmenting paths of length L (this can be done by depth-first search in $O(m)$ time) and performs the augmentations, thus completing one augmentation phase. It can be shown that each augmentation phase increases the length of the shortest augmenting path. Standard analysis of the run-time for general bipartite graphs is based on the observation that once $\sqrt{n}$ augmentations have been performed, the constructed matching necessarily has size at most $\sqrt{n}$ smaller than the maximum matching.

We denote the graph obtained by sampling edges of G independently with probability $p=\frac{c\ln n}{d}$ for a constant $c>0$ by G^*, and let G^{**} be a graph obtained by adding an arbitrary set of edges of G to G^*. For $A\subseteq V(G)$ denote the set of edges in the cut $(A,V(G)\setminus A)$ in G by $\delta(A)$ and the set of edges in the same cut in G^* by $\delta^*(A)$. Similarly, we denote the vertex neighbourhood of A in G by $\Gamma(A)$ and the vertex neighbourhood in G^* by $\Gamma^*(A)$. We consider the alternating level graph in G^* and prove that whp for any partial matching of size smaller than $n-2n/d$ for each $1\le j\le L$ either $|B_{j-1}\cup B_j\cup B_{j+1}|=\Omega(d)$ or B_j expands by at least a factor of $\ln n$ in either forward or backward direction ($|B_{j+1}|\ge(\ln n)|B_j|$ or $|B_{j-1}|\ge(\ln n)|B_j|$). This implies that $L=O\left(\frac{n\ln d}{d\ln\ln n}\right)$, thus yielding the same bound on the length of the shortest augmenting path by virtue of Corollary 5.6. The main technical result of this subsection is

Lemma 5.7. *Let the set of edges E^* be obtained by sampling edges of a bipartite d-regular graph $G=(P,Q,E)$ on $2n$ vertices uniformly with probability p. Let $G^*=(P,Q,E^*)$ and $G^{**}=(P,Q,E^*\cup E^{**})$, where E^{**} is an arbitrary subset of E. There exists a constant $c>0$ such that if $p\ge\frac{c\ln n}{d}$,* *then whp for any partial matching in G^* of size smaller than $n-2n/d$ there exists an augmenting path of length $O\left(\frac{n\ln d}{d\ln\ln n}\right)$.*

The following expansion property of the graph G^* will be used to prove Lemma 5.7:

Lemma 5.8. *Define $\gamma(t)=(1-\exp(-t))/t$. For all $\epsilon, t>0$ there exists a constant $c>0$ that depends on t and ϵ such that if G^* is obtained by sampling the edges of G independently with probability $p>\frac{c\ln n}{d}$, then whp for every set $A\subseteq P$, $|A|\le t/p$*

$$|\Gamma^*(A)|\geq(1-\epsilon)dp\gamma(t)|A|.$$

The corresponding claim also holds for $B\subseteq Q$, $|B|\le t/p$.

Proof. Consider a set $A\subseteq P$, $|A|\le t/p$. For $b\in\Gamma(A)$ denote the indicator variable corresponding to the event that at least one edge incident on b and going to A is sampled by X_b, i.e. $X_b=I_{\{b\in\Gamma^*(A)\}}$. Denote the number of edges between b and vertices of A by k_b. We have

$$\mathbf{Pr}[X_b=1]=1-(1-p)^{k_b}\geq 1-\exp(-k_bp)\geq k_bp\gamma(t),$$

since $k_bp\le t$ and $e^{-x}\le 1-\gamma(t)x$ for $x\in[0,t]$.

Hence,

$$\mathbf{E}\left[\sum_{b\in B}X_b\right]\geq\gamma(t)p\sum_{b\in B}k_b\geq p|\delta(A)|\gamma(t). \tag{10}$$

There are at most n^s subsets A of P of size s and $|\delta(A)|=d|A|$ for all A, so we obtain using Chernoff bounds and the union bound

$$\begin{aligned}\mathbf{Pr}\left[\exists\, A\subseteq P, |A|\leq t/p\colon |\Gamma^*(A)|<(1-\epsilon)pd|A|\gamma(t)\right]\\ <\sum_{s=1}^{t/p}n^s\exp\left(-\epsilon^2pds\gamma(t)\right)\\ =\sum_{s=1}^{n}\exp\left(s(1-c\gamma(t))\ln n\right)=O(n^{2-c\gamma(t)}),\end{aligned}$$

where we summed a geometric sequence with ratio $n^{1-c\gamma(t)}$ in the last step, assuming that $c>(2+r)/\gamma(t)$. ∎

Proof of Lemma 5.7. Let $p = \frac{c' \ln n}{\epsilon^2 d}$ for a sufficiently large constant $c' > 0$ and some $\epsilon > 0$ that we will later fix to an absolute constant.

First note that since the partial matching is of size strictly less than $n - 2n/d$, by Corollary 5.6 there exists an augmenting path with respect to the partial matching.

In order to upperbound the length of the shortest augmenting path, we will show that for each j, at least one of the following is true:

1. $|B_j| \geq d/500$;
2. $|B_{j+1}| \geq d/500$;
3. $|B_{j+1}| \geq (\ln n)|B_j|$;
4. $|B_{j-1}| \geq d/500$;
5. $|B_{j-1}| \geq (\ln n)|B_j|$.

It then follows that for each j there exists j' such that $|j - j'| \leq 1 + \log_{\ln n} d$ and $|B_{j'}| \geq d/500$. Hence, there cannot be more than $O\left(\frac{n \ln d}{d \ln \ln n}\right)$ levels in the alternating level graph, so there always exists an augmenting path of length $O\left(\frac{n \ln d}{d \ln \ln n}\right)$.

For each $1 \leq j \leq L$, where L is the number of levels in the alternating level graph, we classify the edges in E^* leaving B_j into three classes: (1) E_F contains edges that go to $P \setminus A^{(j)}$, (2) E_M contains edges that go to A_j, and (3) E_R contains edges that go to A_{j-1}. The degree of every node in G^* is between $(1-\epsilon)pd$ and $(1+\epsilon)pd$ whp by Chernoff bounds as long as the constant c' in $p = \frac{c' \ln n}{\epsilon^2 d}$ is sufficiently large. Thus, at least one of E_F, E_M, E_R has at least $(1-\epsilon)pd|B_j|/3$ edges. We now consider each of these possibilities.

Case (A): First suppose that E_F contains at least $(1-\epsilon)pd|B_j|/3$ edges, i.e. $|W(B^{(j)}, A^{(j)})| \geq (1-\epsilon)pd|B_j|/3$. Note that since the partial matching has size smaller than $n - 2n/d$ by assumption, we have that $|A^{(j)}| \geq |B^{(j)}| + 2n/d$. Hence, by Corollary 5.6 applied to $W(A^{(j)}, B^{(j)})$ the number of edges going from A_j to B_{j+1} is at least

$$|W(A^{(j)}, B^{(j)})| > \frac{1-3\epsilon}{1+3\epsilon}|W(B^{(j)}, A^{(j)})|$$
$$\geq \frac{(1-3\epsilon)(1-\epsilon)}{1+3\epsilon} pd|B_j|/3 = \frac{(1-3\epsilon)(1-\epsilon)}{1+3\epsilon} pd|A_j|/3,$$

where we used the fact that $|A_j| = |B_j|$ be definition.

Suppose first that $|A_j| < 1/(5p)$. Then by Lemma 5.8 one has that $|\Gamma^*(A_j)| \geq (1-\epsilon)\gamma(1/5)pd|A_j|$. Let $\beta^* = 1 + \epsilon - \frac{(1-3\epsilon)(1-\epsilon)}{3(1+3\epsilon)}$. Observe that since one edge going out of A_j yields at most one neighbor, at most $(1+\epsilon)pd|A_j| - \frac{(1-3\epsilon)(1-\epsilon)}{1+3\epsilon} pd|A_j|/3 = \beta^* pd|A_j|$ neighbours of vertices of A_j

are outside B_{j+1}. Suppose that $\epsilon \leq 1/17$. We then have that B_{j+1} contains at least $((1-\epsilon)\gamma(1/5)-\beta^*)pd|A_j| > 0.011pd|A_j|$ neighbours of $|A_j|$. Since $0.011pd|A_j| = 0.011(c\ln n)/\epsilon^2 > (\ln n)|A_j|$ when the constant c is sufficiently large, we get $|B_{j+1}| \geq (\ln n)|A_j| = (\ln n)|B_j|$ (this corresponds to case 3 above).

If $|A_j| \geq 1/(5p)$, using a simple averaging argument one can find $A' \subseteq A_j$ such that $|A'| = \lfloor 1/(5p) \rfloor$ and at least $\frac{(1-3\epsilon)(1-\epsilon)}{1+3\epsilon}pd|A'|/3$ edges going out of A' go to B_{j+1}, which implies by the same argument that $|B_{j+1}| \geq 0.011pd|A'| \geq d/500$ (this corresponds to case 2 above).

Case (B): Suppose that E_M contains at least $(1-\epsilon)pd|B_j|/3$ edges. Then by the same argument as in the previous paragraph we have that $|B_j| \geq (\ln n)|A_j|$ if $|A_j| \leq 1/(5p)$. This is impossible when $\ln n > 1$ since $|A_j| = |B_j|$. Hence, $|B_j| \geq d/500$ by same argument as above (this corresponds to case 1 above).

Case (C): Suppose that E_R contains at least $(1-\epsilon)pd|B_j|/3$ edges. By the same argument as above we have that either $|B_{j-1}| \geq d/500$ (this corresponds to case 5 above) or $B_{j-1} \geq (\ln n)|B_j|$ (this corresponds to case 4 above).

This completes the proof. ∎

We are now ready to prove the main result of this section.

Theorem 5.9. *Let the graph $G'' \cup G'''$ be obtained from G using steps* ***S1*** *and* ***S2*** *in the algorithm of section 4. Then step* ***S3*** *takes $O\left(\frac{n^2\ln^2 n}{d\ln\ln n}\right)$ time whp, giving an overall run time of $O\left(\frac{n^2\ln^3 n}{d}\right)$ for the entire algorithm whp.*

Proof. Our first step is to note that the set of edges of G included in the graph $G'' \cup G'''$ satisfies the preconditions of Lemma 5.7. Indeed, this is because G''' is a uniform sample of edges of G with probability $(c_2 \ln n)/d$ (see definition of step **S2**), and hence the preconditions are satisified as long as the constant c_2 is chosen sufficiently large.

We now bound the time taken by the augmentation process in step **S3**. We analyze the runtime in two stages: (1) finding a matching of size n–$2n/d$, and (2) extending the matching of size $n-2n/d$ to a perfect matching. By Lemma 5.7 the maximum number of layers in an alternating level graph in $G'' \cup G'''$, and hence the length of the shortest augmenting path, is $O\left(\frac{n\ln d}{d\ln\ln n}\right)$. As each augmentation phase takes time proportional to the number of edges in the graph, this implies that the first stage takes $O\left(\frac{n^2\ln^2 n}{d\ln\ln n}\right)$.

Finally, note that each augmentation phase increases the size of the matching by at least 1, and thus $O(n/d)$ augmentation suffice to extend the matching constructed in the first stage to a perfect matching. This takes

$O\left(\frac{n^2\ln n}{d}\right)$ time, so the runtime is $O\left(\frac{n^2\ln^2 n}{d\ln\ln n}\right)$ for step **S3**, and $O\left(\frac{n^2\ln^3 n}{d}\right)$ overall. ∎

Remark 5.2. Theorem 5.9 as well as Lemma 5.7 can be slightly altered to show that the runtime of the Hopcroft-Karp algorithm on the subsampled graph from [8] is $O\left(\frac{n^3\ln^2 n}{d^2\ln\ln n}\right)$. This shows that the approach in [8] yields an $\tilde{O}(n^{5/3})$ algorithm, which is better than $O(n^{1.75})$ stated in [8].

Theorem 5.10. *For any function $d(n)\geq 2\sqrt{n}$ there exists an infinite family of $d(n)$-regular graphs with $2n+o(n)$ vertices such that whp the algorithm in section 4 performs $\Omega(n/d)$ augmentations in the worst case.*

Proof. In what follows we omit the dependence of d on n for brevity. Define $H^{(k)}=(U,V,E)$, $0\leq k\leq d$, to be a $(d-k)$-regular bipartite graph with $|U|=|V|=d$. The graph G consists of t copies of $H^{(k)}$, which we denote by $\{H_j\}_{j=1}^t$, where $H_j=H^{(t-j+1)}$, and $2t$ vertices $u_1,\dots,u_t$ and $v_1,\dots,v_t$. Each of $u_1,\dots,u_t$ is connected to all d vertices in the V-part of H_1, and for $1\leq j\leq t$, the vertex v_j is connected to all vertices in the U-part of H_j. The remaining connections are established by adding $t-j$ edge-disjoint perfect matchings between the U part of H_j and the V part of H_{j+1} for all $1\leq j<t$. Set $t=n/d\leq\sqrt{n}/2\leq d/4$. Note that the strength of edges in H_j is at least $d/4$, so whp there exists a perfect matching in subgraph of H_j generated by the sampling steps **S1** and **S2**, for $1\leq j\leq t$. Suppose that at the first iteration of the Hopcroft-Karp algorithm a perfect matching is found in each H_j, thus leaving unmatched the vertices $u_1,\dots,u_t$ and $v_1,\dots,v_t$. Then from this point on, the shortest augmenting path for each pair (u_j,v_j) has length j, and each augmentation phase of the Hopcroft-Karp algorithm will increase the size of the matching by 1. Hence, it takes t augmentations to find a perfect matching. The number of vertices is $2(d+1)t=2n+o(n)$. ∎

6. Perfect matchings in doubly stochastic matrices

An $n\times n$ matrix A is said to be *doubly stochastic* if every element is non-negative, and every row-sum and every column-sum is 1. The celebrated Birkhoff-von Neumann theorem says that every doubly stochastic matrix is a convex combination of permutation matrices (*i.e.*, matchings). Surprisingly, the running time of computing this convex combination (known as a Birkhoff-von Neumann decomposition) is typically reported as $O(m^2\sqrt{n})$, even though much better algorithms can be easily obtained using existing techniques or very simple modifications. We list these running times here since there does not seem to be any published record[5]. After listing the running times that can be obtained using existing techniques, we will show how proportionate uncrossings can be applied to this problem to obtain a slight improvement.

1. An $O(m^2)$-time algorithm for finding a Birkhoff-von Neumann decomposition can be obtained by finding a perfect matching in the existing graph using augmenting paths (in time $O(mn)$), assigning this matching a weight which is the weight of the smallest edge in the matching, subtracting this weight from every edge in the matching (causing one or more edges to be removed from the support of A), and continuing the augmenting path algorithm without restarting. When a matching is found, if we remove k edges, then we need to find only k augmenting paths (finding each augmenting path takes time $O(m)$) to find another matching, which leads to a total time of $O(m^2)$.
2. Let b be the maximum number of significant bits in any entry of A. An $O(mb)$-time algorithm for finding a single perfect matching in the support of a doubly stochastic matrix can be easily obtained using the technique of Gabow and Kariv [7]: repeatedly find Euler tours in edges where the lowest order bit (say bit j) is 1, and then increase the weight of all edges going from left to right by 2^{-j} and decrease the weight of all edges going from right to left by the same amount, where the directionality of edges corresponds to an arbitrary orientation of the Euler tour; this eliminates bit j while preserving the doubly stochastic property and without increasing the support.
3. An $O(mnb)$-time algorithm to compute the Birkhoff-von Neumann decomposition can be obtained using the edge coloring algorithm of Gabow and Kariv [7].

We now show how our techniques lead to an $O(m\ln^3 n + n^{1.5}\ln n)$-time algorithm for finding a single perfect matching in the support of a doubly stochastic matrix. In realistic scenarios, this is unlikely to be better than (2) above, and we present this primarily to illustrate another application of our proportionate uncrossing technique. First, define a weighted bipartite graph $G = (P, Q, E)$, where $P = \{u_1, u_2, \ldots, u_n\}$ corresponds to rows of A, $Q = \{v_1, v_2, \ldots, v_n\}$ corresponds to columns of A, and $(u_i, v_j) \in E$ iff $A_{i,j} > 0$. Define a weight function w on edges, with $w(u_i, v_j) = A_{i,j}$. Let $\mathcal{R}$ be the collection of all pairs $(A, B), A \subseteq P, B \subseteq Q, |P| > |Q|$. Since A is doubly stochastic, the collection $\mathcal{R}$ is (1/2)-thick with respect to (G, w, E). Let $\mathcal{T}$ be a (1/2)-uncrossing of $\mathcal{R}$. Performing a Benczúr-Karger sampling on G

[5] This list was compiled by Bhattacharjee and Goel and is presented here to provide some context rather than as original work.

will guarantee (with high probability) that at least one edge is sampled from every witness set in $W(\mathcal{T})$, and hence running the Hopcroft-Karp algorithm on the sampled graph will yield a perfect matching with high probability. The running time of $O(m \ln^3 n + n^{1.5} \ln n)$ is just the sum of the running times of Benczúr-Karger sampling for weighted graphs [2] and the Hopcroft-Karp matching algorithm [9].

Acknowledgments. We would like to thank Rajat Bhattacharjee for many useful discussions on precursors to this work, and Michel Goemans for suggesting an alternate proof mentioned in Remark 3.1. We would also like to thank the anonymous referees for detailed comments that helped improve the presentation of the paper.

References

[1] G. Aggarwal, R. Motwani, D. Shah, and A. Zhu: Switch scheduling via randomized edge coloring, *FOCS,* 2003.

[2] A. A. Benczúr, and D. R. Karger: Approximating s-t minimum cuts in $\tilde{O}(n^2)$ time, *Proceedings of the 28th annual ACM symposium on Theory of computing,* 1996.

[3] G. Birkhoff: Tres observaciones sobre el algebra lineal, *Univ. Nac. Tucumán Rev. Ser. A* **5** (1946), 147–151.

[4] B. Bollobás: *Modern graph theory,* Springer, 1998.

[5] R. Cole and J. E. Hopcroft: On edge coloring bipartite graphs. *SIAM J. Comput.* **11** (1982), 540–546.

[6] R. Cole, K. Ost, and S. Schirra: Edge-coloring bipartite multigraphs in $O(E \log D)$ time, *Combinatorica* **21** (2001), 5–12.

[7] H. N. Gabow and O. Kariv: Algorithms for edge coloring bipartite graphs and multigraphs, *SIAM J. Comput.* **11** (1982), 117–129.

[8] A. Goel, M. Kapralov, and S. Khanna: Perfect matchings via uniform sampling in regular bipartite graphs, *Proceedings of the Nineteenth Annual ACM -SIAM Symposium on Discrete Algorithms,* 2009.

[9] J. E. Hopcroft and R. M. Karp: An $n^{\frac{5}{2}}$ algorithm for maximum matchings in bipartite graphs, *SIAM J. Comput.* **2** (1973), 225–231.

[10] D. Karger: Random sampling in cut, flow, and network design problems, *Mathematics of Operations Research (Preliminary version appeared in the Proceedings of the 26th annual ACM symposium on Theory of computing)* **24** (1999), 383–413.

[11] D. Karger and M. Levine: Random sampling in residual graphs, *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing,* 2002.

[12] D. R. Karger and C. Stein: A new approach to the minimum cut problem, *J. ACM* **43** (1996), 601–640.

[13] D. König: Uber graphen und ihre anwendung auf determinententheorie und mengenlehre, *Math. Annalen* **77** (1916), 453–465.

[14] R. Motwani: Average-case analysis of algorithms for matchings and related problems, *Journal of the ACM(JACM)* **41** (1994), 1329–1356.

[15] R. Motwani and P. Raghavan: *Randomized Algorithms,* Cambridge University Press, 1995.

[16] A. Schrijver: Bipartite edge coloring in $O(\Delta m)$ time, *SIAM J. on Comput.* **28** (1999), 841–846.

[17] J. von Neumann: A certain zero-sum two-person game equivalent to the optimal assignment problem, *Contributions to the optimal assignment problem to the Theory of Games* **2** (1953), 5–12.

A. Proof of Lemma 2.3

Consider any (A,B) where $|A|>|B|, A\subseteq P, B\subseteq Q$. Define $A_i=P_i\cap A$ and $B_i=Q_i\cap B$. Fix an i such that $|A_i|>|B_i|$; such an i is guaranteed to exist. By the definition of relevance, there exists a pair $(X,Y)\in\mathcal{R}$ such that $X\subseteq A_i$, and $W(X,Y)\cap E_R\subseteq W(A_i,B_i)\cap E_R$. By the assumption in the theorem, there exists an edge $(u,v)\in E^*\cap E_R\cap W(X,Y)$. Since $W(X,Y)\cap E_R\subseteq W(A_i,B_i)\cap E_R$, it follows that $(u,v)\in E^*\cap E_R\cap W(A_i,B_i)$. This edge is in G^*, and goes from A_i to $Q_i\setminus B_i$, *i.e.*, from A_i to $Q_i\setminus(Q_i\cap B)$, and hence, from A to $Q\setminus B$. Since the only assumption on (A,B) was that $|A|>|B|$, we can now invoke Hall's theorem to claim that G^* has a perfect matching. ∎

B. Proof of Theorem 2.6

As mentioned before, the proof is along very similar lines to that of the Benczúr-Karger sampling theorem, but does not follow in a black-box fashion and is presented here for completeness. The proof relies on the following result due to Karger and Stein [12]:

Lemma B.1. *Let $H(V,E)$ be an undirected graph on n vertices such that each edge e has an associated non-negative weight $\tilde{p}_e$. Let s^* be the value of minimum cut in H under the weight function $\tilde{p}_e$. Then for any $\alpha\geq 1$, the number of cuts in H of weight at most αs^* is less than $n^{2\alpha}$.*

Proof of Theorem 2.6. We will choose $c=5$. The first part of the proof shows that it is sufficient to bound a certain expression that involves only cuts. The second part then bounds this expression.

For the first part, let $\mu(X)=\sum_{e\in X}p_e$ denote the expected number of edges chosen from X by the sampling process. If a set $X\in\mathcal{X}$ contains an edge e with $p_e=1$, then that edge will definitely be chosen, and that set does not contribute to

$$\sum_{X\in\mathcal{X}}\Pr[\text{No edge in } X \text{ is chosen in } H']$$

and can be removed from $\mathcal{X}$. Hence, assume without loss of generality that $p_e < 1$ for every edge in $\bigcup_{X\in\mathcal{X}} X$. Define $\tilde{\mu}(X) = \sum_{e\in X}\left(\frac{c\ln n}{s_e}\right)$. Now for any set $X \in \mathcal{X}$,

$$\Pr[\text{No edge in } X \text{ is chosen in } H'] = \prod_{e\in X}(1-p_e) \le \prod_{e\in X} e^{-p_e} \le e^{-\mu(X)},$$

where

$$\mu(X) = \frac{c\ln n}{\gamma}\sum_{e\in X}\frac{1}{s_e} > (c\ln n)\sum_{e\in f(X)}\frac{1}{s_e} = \tilde{\mu}(f(X)).$$

Since f is a one-one function, it is sufficient to provide an upper-bound on $\sum_{C\in\mathcal{C}} e^{-\tilde{\mu}(C)}$.

For the second part, let $\tilde{\mu}_1, \tilde{\mu}_2, \ldots, \tilde{\mu}_{2^n-2}$ be a non-decreasing sorted sequence corresponding to the multi-set $\{\tilde{\mu}(C) \colon C \in \mathcal{C}\}$. Define $q_i = e^{-\tilde{\mu}_i}$. Consider an arbitrary cut C. Any edge in C can have strength at most $|C|$, and hence $\tilde{\mu}(C) \ge c\ln n$, and therefore, $q_1 \le n^{-c}$. So the sum of q_i for the first n^2 cuts in the sequence is bounded by n^{-c+2}. We now focus on the remaining cuts. By Lemma B.1, we know that for any $\alpha \ge 1$, we have $\tilde{\mu}_{n^{2\alpha}} \ge \alpha\tilde{\mu}_1$. Hence

$$\tilde{\mu}_k \ge \frac{\ln k}{2\ln n}\tilde{\mu}_1,$$

which in turn implies that $q_k \le k^{-c/2}$. Thus

$$\sum_{X\in\mathcal{X}} \Pr[\text{No edge in } X \text{ is chosen in } H']$$

$$\le \sum_{C\in\mathcal{C}} e^{-\tilde{\mu}(C)} \le \sum_{k=1}^{n^2} q_k + \sum_{k>n^2} q_k \le n^{-c+2} + \sum_{k>n^2} k^{-c/2} = O(n^{-c+2}),$$

giving us the desired result when we choose $c = 5$. ∎

C. Proof of Lemma 2.8

Assume by way of contradiction that no such integer j exists for some pair of multisets S_1 and S_2. Let K be the largest integer in $S_1 \cup S_2$, and let α_i and β_i denote the number of occurrences of i in the multisets S_1 and S_2 respectively. Then for all $j \ge 1$, we have

$$\sum_{i=j}^{K}\frac{\alpha_i}{i} \le \gamma\left(\sum_{i=j}^{K}\frac{\beta_i}{i}\right).$$

Summing the above inequality for all $j \in \{1..K\}$, we get

$$\sum_{i=1}^{K} \alpha_i \leq \gamma \left(\sum_{i=1}^{K} \beta_i \right),$$

which is a contradiction since $|S_1| > \gamma |S_2|$ by assumption. ∎

Ashish Goel

Departments of Management Science
and Engineering and (by courtesy)
Computer Science
Stanford University
ashishg@stanford.edu

Michael Kapralov

Institute for Computational
and Mathematical Engineering
Stanford University
kapralov@stanford.edu

Sanjeev Khanna

Department of Computer and Information Science
University of Pennsylvania
Philadelphia PA
sanjeev@cis.upenn.edu