

Perfect Matchings via Uniform Sampling in Regular Bipartite Graphs

ASHISH GOEL AND MICHAEL KAPRALOV

Stanford University

AND

SANJEEV KHANNA

University of Pennsylvania

Abstract. In this article we further investigate the well-studied problem of finding a perfect matching in a regular bipartite graph. The first nontrivial algorithm, with running time $O(mn)$, dates back to König's work in 1916 (here $m = nd$ is the number of edges in the graph, $2n$ is the number of vertices, and d is the degree of each node). The currently most efficient algorithm takes time $O(m)$, and is due to Cole et al. [2001]. We improve this running time to $O(\min\{m, \frac{n^{2.5} \ln n}{d}\})$; this minimum can never be larger than $O(n^{1.75} \sqrt{\ln n})$. We obtain this improvement by proving a uniform sampling theorem: if we sample each edge in a d -regular bipartite graph independently with a probability $p = O(\frac{n \ln n}{d^2})$ then the resulting graph has a perfect matching with high probability. The proof involves a decomposition of the graph into pieces which are guaranteed to have many perfect matchings but do not have any small cuts. We then establish a correspondence between potential witnesses to nonexistence of a matching (after sampling) in any piece and cuts of comparable size in that same piece. Karger's sampling theorem [1994a, 1994b] for preserving cuts in a graph can now be adapted to prove our uniform sampling theorem for preserving perfect matchings. Using the $O(m\sqrt{n})$ algorithm (due to Hopcroft and Karp [1973]) for finding maximum matchings in bipartite graphs on the sampled graph then yields the stated running time. We also provide an infinite family of instances to show that our uniform sampling result is tight up to polylogarithmic factors (in fact, up to $\ln^2 n$).

27

The research of A. Goel is supported NSF ITR grant 0428868, NSF CAREER award 0339262, and a grant from the Stanford-KAUST alliance for academic excellence. The research of M. Kapralov is supported by a Stanford Graduate Fellowship. The research of S. Khanna, is supported in part by a Guggenheim Fellowship, an IBM Faculty Award, and by NSF Award CCF-0635084.

Author's addresses: A. Goel, Departments of Management Science and Engineering and (by courtesy) Computer Science, Stanford University, Palo Alto, CA, e-mail: ashishg@stanford.edu; M. Kapralov, Institute for Computational and Mathematical Engineering, Stanford University, Palo Alto, CA, e-mail: kapralov@stanford.edu; S. Khanna, Department of Computer and Information Science, University of Pennsylvania, Philadelphia PA, e-mail: sanjeev@cis.upenn.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2010 ACM 1549-6325/2010/03-ART27 \$10.00

DOI 10.1145/1721837.1721843 <http://doi.acm.org/10.1145/1721837.1721843>

Categories and Subject Descriptors: G.2.2 [Discrete Mathematics]: Graph Theory—*Graph algorithms*; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*Computations on discrete structures*

General Terms: Algorithms

Additional Key Words and Phrases: Perfect matching, regular bipartite graphs

ACM Reference Format:

Goel, A., Kapralov, M., and Khanna, S. 2010. Perfect matchings via uniform sampling in regular bipartite graphs. *ACM Trans. Algor.* 6, 2, Article 27 (March 2010), 13 pages.
DOI = 10.1145/1721837.1721843 <http://doi.acm.org/10.1145/1721837.1721843>

1. Introduction

A bipartite graph $G = (U, V, E)$ with vertex set $U \cup V$ and edge set $E \subseteq U \times V$ 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 U (as a consequence of regularity, U and V have the same size). Regular bipartite graphs have been the subject of much study. Random regular bipartite graphs represent some of the simplest examples of expander graphs [Motwani and Raghavan 1995]. These graphs are also used to model scheduling, routing in switch fabrics, and task-assignment problems (sometimes via edge coloring, as described shortly) [Aggarwal et al. 2003; Cole et al. 2001].

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 [Bollobás 1998]. Finding a matching in a regular bipartite graph is a well-studied problem, starting with the algorithm of König in 1916, which is now known to run in time $O(mn)$ [König 1916]. The well-known bipartite matching algorithm of Hopcroft and Karp [1973] can be used to obtain a running time of $O(m\sqrt{n})$. An algorithm of complexity $O(n^\omega)$, where ω is the matrix multiplication constant, was given by Mucha and Sankowski [2004]. In graphs where d is a power of 2, the following simple idea, due to Gabow and Kariv [1982], 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 (e.g., by Cole and Hopcroft [1982], and then by Schrijver [1999] to $O(md)$), Cole et al. [2001] gave an $O(m)$ algorithm for the case of general d .

The main interest of Cole et al. [2001] was in edge coloring of general bipartite graphs of maximum degree d , where finding perfect matchings in regular bipartite graphs is an important subroutine. Finding perfect matchings in regular bipartite graphs is also closely related to the problem of finding a Birkhoff von Neumann decomposition of a doubly stochastic matrix [Birkhoff 1946; von Neumann 1953].

In this article we present an algorithm for finding a perfect matching in a regular bipartite graph that runs in time $O(\min\{m, \frac{n^{2.5} \ln n}{d}\})$. It is easy to see that this minimum can never be larger than $O(n^{1.75} \sqrt{\ln n})$. This is a significant improvement over the running time of Cole et al. [2001] when the bipartite graph is relatively dense. We first prove (Theorem 2.1 in Section 2) that if we sample the edges of a

regular bipartite graph independently and uniformly at rate $p = O(\frac{n \ln n}{d^2})$, then the resulting graph has a perfect matching with high probability. The resulting graph has $O(mp)$ edges in expectation, and running the bipartite matching algorithm of Hopcroft and Karp [1973] gives an expected running time of $O(\frac{n^{2.5} \ln n}{d})$. Since we know this running time in advance, we can choose the better of m and $\frac{n^{2.5} \ln n}{d}$ in advance. It is worth noting that uniform sampling can easily be implemented in $O(1)$ time per sampled edge assuming that the data is given in adjacency list format, with each list stored in an array, and assuming that $\log n$ bit random numbers can be generated in one time-step.¹

We believe that our sampling result is also independently interesting as a combinatorial fact. The proof of our sampling theorem relies on a sequential decomposition procedure that creates a vertex-disjoint collection of subgraphs, each subgraph containing many perfect matchings on its underlying vertex set. We then show that if we uniformly sample edges in each decomposed subgraph at a suitably chosen rate, with high probability at least one perfect matching survives in each decomposed subgraph. This is established by using Karger's sampling theorem [Karger 1994a, 1994b] in each subgraph. An effective use of Karger's sampling theorem requires the min-cuts to be large, a property that is not necessarily true in the original graph. For instance, G could be a union of two disjoint d -regular bipartite graphs, in which case the min-cut is 0; nonpathological examples are also easy to obtain. However, our serial decomposition procedure ensures that the min-cuts are large in each decomposed subgraph. We then establish a 1-1 correspondence between possible Hall's theorem counterexamples in each subgraph and cuts of comparable size in that subgraph. Since Karger's sampling theorem [1994a, 1994b] is based on counting cuts of a certain size, this coupling allows us to claim (with high probability) that no possible counterexample to Hall's theorem exists in the sampled graph. On a related note, Benczúr [1997] presented another sampling algorithm which generates $O(n \ln n)$ edges that approximate *all cuts*; however, this sampling algorithm, as well as recent improvements [Spielman and Teng 2004; Spielman and Srivastava 2008] take $\tilde{\Omega}(m)$ time to generate the sampled graph. Hence these approaches do not directly help in improving upon the already known $O(m)$ running time for finding perfect matchings in d -regular bipartite graphs.

The sampling rate we provide may seem counterintuitive; a superficial analogy with Karger's sampling theorem [1994a, 1994b] or Benczúr's work [1997] might suggest that sampling a total of $O(n \ln n)$ edges should suffice. We show (Theorem 4.1, Section 4) that this is not the case. In particular, we present a family of graphs where uniform sampling at rate $o(\frac{n}{d^2 \ln n})$ results in a vanishingly low probability that the sampled subgraph has a perfect matching. Thus, our sampling rate is tight up to factors of $O(\ln^2 n)$. This lower bound suggests two promising directions for further research: designing an efficiently implementable nonuniform sampling scheme, and designing an algorithm that runs faster than Hopcroft-Karp's algorithm for near-regular bipartite graphs (since the degree of each vertex in the sampled subgraph will be concentrated around the expectation).

¹ Even if we assume that only one random bit can be generated in one time-step, the running time of our algorithm remains unaltered since the Hopcroft-Karp algorithm incurs an overhead of \sqrt{n} per sampled edge, anyway.

2. Uniform Sampling for Perfect Matchings: An Upper Bound

In this section, we will establish our main sampling theorem stated shortly. We will then show in Section 3 that this theorem immediately yields an $O(n^{1.75}\sqrt{\ln n})$ -time algorithm for finding a perfect matching in regular bipartite graphs.

THEOREM 2.1. *There exists a constant c such that given a d -regular bipartite graph $G(U, V, E)$, a subgraph G' of G generated by sampling the edges in G uniformly at random with probability $p = \frac{cn \ln n}{d^2}$ contains a perfect matching with high probability.*

Our proof is based on a decomposition procedure that partitions the given graph into a vertex-disjoint collection of subgraphs such that (i) the minimum cut in each subgraph is large, and (ii) each subgraph contains $\Omega(d)$ perfect matchings on its vertices. We then show that for a suitable choice of sampling rate, with high probability at least one perfect matching survives in each subgraph. The union of these perfect matchings then gives us a perfect matching in the original graph. We emphasize here that the decomposition procedure is merely an artifact for our proof technique. Note that the theorem is trivially true when $d \leq \sqrt{n \log n}$. So in what follows we assume that $d > \sqrt{n \log n}$.

2.1. HALL'S THEOREM WITNESS SETS. Let $G(U, V, E)$ be a bipartite graph. We denote by $V(G)$ the vertex set of G . For any set $S \subseteq V(G)$, let $\delta_G(S)$ denote the set of edges crossing the boundary of S in G . Also, for any set $S \subseteq V(G)$, we denote by $\Gamma_G(S)$ the set of vertices that are adjacent to vertices in S .

A pair (A, B) with $A \subseteq U$ and $B \subseteq V$ is said to be a *relevant pair* to Hall's theorem if $|A| > |B|$. Given a relevant pair (A, B) , we denote by $E(A, B)$ the set of edges in $E \cap (A \times (V \setminus B))$. We refer to the set $E(A, B)$ as a *witness edge set* if (A, B) is a relevant pair. Also, for any two sets $A, A' \subseteq U$ we denote by $A \oplus A'$ the set $(A \setminus A') \cup (A' \setminus A)$. In what follows we will be using Hall's theorem, which we state here for convenience of the reader:

THEOREM 2.2 (HALL'S THEOREM, REFER TO BOLLOBÁS [1998]). *A bipartite graph $G(U, V, E)$ contains a matching that includes every vertex in U iff $|\Gamma_G(S)| \geq |S|$ for all $S \subset U$.*

Note that if $|U| = |V|$, then any matching that includes every vertex in U is also a perfect matching in G . Since $|U| = |V|$ in a d -regular graph, by Hall's theorem, to prove Theorem 2.1 it suffices to show that with high probability in the sampled graph G' , at least one edge is chosen from each witness set. We will focus on a subclass of relevant pairs, referred to as minimal relevant pairs. A relevant pair (A, B) is *minimal* if there does not exist another relevant pair (A', B') with $A' \subset A$ and $E(A', B') \subseteq E(A, B)$. A witness edge set corresponding to a minimal relevant pair is called a *minimal witness set*, respectively. If a graph G has a perfect matching, every minimal witness set must be nonempty. It also follows from Hall's theorem that any balanced subgraph of G that includes at least one edge from every minimal witness set must have a perfect matching. We refer to a bipartite graph as balanced if it contains the same number of vertices in each side.

A key idea underlying our proof is a mapping from minimal witness sets in G to *distinct* cuts in G . In particular, we will map each minimal witness set $E(A, B)$ to the cut $\delta_G(A \cup B)$. The next theorem shows that this is a one-to-one mapping.

THEOREM 2.3. *Let $G(U, V, E)$ be a bipartite graph that has at least one perfect matching. If (A, B) and (A', B') are minimal relevant pairs in G with $E(A, B) \neq E(A', B')$, then $\delta_G(A \cup B) \neq \delta_G(A' \cup B')$.*

PROOF. Assume by way of contradiction that there exist minimal relevant pairs (A, B) and (A', B') in G with $E(A, B) \neq E(A', B')$ but $\delta_G(A \cup B) = \delta_G(A' \cup B')$. Then the following conditions must be satisfied for any edge $(u, v) \in E$.

- (A1) If $u \in A \oplus A'$ then $v \in B \oplus B'$. To see this, assume without loss of generality that $u \in A \setminus A'$, and then note that if $v \in B \cap B'$, then $(u, v) \in \delta_G(A' \cup B')$ but $(u, v) \notin \delta_G(A \cup B)$, which is a contradiction. Similarly, if $v \in V \setminus (B \cup B')$, then $(u, v) \in \delta_G(A \cup B)$ but $(u, v) \notin \delta_G(A' \cup B')$, which is again a contradiction.
- (A2) If $u \in (A \cap A')$ then $v \notin B \oplus B'$. To see this, without loss of generality, assume that $v \in (B \setminus B')$. Then $(u, v) \in \delta_G(A' \cup B')$ but $(u, v) \notin \delta_G(A \cup B)$. This is a contradiction.

In what follows, we slightly abuse the notation and given any (not necessarily relevant) pair (C, D) with $C \subseteq U$ and $D \subseteq V$, we denote by $E(C, D)$ the set of edges in $E \cap (C \times (V \setminus D))$. As an immediate corollary of the properties A1 and A2, we now obtain the following containment results.

- (B1) $E(A \setminus A', B \setminus B') \subseteq E(A, B)$. This follows directly from property A1 given before.
- (B2) $E(A \cap A', B \cap B') \subseteq E(A, B)$. This follows directly from property A2 given before.

We now consider three possible cases based on the relationship between A and A' , and establish a contradiction for each case.

Case 1. $A \cap A' = \emptyset$. By property A1, if $u \in A \cup A'$ then $v \in B \cup B'$. In other words, there are no edges from $A \cup A'$ to vertices outside $B \cup B'$. Since $|A \cup A'| = |A| + |A'| > |B| + |B'|$, this contradicts our assumption that G has at least one perfect matching.

Case 2. $A = A'$. For any edge (u, v) with $u \in A$, property A2 shows that $v \notin B \oplus B'$. Then $E(A, B) = E(A', B')$. A contradiction.

Case 3. $A \cap A' \neq \emptyset$ and $A \neq A'$. Assume without loss of generality that $A \setminus A' \neq \emptyset$. Since $|A| > |B|$, it must be that either $|A \setminus A'| > |B \setminus B'|$ or $|A \cap A'| > |B \cap B'|$. If $|A \setminus A'| > |B \setminus B'|$, then $(A \setminus A', B \setminus B')$ is a relevant pair, and by B1, it contradicts the fact that (A, B) is a minimal relevant pair. If $|A \cap A'| > |B \cap B'|$, then $(A \cap A', B \cap B')$ is a relevant pair set, and by B2, it contradicts the fact that (A, B) is a minimal relevant pair. \square

2.2. A DECOMPOSITION PROCEDURE. Given a d -regular bipartite graph on n vertices, we will first show that it can be partitioned into $k = O(n/d)$ vertex-disjoint graphs $G_1(U_1, V_1, E_1), G_2(U_2, V_2, E_2), \dots, G_k(U_k, V_k, E_k)$ such that each graph G_i satisfies the following properties.

- (1) The size of a minimum cut in $G_i(U_i, V_i, E_i)$ is strictly greater than $\alpha = \frac{d^2}{4n}$.
- (2) $|\delta_G(U_i \cup V_i)| \leq d/2$ (hence G_i contains at least $d/2$ edge-disjoint perfect matchings).

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

- (1) Find a smallest nonempty proper subset $X_i \subseteq V(H_i)$ such that $|\delta_{H_i}(X_i)| \leq 2\alpha$. Let M_i denote the number of edges in the cut $\delta_{H_i}(X_i)$. If no such set X_i exists, we define G_i to be the graph H_i , and terminate the decomposition procedure.
- (2) Define G_i to be the subgraph of H_i induced by the vertices in X_i , that is, $X_i = U_i \cup V_i = V(G_i)$. 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).

Note that if the minimum cut of G is greater than 2α , then the procedure terminates after the first step, and the decomposition trivially satisfies both P1 and P2. So we focus shortly on the case when step (2) is executed at least once.

We now prove the following properties of the decomposition procedure.

THEOREM 2.4. *The decomposition procedure outlined earlier satisfies properties P1 and P2.*

PROOF. We start by proving that property P1 is satisfied. Suppose that there exists a cut $(C, V(G_i) \setminus C)$ in G_i of value at most α , that is, $|\delta_{G_i}(C)| \leq \alpha$ (note that one could have $C \cap U \neq \emptyset$ and $C \cap V \neq \emptyset$). Let $D = V(G_i) \setminus C$. We have $|\delta_{H_i}(C) \setminus \delta_{G_i}(C)| + |\delta_{H_i}(D) \setminus \delta_{G_i}(D)| \leq 2\alpha$ by the choice of X_i in (1). Suppose without loss of generality that $|\delta_{H_i}(C) \setminus \delta_{G_i}(C)| \leq \alpha$. Then $|\delta_{H_i}(C)| \leq 2\alpha$ and $C \subset X_i$, which contradicts the choice of X_i as the smallest cut of value at most 2α in step (1) of the procedure.

It remains to show that $|\delta_G(U_i \cup V_i)| \leq d/2$ for all i . In order to establish this property, it suffices to show that $\sum_{i=1}^k M_i \leq d/2$ (recall that $M_i = |\delta_{H_i}(X_i)|$).

We prove the following statements by induction on k , the number of times step (2) in the preceding decomposition procedure has been executed thus far.

- (1) $|V(G_k)| = |U_k \cup V_k| \geq 2d$;
- (2) $\sum_{i=1}^k M_i \leq d/2$;
- (3) $k + 1 \leq n/d$.

Base: $k = 1$. Since $2\alpha = \frac{d^2}{2n} \leq d/2$, we have $M_1 \leq d/2$, which establishes (2). We now prove (1), that is, we show that $G_1(U_1, V_1, E_1)$ has at least $2d$ vertices. Consider any vertex $u \in U_1$. Let $\Gamma_{G_1}(u) \subseteq V_1$ be the neighbors of u in G_1 . Clearly,

$$|\delta_G(u) \cap \delta_G(X_1)| + \sum_{v \in \Gamma_{G_1}(u)} |\delta_G(v) \cap \delta_G(X_1)| \leq |\delta_G(X_1)| \leq 2\alpha \leq d/2.$$

If all terms are positive then we have

$$d/2 \geq |\delta_G(u) \cap \delta_G(X_1)| + |\Gamma_{G_1}(u)|.$$

This is a contradiction since the right-hand side is d , the number of neighbors of u . So we have $|\delta_G(v) \cap \delta_G(X_1)| = 0$ for some $v \in \Gamma_{G_1}(u)$, implying that all neighbors of v are inside U_1 , so $|U_1| \geq d$. A similar argument shows that $|V_1| \geq d$, so $|X_1| \geq 2d$. By the same argument, $|V(H_2)| \geq 2d$, which establishes (3).

Inductive step: $k - 1 \rightarrow k$. Suppose that the algorithm constructs G_k . Since $k \leq n/d$ by the inductive hypothesis, we have $\sum_{i=1}^k M_i \leq (n/d)(2\alpha) = \frac{n}{d}(\frac{d^2}{2n}) \leq d/2$, which establishes (2). Consider the cut $(X_k, V(H_k) \setminus X_k)$ of H_k .

Every edge in $\delta_G(X_k)$ has one endpoint in X_k and the other in either $V(H_k) \setminus X_k$ or $V(G_i) = X_i$ for some $i < k$. Thus

$$\delta_G(X_k) \subseteq \delta_{H_k}(X_k) \cup \bigcup_{i=1}^{k-1} \delta_{H_i}(X_i).$$

Thus

$$|\delta_G(X_k)| \leq M_k + \sum_{i=1}^{k-1} M_i.$$

By induction $k \leq n/d$, so we have

$$|\delta_G(X_k)| \leq \sum_{i=1}^k M_i \leq (n/d)(2\alpha) \leq (n/d) \frac{d^2}{2n} \leq d/2.$$

An argument similar to the base case can be used to show that $|X_k| \geq 2d$ as well as $|V(H_k) \setminus X_k| \geq 2d$, establishing (1). Since at every decomposition step $j \leq k$ at least $2d$ vertices were removed from the graph, we have $k + 1 \leq n/d$, which establishes (3). \square

2.3. PROOF OF THEOREM 2.1. We now argue that if the graph G' is obtained by uniformly sampling the edges of G with probability $p = \Theta(\frac{\ln n}{\alpha})$, then with high probability G' contains a perfect matching.

It suffices to show that in each graph G_i obtained in the decomposition procedure, every minimal witness set is hit with high probability in the sampled graph (that is, at least one edge in each minimal witness set is chosen in the sampled graph). This ensures that at least one perfect matching survives inside each G_i . A union of these perfect matchings then gives us a perfect matching of G in the sampled graph G' .

Fix a graph $G_i(U_i, V_i, E_i)$. Let (A, B) be a relevant pair in G_i . Using the fact that our starting graph G is d -regular, we observe that $|E(A, B)| \geq d + |E(B, A)|$, and obtain

$$|\delta_G(A \cup B)| \leq 2|E(A, B)| - d.$$

Let m_A, m_B denote the number of edges in G that connect nodes in A and B , respectively, to nodes outside G_i . Then

$$|\delta_{G_i}(A \cup B)| \leq 2|E(A, B)| - d - m_A - m_B.$$

By property P2, since $|\delta_G(U_i \cup V_i)| \leq d/2$, it follows that $|E(A, B) \cap E_i| \geq |E(A, B)| - d/2$. Also, by definition, $|E(A, B) \cap E_i| \geq |E(A, B)| - m_A - m_B$. Combining, we obtain

$$|\delta_{G_i}(A \cup B)| \leq 2|E(A, B) \cap E_i| - d/2.$$

Thus the set $E(A, B) \cap E_i$ contains at least half as many edges as the the cut $\delta_{G_i}(A \cup B)$. We will now use the following sampling result due to Karger [1994b].

THEOREM 2.5. (KARGER 1994b). *Let G_i be an undirected graph on at most n vertices, and let κ be the size of a minimum cut in G_i . There exists a positive constant c such that for any $\epsilon \in (0, 1)$, if we sample the edges in G_i uniformly with*

probability at least $p = c(\frac{\ln n}{\kappa \epsilon^2})$, then every cut in G_i is preserved to within $(1 \pm \epsilon)$ of its expected value with probability at least $1 - 1/n^{\Omega(1)}$.

Thus the sampling probability, needed to ensure that all cuts are preserved close to their expected value, is inversely related to the size of a minimum cut in the graph. We now use the previous theorem to prove that at least one perfect matching survives in each graph G_i when edges are sampled with probability as specified in Theorem 2.1.

By Property P1, we know that the size of a minimum cut in G_i is at least $\alpha = d^2/4n$. Fix an $\epsilon \in (0, 1)$. The preceding theorem implies that if we sample edges in G_i with probability $p = \Theta(\frac{\ln n}{\alpha \epsilon^2})$, then for every relevant pair (A, B) , with high probability the sampled graph contains $(1 \pm \epsilon)p|\delta_{G_i}(A \cup B)| = \Omega(\ln n)$ edges from the set $\delta_{G_i}(A \cup B)$.

Note that the set $\delta_{G_i}(A \cup B)$ is not necessarily a Hall's theorem witness edge set. However, by Theorem 2.3, we know that for every left (right) minimal witness edge set $E(A, B) \cap E_i$, we can associate a distinct cut, namely $\delta_{G_i}(A \cup B)$, of size at most twice $|E(A, B) \cap E_i|$. We now show that this correspondence can be used to directly adapt Karger's proof of Theorem 2.5 to claim that every witness edge set in G_i is preserved to within $(1 \pm \epsilon)$ of its expected value. We remind the reader that the proof of Karger's theorem is based on an application of union bound over all cuts in the graph. In particular, it is shown that the number of cuts of size at most β times the minimum cut size is bounded by $n^{2\beta}$. Then, for the sampling rate given in Theorem 2.5, Chernoff bounds are used to claim that the probability that a cut of size β times the minimum cut deviates by $(1 \pm \epsilon)$ from its expected value is at most $1/n^{\Omega(\beta)}$. The theorem follows by combining these two facts.

Within any piece of the decomposition, let c_i be the number of cuts of size i and let w_i be the number of minimal witness sets of size i . We know by the preceding correspondence argument that every Hall's theorem minimal witness set of size i corresponds to a cut of size at most $2i$, and at most one minimal witness set corresponds to the same cut.

Now, given a sampling probability p , the probability that none of the edges in some minimal witness set is sampled is at most $\sum_i w_i(1-p)^i$, which is at most $\sum_i c_i(1-p)^{i/2}$. Therefore the probability that there is no matching in this piece can be at most twice the expression used in Karger's theorem to bound the probability that there exists a cut from which no edge is sampled when the sampling rate is q , where $1-q = (1-p)^{1/2}$, or $p = 2q - q^2$. Hence, it is sufficient to use a sampling rate which is twice that required by Karger's sampling theorem to conclude that a perfect matching survives with probability at least $1 - 1/n^{\Omega(1)}$ in any given piece of the decomposition. The union bound over all pieces of decomposition can be handled by increasing the constant in the sampling probability.

Even though we don't use it in this article, the following remark is interesting and is worth making explicitly. The remark follows from the additional observation that Karger's proof [Karger 1994b] of Theorem 2.5 uses Chernoff bounds for each cut, and these bounds remain the same if we use minimal witness sets which are at least half the size of the corresponding cuts, and then sample with twice the probability.

Remark 2.6. There exists a positive constant c' such that for any $\epsilon \in (0, 1)$, if we sample the edges in G uniformly with probability at least $p = c'(\frac{\ln n}{\alpha \epsilon^2})$, then every minimal witness set in every piece G_i is preserved to within $(1 \pm \epsilon)$ of its

expected value with probability at least $1 - 1/n^{\Omega(1)}$. Here $\alpha = d^2/(4n)$, as defined before.

Putting everything together, the sampled graph G' will have a perfect matching with high probability as long as we sample the edges with probability $p > \frac{c \ln n}{\alpha}$ for a sufficiently large constant c , thus completing the proof of Theorem 2.1. We have made no attempt to optimize the constants in this proof (an upper bound of $\frac{8 \ln n}{\alpha}$ follows from the previous reasoning). In fact, in an implementation, we can use geometrically increasing sampling rates until either the sampled graph has a perfect matching, or the sampling rate becomes so large that the expected running time of the Hopcroft and Karp [Hopcroft and Karp 1973] algorithm is $\Omega(m)$.

3. A Faster Algorithm for Perfect Matchings in Regular Bipartite Graphs

We now show that the sampling theorem from the preceding section can be used to obtain a faster randomized algorithm for finding perfect matchings in d -regular bipartite graphs.

THEOREM 3.1. *There exists an $O(\min\{m, \frac{n^{2.5} \ln n}{d}\})$ expected time algorithm for finding a perfect matching in a d -regular bipartite graph with $2n$ vertices and $m = nd$ edges.*

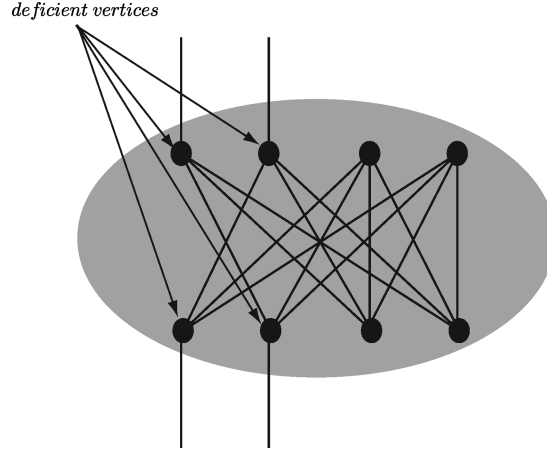
PROOF. Let G be a d -regular bipartite graph with $2n$ vertices and $m = nd$ edges. If $d \leq n^{3/4} \sqrt{\ln n}$, we use the $O(m)$ -time algorithm of Cole et al. [2001] for finding a perfect matching in a d -regular bipartite graph. It is easy to see that $m \leq \frac{n^{2.5} \ln n}{d}$ in this case.

Otherwise, we sample the edges in G at a rate of $p = \frac{cn \ln n}{d^2}$ for some suitably large constant c ($c = 32$ suffices by the reasoning from the previous section), and by Theorem 2.1, the sampled graph G' contains a perfect matching with high probability. The expected number of edges, say m' , in the sampled graph G' is $O(\frac{n^2 \ln n}{d})$. We can now use the algorithm of Hopcroft and Karp [Hopcroft and Karp 1973] to find a maximum matching in the bipartite graph G' in expected time $O(m' \sqrt{n})$. The sampling is then repeated if no perfect matching exists in G' . This takes $O(\frac{n^{2.5} \ln n}{d})$ expected running time. Hence, the algorithm takes $O(\min\{m, \frac{n^{2.5} \ln n}{d}\})$ expected time. \square

Note that by aborting the computation whenever the number of sampled edges is more than twice the expected value, the preceding algorithm can be easily converted to a Monte-Carlo algorithm with a worst-case running time of $O(\min\{m, \frac{n^{2.5} \ln n}{d}\})$ and a probability of success $1 - o(1)$. Finally, it is easy to verify that the stated running time never exceeds $O(n^{1.75} \sqrt{\ln n})$.

4. Uniform Sampling for Perfect Matchings: A Lower Bound

We now present a construction that shows that the uniform sampling rate of Theorem 2.1 is optimal to within a factor of $O(\ln^2 n)$. As before, G' denotes the graph obtained by sampling the edges of a graph G uniformly with probability p .

FIG. 1. Graph $H^{(k)}$ for $k = 2$ and $d = 4$.

THEOREM 4.1. *Let $d(n)$ be a nondecreasing positive-integer-valued function such that for some fixed integer n_0 , it always satisfies one of the following two conditions for all $n \geq n_0$: (a) $d(n) \leq \sqrt{n/\ln n}$, or (b) $\sqrt{n/\ln n} < d(n) \leq n/\ln n$. Then there exists a family of $d(n)$ -regular bipartite graphs G_n with $2n + o(n)$ vertices such that the probability that the graph G'_n , obtained by sampling edges of G_n with probability p , has a perfect matching goes to zero faster than any inverse polynomial function in n if $p = o(1)$ when $d(n)$ satisfies preceding condition (a), and if*

$$p = o\left(\frac{n}{(d(n))^2 \ln n}\right)$$

when $d(n)$ satisfies preceding condition (b).

PROOF. Note that the theorem asserts that essentially no sampling can be done when $d(n) \leq \sqrt{n/\ln n}$. We shall omit the dependence on n in $d(n)$ to simplify notation.

Define $H^{(k)} = (U, V, E)$, $0 \leq k \leq d$, to be a bipartite graph with $|U| = |V| = d$ such that k vertices in each of U and V have degree $(d - 1)$ and the remaining vertices have degree d . We will call the vertices of degree $(d - 1)$ *deficient*. Clearly, for any $0 \leq k \leq d$, the graph $H^{(k)}$ exists: starting with a d -regular bipartite graph on $2d$ vertices, we can remove an arbitrary subset of k edges that belong to a perfect matching in the graph ($H^{(k)}$ with $k = 2$ and $d = 4$ is shown in Figure 1). In the following construction, we will use copies of $H^{(k)}$ as building blocks to create our final instance. In doing so, only the set of deficient vertices in a copy of $H^{(k)}$ will be connected to (deficient) vertices in other copies in our construction.

We now define a d -regular bipartite graph G_n . Let $\gamma = \lceil \frac{d^2 \ln n}{n} \rceil$ (note that $\gamma \leq d$ since $d \leq n/\ln n$). We choose $W = \lceil \frac{d}{\gamma} \rceil$, $k_j = \gamma$ for $1 \leq j < W$, and $k_W = d - \gamma(W - 1) \leq \gamma$. We also define $K(n) = \lceil \ln n \rceil$ if $d(n) \geq \sqrt{n/\ln n}$ and $K(n) = \lceil \frac{n}{d^2} \rceil$ otherwise.




FIG. 2. Illustration of the family of graphs that yields the lower bound.

The graph G_n consists of $K(n) \cdot W$ copies of $H^{(k)}$ that we index as $\{H_{i,j}\}_{1 \leq i \leq K(n), 1 \leq j \leq W}$. The subgraph $H_{i,j}$ is a copy of $H^{(k_j)}$, where k_j is as defined earlier. Note that the sum of the number of deficient vertices over each of the parts of $H_{i,j}$, $1 \leq j \leq W$, equals d for all fixed i . Moreover, the number of deficient vertices in $H_{i,j}$ is the same for all i when j is held fixed.

We now introduce two distinguished vertices u and v and add additional edges as follows.

- (1) For every $1 \leq i < K(n)$ and for every $1 \leq j \leq W$, all deficient vertices in part V of $H_{i,j}$ are matched to the deficient vertices in part U of $H_{i+1,j}$ (that is, we insert an arbitrary matching between these two sets of vertices).
- (2) All deficient vertices in part U of $H_{1,j}$ for $1 \leq j \leq W$ are connected to u .
- (3) All deficient vertices in part V of $H_{K(n),j}$ for $1 \leq j \leq W$ are connected to v .

Essentially, we are connecting the graphs $H_{i,j}$ for fixed j in series via their deficient vertices, and then connecting the left ends of these chains to the distinguished vertex u and the right ends of the chains to the distinguished vertex v . The construction is illustrated in Figure 2.

We note that the graph G_n constructed as described before is a d -regular bipartite graph with $2dK(n)W + 2 = 2n + o(n)$ vertices.

Consider the sampled graph G'_n . Suppose G'_n has a perfect matching M . In the matching M , if u is matched to a vertex in part U of $H'_{1,j}$ for some $1 \leq j \leq W$, then there must be a vertex in part V of $H'_{1,j}$ that is matched to a vertex in part U of $H'_{2,j}$. Proceeding in the same way, one concludes that for every i , $1 \leq i < K(n)$ there must be a vertex in part V of $H'_{i,j}$ that is matched to a vertex in part U of $H'_{i+1,j}$. Finally, vertex v must be matched to a vertex in part V of $H'_{K(n),j}$. This implies that the sampled graph G'_n can have a perfect matching only if at least one edge survives in G'_n between every pair of adjacent elements in the sequence that follows:

$$u \rightarrow H_{1,j} \rightarrow H_{2,j} \rightarrow \cdots \rightarrow H_{K(n)-1,j} \rightarrow H_{K(n),j} \rightarrow v.$$

Now suppose that we sample edges uniformly with probability p . It follows from the construction of G_n that for any fixed j , the probability that at least one edge survives between every pair of adjacent elements in the sequence $u \rightarrow H_{1,j} \rightarrow H_{2,j} \rightarrow \cdots \rightarrow H_{K(n)-1,j} \rightarrow H_{K(n),j} \rightarrow v$ is equal to

$$(1 - (1 - p)^{k_j})^{K(n)+1} \leq (pk_j)^{K(n)+1}.$$

Hence, the probability that at least one such path survives in G'_n is at most

$$W \left(p \max_{1 \leq j \leq W} k_j \right)^{K(n)+1}$$

by the union bound.

When $d(n) \leq \sqrt{n/\ln n}$, we have $\gamma = 1$, $W = d$, $k_j = 1$ and $K(n) = \lceil n/d^2 \rceil$. So the bound transforms to

$$W p^{K(n)+1} = d p^{\lceil n/d^2 \rceil + 1}, \quad (1)$$

which goes to zero faster than any inverse polynomial function in n when $p = o(1)$ since $K(n) = \lceil n/d^2 \rceil = \Omega(\ln n)$.

When $d \geq \sqrt{n/\ln n}$, we have $k_j \leq \gamma$ where $\gamma = \lceil \frac{d^2 \ln n}{n} \rceil$, $W = \lceil \frac{d}{\gamma} \rceil$ and $K(n) = \lceil \ln n \rceil$. Hence, the bound becomes

$$W (p\gamma)^{K(n)+1} = \left\lceil \frac{d}{\gamma} \right\rceil (p\gamma)^{\lceil \ln n \rceil + 1}, \quad (2)$$

which goes to zero faster than any inverse polynomial function in n when $p = o(\frac{n}{d^2 \ln n})$. This completes the proof of the theorem. \square

The construction given in Theorem 4.1 shows that the sampling upper bound for preserving a perfect matching proved in Theorem 2.1 is tight up to a factor of $O(\ln^2 n)$.

ACKNOWLEDGMENTS. We thank R. Bhattacharjee for many helpful discussions in the early stages of this work. We also thank the anonymous referees whose comments have helped improve the presentation.

REFERENCES

- AGGARWAL, G., MOTWANI, R., SHAH, D., AND ZHU, A. 2003. Switch scheduling via randomized edge coloring. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS'03)*. IEEE Computer Society, 502.
- BENCZÚR, A. 1997. Cut structures and randomized algorithms in edge-connectivity problems. Ph.D. thesis, Massachusetts Institute of Technology.
- BIRKHOFF, G. 1946. Tres observaciones sobre el algebra lineal. *Univ. Nac. Tucumán Rev. Ser. A* 5, 147–151.
- BOLLOBÁS, B. 1998. *Modern Graph Theory*. Springer.
- COLE, R., AND HOPCROFT, J. 1982. On edge coloring bipartite graphs. *SIAM J. Comput.* 11, 3, 540–546.
- COLE, R., OST, K., AND SCHIRRA, S. 2001. Edge-Coloring bipartite multigraphs in $O(E \log D)$ time. *Combinatorica* 21, 1, 5–12.
- GABOW, H., AND KARIV, O. 1982. Algorithms for edge coloring bipartite graphs and multigraphs. *SIAM J. Comput.* 11, 1, 117–129.
- HOPCROFT, J. AND KARP, R. 1973. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM J. Comput.* 2, 4, 225–231.
- KARGER, D. R. 1994a. Random sampling in cut, flow, and network design problems. In *Proceedings of the 26th Annual ACM Symposium on Theory of Computing (STOC'94)*. ACM, New York, 648–657.
- KARGER, D. R. 1994b. Using randomized sparsification to approximate minimum cuts. In *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'94)*. 424–432.
- KÖNIG, D. 1916. Über graphen und ihre anwendung auf determinantentheorie und mengenlehre. *Math. Annalen* 77, 453–465.
- MOTWANI, R., AND RAGHAVAN, P. 1995. *Randomized Algorithms*. Cambridge University Press.
- MUCHA, M., AND SANKOWSKI, P. 2004. Maximum matchings via gaussian elimination. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS'04)*. IEEE Computer Society, 248–255.
- SCHRIJVER, A. 1999. Bipartite edge coloring in $O(\Delta m)$ time. *SIAM J. Comput.* 28, 841–846.
- SPIELMAN, D. A., AND SRIVASTAVA, N. 2008. Graph sparsification by effective resistances. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC'08)*. ACM, New York, 563–568.
- SPIELMAN, D. A., AND TENG, S.-H. 2004. Nearly-Linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC'04)*. ACM, New York, 81–90.
- VON NEUMANN, J. 1953. A certain zero-sum two-person game equivalent to the optimal assignment problem. *Contrib. Optimal Assignm. Problem Theory Games* 2, 5–12.

RECEIVED NOVEMBER 2008; REVISED MARCH 2009; ACCEPTED APRIL 2009