

Counting the Number of Eulerian Orientations ^{*}

Zhenghui Wang

March 16, 2011

1 Introduction

Consider an undirected Eulerian graph, a graph in which each vertex has even degree. An *Eulerian orientation* of the graph is an orientation of its edges such that for each vertex v , the number of incoming edges of v equals to outgoing edges of v , i.e. $d_{in}(v) = d_{out}(v)$. Let $\mathcal{P}_0$ denote the set of all Eulerian orientations of graph G . In this paper, we are concerned with the questions of *sampling* uniformly and *counting* the set of $\mathcal{P}_0$ of an arbitrary graph G .

The significance of counting the number of Eulerian orientation was raised in several different fields. In statistical physics, the crucial partition function “ Z_{ICE} ”, in the so-called “ice-type model” [5], is equal to the number of Eulerian orientations of some underlying Eulerian graph. It has been also observed [2, 4] that the counting problem for Eulerian orientations corresponds to evaluating the Tutte polynomial at the point $(0, -2)$. In addition, If the Eulerian graph G represents the topology of a network, an Eulerian orientation is a unidirectional configuration of the network, which is a maximum flow without source and sink. Consequently, counting the number of Eulerian orientations is equivalent to counting the number of maximum flows around the network.

Finding an Eulerian orientation of a graph G can be accomplished in polynomial time, but sampling and counting the set of all Eulerian orientations is intractable, unless unexpected collapses of complexity classes occur [1]. This behavior of a problem was first observed in the case of perfect matchings of bipartite graphs. For perfect matchings, construction is well known to be solvable in polynomial time, but exact enumeration is complete for the class $\#P$ [3]. The latter hardness result directed efforts toward obtaining efficient randomized approximations. Now, sampling and approximate counting can be achieved via randomized schemes which run in time polynomial in the size of the input graph, the inverse of the desired approximation accuracy, and the ratio of the total number of near-perfect over perfect matchings $|\mathcal{M}_{n-1}|/|\mathcal{M}_n|$ of the input graph.

In this paper, we show that sampling and approximately counting Eulerian orientations reduce to sampling and approximately counting perfect matchings for a class of graphs where $|\mathcal{M}_{n-1}|/|\mathcal{M}_n| = \mathcal{O}(n^4)$. Thus, we obtain efficient solutions for Eulerian orientations. This reduction is a fully combinatorial argument without resorting to any theory of Markov chains.

^{*}This paper is a project of Math 100 in winter term 2011 at Dartmouth College, which explains some results of [1].

2 Preliminaries

For some set $S \neq \emptyset$, a (d, δ) -sampling scheme for S is an algorithm that, with probability at least $1 - \delta$, outputs a element $s \in S$ such that

$$\sum_{x \in S} \left| \Pr\{s = x\} - \frac{1}{|S|} \right| < d.$$

An (ϵ, δ) -approximation scheme for S is an algorithm that outputs an estimate $\hat{S}$ for $|S|$ such that

$$\Pr \left\{ \left| |S| - \hat{S} \right| \leq \epsilon |S| \right\} \geq 1 - \delta.$$

Our aim is to obtain *efficient* (d, δ) -sampling and (ϵ, δ) -approximation schemes for the set $\mathcal{P}_0$ of Eulerian orientations of any Eulerian graph $G = (V, E)$, where $|V| = n$. By efficient, we mean running times polynomial in n , ϵ^{-1} , $\log d^{-1}$ and $\log \delta^{-1}$.

Let $G' = (V'_1, V'_2, E')$ be a bipartite graph, where $|V'_1| = |V'_2| = 2n'$. Let $\mathcal{M}_{n'}$ and $\mathcal{M}_{n'-1}$ be the sets of perfect matchings and near-perfect matchings of G' respectively (a near-perfect matching is a matching of size $n' - 1$). We have following (d, δ) -sampling and (ϵ, δ) -approximation schemes for perfect matchings.

Fact 2.1 ([7, 8] and improvement in [9]). *For the class of graphs G' as above, there is a (d, δ) -sampling scheme for the set of perfect matchings $\mathcal{M}_{n'}$. The scheme runs in time polynomial in $n', \log d^{-1}$, $\log \delta^{-1}$, and the ratio $|\mathcal{M}_{n'-1}|/|\mathcal{M}_{n'}|$.*

Fact 2.2 ([7, 8] and improvement in [9, 10]). *For the class of graphs G' as above, there is a (ϵ, δ) -approximation scheme for the set of perfect matchings $\mathcal{M}_{n'}$. The scheme runs in time polynomial in n', ϵ^{-1} , $\log \delta^{-1}$, and the ratio $|\mathcal{M}_{n'-1}|/|\mathcal{M}_{n'}|$.*

3 Sampling and Approximate Counting: the Reduction

The essential problem of finding an Eulerian orientation is to choose one direction for each edge to balance the number of incoming edges and outgoing edges of each vertex. Intuitively, this problem can be rephrased in the language of matchings in this way: for each edge $e = \{u, v\}$ in G , we create a vertex $w_e \in V'_2$ and two vertices $x_{v,e}$ and $x_{u,e}$ in V'_1 representing two possible direction for edge e . Let vertex w_e can be matched only to $x_{v,e}$ and $x_{u,e}$. In the perfect matchings, exact one of $x_{v,e}$ and $x_{u,e}$ matches w_e . This tells which direction is chosen for edge e . In order to match the other unmatched vertex of $x_{v,e}$ and $x_{u,e}$, we need to add some dummy nodes in V'_2 .

Formally, let $G = (V, E)$, where $|V| = n$ and $|E| = m$, be the undirected Eulerian graph with a set of Eulerian orientations $\mathcal{P}_0$. Let $d(v)$ denote the degree of vertex $v \in V$. The reduction from Eulerian orientations to perfect matchings is as follows (see also Figure 1 in [1]). The bipartite graph $G' = (V'_1, V'_2, E')$ has vertex bipartition

$$V'_1 = \bigcup_{v \in V} X_v, \quad \text{where } X_v = \bigcup_{e \in E: e=\{v,u\}} \{x_{v,e}\},$$

$$V'_2 = \bigcup_{v \in V} Y_v \bigcup_{e \in E} \{w_e\}, \quad \text{where } Y_v = \left\{ y_{v,i}, 1 \leq i \leq \frac{d(v)}{2} \right\},$$

and edges

$$E' = \bigcup_{e \in E: e=\{v,u\}} \{\{x_{v,e}, w_e\}, \{x_{u,e}, w_e\}\} \bigcup_{v \in V} X_v \times Y_v.$$

X_v is the set of $d(v)$ copies of vertex v and Y_v is the set of $d(v)/2$ dummy nodes. Clearly, $n' = |V'_1| = |V'_2| = \sum_{v \in V} d(v) = 2m$, hence the size of G' is polynomial in the size of G .

Lemma 3.1. *For any Eulerian graph G , the set $\mathcal{M}_{n'}$ of perfect matchings of G' can be partitioned so that the partition classes are in one-to-one relation with the set $\mathcal{P}_0$ of Eulerian orientation of G . Furthermore, each partition class has cardinality $\prod_{v \in V} (d(v)/2)!$.*

Proof. Each perfect matching of G' can be associated with a unique Eulerian orientation of G as follows (see Figure 1(b) in [1]). For each edge $e = \{v, u\}$ of G , there is a pair of edges $\{x_{v,e}, w_e\}$ and $\{x_{u,e}, w_e\}$ in G' , exactly one of which must be in the perfect matching to cover w_e . If $\{x_{v,e}, w_e\}$ is in the perfect matching of G' , then we direct $\{v, u\}$ toward v in the orientation of G , otherwise we direct $\{u, v\}$ toward u . Thus, a unique orientation of G is obtained. To see that this is indeed Eulerian, note that for each v , all $d(v)/2$ vertices in Y_v must be matched to $d(v)/2$ of the $d(v)$ vertices in X_v , which balances the incoming and outgoing edges of the corresponding orientation. Hence perfect matchings can be partitioned according to the Eulerian orientation that they are associated with. Following similar reasoning, it is easy to see that each Eulerian orientation that gives rise to a nonempty partition class has cardinality exactly $\prod_{v \in V} (d(v)/2)!$, since one the w_e 's are matched with half of the X_v 's for each v , there are $(d(v)/2)!$ ways of matching the other half of the X_v 's with Y_v 's. $\square$

Now by Facts 2.1 and 2.2, in order to sample and approximately count the set $\mathcal{P}_0$ of Eulerian orientations of G , it suffices to construct the reduction graph G' and sample and approximately count the set of perfect matchings $\mathcal{M}_{n'}$ of G' . In particular, when the (d, δ) -sampling scheme of the set of perfect matchings of G' output a perfect matching of G' , the sampling scheme for Eulerian orientations will output the corresponding Eulerian orientation of G , and this can be easily seen to be a (d, δ) -sampling scheme for $\mathcal{P}_0$. Similarly, when the (ϵ, δ) -approximation scheme for $\mathcal{M}_{n'}$ outputs an estimate $\hat{\mathcal{M}}_{n'}$ for $|\mathcal{M}_{n'}|$, the approximation scheme for $\mathcal{P}_0$ outputs the estimate $\hat{\mathcal{P}}_0 = \hat{\mathcal{M}}_{n'} / \prod_{v \in V} (d(v)/2)!$, and this is an (ϵ, δ) -approximation scheme for $|\mathcal{P}_0|$.

By Fact 2.1 and 2.2, the running times are polynomial in n , ϵ^{-1} , $\log d^{-1}$, $\log \delta^{-1}$, and the ratio $|\mathcal{M}_{n'-1}|/|\mathcal{M}_{n'}|$. In the following lemma, we relate the ratio $|\mathcal{M}_{n'-1}|/|\mathcal{M}_{n'}|$ that determines the efficiency of the running time to a certain ratio concerning orientations. In particular, by analogy to near-perfect matchings, we say that an orientation of G is *near-Eulerian* if, for two vertices u and v we have $d_{\text{out}}(u) = d_{\text{in}}(u) + 2$ and $d_{\text{out}}(v) = d_{\text{in}}(v) - 2$, while for all other vertices the in-degree is equal to the out-degree. Let $\mathcal{P}_1$ denote the set of near-Eulerian orientations of G .

Lemma 3.2. *For any Eulerian graph G , $|\mathcal{M}_{n'-1}|/|\mathcal{M}_{n'}| = \mathcal{O}(n^2 |\mathcal{P}_1|/|\mathcal{P}_0|)$.*

The proof is similar to that of Lemma 3.1. In Lemma 3.1, we relate the set $\mathcal{M}_{n'}$ of G' to the set of $\mathcal{P}_0$ of G and obtain $|\mathcal{M}_{n'}| = |\mathcal{P}_0| \prod_{v \in V} (d(v)/2)!$. We can relate the set $\mathcal{M}_{n'-1}$ of G'

to the set of $\mathcal{P}_1$ of G in a similar way and upper bound $|\mathcal{M}_{n'-1}|$ in terms of $|\mathcal{P}_1|$. Combining all these above yields this lemma.

By all lemmas and facts before, we can efficiently sample and approximately count Eulerian orientations for the class of graph G , provided that $|\mathcal{P}_1|/|\mathcal{P}_0|$ can be upper bounded by a polynomial in n . In the next section, we will obtain the desired polynomial bound on the ratio $|\mathcal{P}_1|/|\mathcal{P}_0|$.

4 Monotonicity Lemma

Before we prove this key lemma, let us first generalize the notion of Eulerian orientation. We call a directed graph $P = (V, H)$ an *orientation* of Eulerian graph $G = (V, E)$ if either $(u, v) \in H$ or $(v, u) \in H$ whenever $\{u, v\} \in E$.

For an orientation $P = (V, H)$ and vertex $v \in V$, the *charge* of v is defined as

$$q(v) = \frac{d_{out}(v) - d_{in}(v)}{2}.$$

The charge of v indicates the number of outgoing edges incident to v that should be reversed in order for v to have equal number of incoming and outgoing edges. Moreover, any orientation P could be viewed as an implementation of certain charge function q . In particular, the set $\mathcal{P}_0$ of Eulerian orientations are all implementations of charge function q with $q(v) = 0$ for all $v \in V$.

Lemma 4.1 (Monotonicity Lemma¹). *For any Eulerian graph $G = (V, E)$, where $|V| = n$, we have $|\mathcal{P}_1| \leq n(n-1)|\mathcal{P}_0|$.*

Proof. Consider only charge function q such that:

1. There exists exactly one vertex u with $q(u) = 1$;
2. There exists exactly one vertex v with $q(v) = -1$;
3. $q(x) = 0$ for all vertex $x \in V - \{u, v\}$.

Clearly, each orientation that implements such a charge function q is near-Eulerian. Let $\mathcal{P}(q)$ be the set of orientations that implement charge function q . We have $\mathcal{P}_1 = \bigcup_q \mathcal{P}(q)$. Since all sets of $\mathcal{P}(q)$'s are disjoint for different q 's, we also have $|\mathcal{P}_1| = \sum_q |\mathcal{P}(q)|$.

We will show that, for any fixed q , $|\mathcal{P}(q)| \leq |\mathcal{P}_0|$. Once this is established, the monotonicity lemma follows by

$$|\mathcal{P}_1| = \sum_q |\mathcal{P}(q)| \leq \sum_q |\mathcal{P}_0| = n(n-1)|\mathcal{P}_0|.$$

The last equality follows that the number of charge function q satisfying those three properties is $n(n-1)$ by the number of choices of u and v .

To prove $|\mathcal{P}(q)| \leq |\mathcal{P}_0|$, we relate orientations in $\mathcal{P}(q)$ with orientations in $\mathcal{P}_0$ so that, in a local but sufficiently strong sense, the number of orientations in $\mathcal{P}_0$ related with each orientation in $\mathcal{P}(q)$ is larger than the number of orientations in $\mathcal{P}(q)$ related with each orientation

¹This lemma is a special case of the original monotonicity lemma in [1], but suffices to obtain the the desired polynomial bound on the ratio $|\mathcal{P}_1|/|\mathcal{P}_0|$. The proof is essentially the same as in [1], but simpler hopefully.

in $\mathcal{P}_0$. The natural way to relate orientations in $\mathcal{P}(q)$ with orientations in $\mathcal{P}_0$ is by reversing a path with endpoints u and v . We decompose an Eulerian orientation into disjoint paths or circuits so that the number of paths to be reversed can be easily accounted for.

For fixed orientation $P = (V, H) \in \mathcal{P}(q) \cup \mathcal{P}_0$, and for each $x \in V - \{u, v\}$, consider one of the $(d(x)/2)!$ ways of pairing incoming and outgoing edges of x (see Figure 2(a) in [1]). Let ϕ_x denote such a pairing. The partition of H into Eulerian elements suggested by some fixed $\phi := \{\phi_x : x \in V - \{u, v\}\}$ is as follows (see also Figure 2(b) in [1]):

- We have (u, v) -paths of the form
 $u = x_1, x_2, \dots, x_l = v$ with $x_i \in V - \{u, v\}$, $1 < i < l$, and $x_{i+1} = \phi_{x_i}(x_{i-1})$.
- We have *free circuits* of the form
 $x_1, x_2, \dots, x_l$ with $x_i \in V - \{u, v\}$, $1 \leq i \leq l$, and $x_{i+1} = \phi_{x_i}(x_{i-1})$.
- We have (u, u) -paths of the form
 $u = x_1, x_2, \dots, x_l = u$ with $x_i \in V - \{u, v\}$, $1 < i < l$, and $x_{i+1} = \phi_{x_i}(x_{i-1})$.
- We have (v, v) -paths of the form
 $v = x_1, x_2, \dots, x_l = v$ with $x_i \in V - \{u, v\}$, $1 < i < l$, and $x_{i+1} = \phi_{x_i}(x_{i-1})$.

It is easy to verify that this is indeed a partition of H , thus each edge belongs to exactly one circuit or path. Furthermore, two distinct pairings $\phi \neq \phi'$ induce distinct partitions.

In addition, for any $P = (V, H) \in \mathcal{P}(q) \cup \mathcal{P}_0$, there are exactly $\prod_{x \in V - \{u, v\}} (d(x)/2)!$ distinct pairings ϕ . Thus, for each such P , we may consider $\prod_{x \in V - \{u, v\}} (d(x)/2)!$ distinct copies of P , and associate with each copy a distinct ϕ . We denote such a pair $\langle P, \phi \rangle$, and identify it with the unique partitioning of H with respect to ϕ . (See Figure 3 in [1]).

Now for $P \in \mathcal{P}(q)$, $P' \in \mathcal{P}_0$ and pairings ϕ for P and ϕ' for P' , we relate $\langle P, \phi \rangle$ with $\langle P', \phi' \rangle$ if and only if the partition of P with respect to ϕ is identical to the partition of P' with respect to ϕ' , except for one (u, v) -path of $\langle P, \phi \rangle$ which, if reversed, is a (v, u) -path of $\langle P', \phi' \rangle$.

Note that if $\langle P, \phi \rangle$ has $2 + c$ (u, v) -paths and hence c (v, u) -paths, then it is related with $2 + c$ distinct $\langle P', \phi' \rangle$'s, and that each such $\langle P', \phi' \rangle$ has $c + 1$ (v, u) -paths, thus it is related with $c + 1$ distinct $\langle P, \phi \rangle$'s. Since $2 + c > c + 1$, all the above imply that

$$|\mathcal{P}(q)| \left(\prod_{x \in V - \{u, v\}} \left(\frac{d(x)}{2} \right)! \right) \leq |\mathcal{P}_0| \left(\prod_{x \in V - \{u, v\}} \left(\frac{d(x)}{2} \right)! \right),$$

which proves $|\mathcal{P}(q)| \leq |\mathcal{P}_0|$ and hence the monotonicity lemma. $\square$

By Facts 2.1 and 2.2, and Lemma 3.1, 3.2 and 4.1, we have the following theorem.

Theorem 4.2. *For the class of Eulerian graphs $G = (V, E)$, where $|V| = n$, there are (d, δ) -sampling and (ϵ, δ) -approximate counting scheme for the set $\mathcal{P}_0$ of Eulerian orientations of G . The running time is polynomial in n , ϵ^{-1} , $\log d^{-1}$, and $\log \delta^{-1}$.*

5 Open Question: Counting Eulerian Circuits

An *Eulerian circuit* of Eulerian graph G is a closed path, with a direction but no distinguished starting point, which traverses each edge exactly once.

The problem of counting Eulerian circuits is $\#P$ -complete [6]. So we desire a (ϵ, δ) -approximation scheme. Unlike the Eulerian orientation problem, the problem of designing a (ϵ, δ) -approximation scheme for Eulerian circuits is still open. See [6] for a good reference of this problem.

References

- [1] M. Mihail and P. Winkler, On the number of Eulerian orientations of a graph, *Algorithmica* 1996, 402–414.
- [2] F. Jaeger, D. L. Vertigan, and D. J. A. Welsh, On the Computational Complexity of Jones and Tutte Polynomials, *Math. Proc. Cambridge Philos. Soc.*, **108**(35) (1990), 35–53.
- [3] L. G. Valiant, The Complexity of Computing the Permanent, *Theoret. Comput. Sci.*, (1979), 189–201.
- [4] D. J. A. Welsh, *On the Tutte Map, Lecture Notes, DIMACS, May 1991/Dagstuhl, June 1991*.
- [5] D. J. A. Welsh, The Computational Complexity of Some Classical Problems from Statistical Physics, in *Disorder in Physical Systems*, Oxford University Press, Oxford, 1990, 307–321.
- [6] G. R. Brightwell and P. Winkler, Counting Eulerian circuits is $\#P$ -complete, *Proc. 7th ALENEX and 2nd ANALCO 2005*, 259–262.
- [7] A. Z. Broder, How Hard Is It To Marry at Random? (On the Approximation of the Permanent), *STOC 1986* 50–58.
- [8] M. R. Jerrum and A. Sinclair, Conductance and the Rapid Mixing Property for Markov Chains: the Approximation of the Permanent Resolved, *STOC 1988*, 235–243.
- [9] P. Diaconis and D. Stroock, Geometric Bounds for Eigenvalues of Markov Chains, *Ann. Appl. Probab.*, **1** (1991), 36–61.
- [10] P. Dagum and M. Luby, Approximating the Permanent of Graphs with Large Factors, *SIAM J. Comput.*, to appear.