

Note on Counting Eulerian Circuits

Graham R. Brightwell * Peter Winkler †

May 2004

CDAM Research Report LSE-CDAM-2004-12

Abstract

We show that the problem of counting the number of Eulerian circuits in an undirected graph is complete for the class $\#P$.

1 Introduction

Every basic text in graph theory contains the story of Euler and the Königsberg bridges, together with the theorem that guarantees the existence of a circuit traversing every edge of a graph exactly once, if and only if the graph is connected and all vertices have even degree.

Later in the text one might find the “Matrix Tree Theorem”, which provides an efficient algorithm for counting the number of spanning trees of a graph; and an application of this and the so-called “BEST” Theorem (see below) to count the number of Eulerian circuits in a *directed* graph. But what about counting Eulerian circuits in an undirected graph?

This problem is clearly in the class $\#P$ (introduced by Valiant [11] in the 1970s), since it is easy to check whether a candidate circuit traverses each edge once. Since there is no known efficient way to count Eulerian circuits, it is natural to suspect that the problem is $\#P$ -complete, and thus presumably very difficult—especially in view of Toda’s result [9], which implies that one call to a $\#P$ oracle suffices to solve any problem in the polynomial hierarchy in deterministic polynomial time.

Researchers have shown a myriad of graphical counting problems to be $\#P$ -complete, including Hamilton circuits [12], acyclic orientations [6], and Eulerian orientations [7]. Yet, the complexity of counting Eulerian circuits remained open (see e.g. [5], Open Problem on p. 5, at

*Department of Mathematics, London School of Economics and Political Science, London, UK

†Bell Labs 2C-365, 700 Mountain Ave., Murray Hill, NJ 07974; and Institute for Advanced Study, Princeton NJ 08540. Research supported in part by ONR grant N00014-03-M-0141.

the end of Section 1.1) for 25 years—rather mysteriously, especially considering the simplicity of the reduction below.

Our approach is to show that, with the help of an oracle which counts Eulerian circuits, a Turing machine can count the number of Eulerian orientations of any given graph in polynomial time. The latter problem was shown to be #P-hard by Mihail and Winkler [7, 8] and for completeness we provide a second proof below.

Both reductions proceed by enumeration modulo various primes, a technique introduced originally by Valiant [11] and utilized later by the authors [1, 2] to settle another long-open complexity problem, counting the linear extensions of a partially ordered set.

To show later that we can reconstruct a number uniquely from its values modulo a small set of primes, it is useful to have a technical lemma such as the following.

Lemma 1.1. *For any $n \geq 4$, the product of the set of primes strictly between n and n^2 is at least $n!2^n$.*

Proof. We use some facts from Hardy and Wright [4], Chapter 22, concerning the functions

$$\vartheta(n) = \log \prod_{p \leq n} p, \text{ where } p \text{ runs over all primes less than } n, \text{ and } \psi(n) = \sum_{i=1}^{\log n / \log 2} \vartheta(n^{1/i}).$$

From [4] we find that $\vartheta(n) < 2n \log 2$ for $n \geq 1$, and that $\psi(n) \geq \frac{1}{4}n \log 2$ for $n \geq 2$.

We are interested in the quantity $V = \vartheta(n^2) - \vartheta(n)$. From the above facts, we have:

$$\begin{aligned} V &\geq \psi(n^2) - \sum_{i=2}^{2 \log n / \log 2} \vartheta(n^{2/i}) - \vartheta(n) \\ &\geq \frac{1}{4}n^2 \log 2 - \frac{2 \log n}{\log 2} \cdot 2n \log 2 - 2n \log 2 \\ &\geq n \log n \geq \log(n!2^n), \end{aligned}$$

at least provided $n \geq 150$. The inequality for $4 \leq n < 150$ is easily verified by direct calculation. □

It is evident that this lemma is not tight: it is possible to replace the n^2 upper limit by $Kn \log n$, for some suitably large K .

2 Circuits, Orientations, Arborescences and Orbs

For us a graph $G = \langle V, E \rangle$ will be finite and undirected, with no loops or multiple edges; if multiple edges are permitted we use the term *multigraph*. A *circuit* C of G is a closed path,

with a direction but no distinguished starting point; it is *Eulerian* if it traverses each $e \in E$ exactly once. Of course the possession of even one Eulerian circuit implies all degrees are even, and it will be convenient for us to denote the degree of a vertex v by $2d_v$ instead of d_v .

An *Eulerian orientation* of G is an orientation of its edges with the property that each vertex has the same number (namely, d_v) of incoming and outgoing arcs. Any Eulerian circuit induces an Eulerian orientation by orienting each edge in accordance with its direction of traversal.

If a particular starting edge is chosen for the Eulerian circuit C , originating say at vertex r , then C also induces a spanning tree $T = \{\text{exit}(v) : v \neq r\}$ where $\text{exit}(v)$ is the last edge incident to v used by C before its final return to r . When oriented according to C , T becomes an in-bound spanning tree, or *arborescence*, rooted at r .

Let us fix a root $r \in V$ and denote by the term *orb* a pair consisting of an Eulerian orientation and an arborescence (for that orientation) rooted at r . If an orb is specified, it is easy to construct a corresponding Eulerian circuit C : simply begin walking from the root r , following any unused outgoing arc from each vertex v , except that the tree arc exiting v is avoided as long as possible. Since there are $(d_v - 1)!$ ways to order the non-tree outgoing arcs from v , and $d_r!$ from r , the number of ways to construct C is precisely $d_r! \prod_{v \neq r} (d_v - 1)!$. However, this over-counts Eulerian circuits (as we have defined them) by a factor of d_r since each circuit passes d_r times through r . Hence orbs and Eulerian circuits are in perfect $\prod_{v \in V} (d_v - 1)!$ -to-1 correspondence, and thus counting Eulerian circuits is equivalent to counting orbs. This result is sometimes known as the “BEST” Theorem after de **B**ruijn, van Aardenne-**E**hrenfest, **S**mith and **T**utte, although the former two should perhaps get additional credit as the original discoverers.

Note that for any *particular* orientation, one can use the BEST Theorem together with the Matrix-Tree Theorem [10] (due to Tutte, but implicit in Kirchhoff’s work of 1847) to count Eulerian circuits. The difficulty in using this approach to count (or even approximate) the number of Eulerian circuits in G is that the number of arborescences in an orientation can vary enormously.

Theorem 2.1. *Counting Eulerian circuits is #P-complete.*

Proof. It suffices to reduce the problem of counting Eulerian orientations to counting orbs in a multigraph; the latter is equivalent to counting orbs in a simple graph since multiple edges can

be subdivided without affecting the number of orbs. Let us therefore assume that an Eulerian graph $G = \langle V = \{1, \dots, n\}, E \rangle$ has been given, and that we have an oracle for counting orbs in any multigraph. We wish to compute the number N of Eulerian orientations of G .

We construct for any odd prime p , a graph G_p whose number of orbs is equivalent to N modulo p . The construction is shockingly simple: each edge e of E is replaced by p parallel edges e_1, \dots, e_p , and a new node 0 is added which is adjacent to every node v of V by two parallel edges, $e(v)$ and $e'(v)$. We take 0 as the root of all orbs.

Figure 1 shows a small G and the resulting G_p when $p = 3$.




Figure 1: Reducing EULERIAN ORIENTATION COUNT TO ORB COUNT

In G_p , the *type* τ of an orb is a function from E to $\{0, 1, \dots, p\} \times \{T, F\}$ which tells how many of e_1, \dots, e_p are oriented from the smaller to the larger-numbered vertex, and whether any is a tree edge (“T” means “yes”, “F” means “no”). A type τ is *special* if $\tau(e) \in \{0, p\} \times \{F\}$ for every $e \in E$.

If an orb belongs to a special type, then the common direction of the parallel edges corresponding to each $e \in E$ provides an orientation of G . This orientation is Eulerian, as otherwise the in-degree and out-degree of some vertex of G_p will differ by at least $2p - 2$. Therefore, in an orb of special type, for each vertex v of G_p , exactly one of $e(v)$ and $e'(v)$ is directed towards the root 0 , and this is the arc that carries the edge of the associated tree directed away from v . Thus, the number of orbs of special type is precisely $2^n \times N$.

On the other hand, we claim that the number of orbs of any non-special type τ is a multiple of p . To see this let e be such that $\tau(e) = (k, X) \notin \{0, p\} \times \{F\}$. Suppose first that $0 < k < p$ and $X = F$; then the orbs of type τ can be partitioned into $\binom{p}{k}$ equal parts according to which of e_1, \dots, e_p are oriented from the smaller to the larger-numbered vertex, and of course $\binom{p}{k}$ is

a multiple of p .

If $0 < k < p$ and $X = T$, then the part sizes are multiples either of $k \binom{p}{k}$ or $(p-k) \binom{p}{k}$ since we must also decide which of the correctly-oriented arcs belongs to the tree.

Finally, if $k \in \{0, p\}$, and $X = T$, we partition according to which of the now-parallel arcs is in the tree and there are p choices.

It follows that the total number of orbs is equivalent to $2^n \times N$ modulo p , and thus we can compute $N \pmod p$. We repeat this process for every prime p between m and $2m$, where $|E| = m$, and apply the Chinese Remainder Theorem to nail N . \square

A similar argument, equally simple, can be used to reduce NOT-ALL-EQUAL 3-SAT COUNT (shown in [3] to be #P-complete) to COUNTING EULERIAN ORIENTATIONS. Given an instance of NOT-ALL-EQUAL 3-SAT COUNT, the graph is provided with a vertex for each literal and another for each clause, plus one spare vertex s . Each clause is connected by a single edge to its three literals and s . Each literal vertex is given p parallel edges to its mate (where p is a prime larger than the number of appearances of any literal), and a total of p other edges, consisting of some number (as mentioned above) to the clauses in which it appears, with the remainder going to s .

Figure 2 shows the construction for a particular 2-clause, 3-variable instance with $p = 3$.




Figure 2: Reducing NOT-ALL-EQUAL 3SAT COUNT to EULERIAN ORIENTATION COUNT

As in the proof above, orientations which fail to align in parallel all p edges associated with any given variable fall into classes of size $0 \pmod p$. Each of the other “special” orientations has the property that for each variable x , either every edge from x to a clause vertex points outward

and every edge from \bar{x} to a clause vertex points inward, or vice-versa. Such an orientation corresponds to a satisfying assignment, since we cannot have three of the four edges incident to a clause vertex pointing the same way. Conversely, given a satisfying assignment, we orient all the edges between x and \bar{x} towards the true literal, orient all other edges incident with a literal vertex away from true literals and towards false literals, and orient the edge between each clause vertex and s in the necessary manner. This ensures that the in-degree is equal to the out-degree at every vertex other than s , and therefore also at s . Hence there is an exact correspondence between special orientations and satisfying assignments, and we proceed as before.

Remark

Still open is the question of whether there is a fully polynomial randomized approximation scheme (“fpras”) for counting Eulerian orientations (as there is, e.g., for Eulerian orientations [8]). We believe that there is, and even that a particular Markov chain whose states are orbs and near-orbs mixes rapidly. We hope and expect that this question will not remain open for another 25 years.

References

- [1] G.R. Brightwell and P. Winkler, Counting linear extensions is #P-complete, extended abstract, *Proc. 23rd ACM Symp. on the Theory of Computing* (1991) 175–181.
- [2] G.R. Brightwell and P. Winkler, Counting linear extensions, *ORDER* **8** #3 (1991) 225–242.
- [3] N. Creignou and M. Hermann, Complexity of generalized satisfiability counting problems, *Information and Computation* **25** #1 (Feb. 1996) 1–12.
- [4] G.H. Hardy and E.M. Wright, *An Introduction to the Theory of Numbers*, 4th Ed., Oxford University Press 1960.
- [5] M. Jerrum, *Counting, Sampling and Integrating: Algorithms and Complexity*, Lectures in Mathematics, ETH Zrich, Birkhuser Verlag, Basel, 2003.
- [6] N. Linial, Hard enumeration problems in geometry and combinatorics, *SIAM J. Alg. Disc. Meth.* **7** No. 2 (1986) 331–335.
- [7] M. Mihail and P. Winkler, On the number of Eulerian orientations of a graph, extended abstract, *Proc. 3rd ACM-SIAM Symposium on Discrete Algorithms* (1992) 138–145.

- [8] M. Mihail and P. Winkler, On the number of Eulerian orientations of a graph, *Algorithmica* **16** (1996) 402–414.
- [9] S. Toda, On the computational power of PP and +P, *Proc. 30th IEEE Symposium on Foundations of Computer Science* (1989) 514–519.
- [10] W.T. Tutte, *Graph Theory*, Addison-Wesley, New York, 1984.
- [11] L.G. Valiant, The complexity of computing the permanent, *Theoret. Comput. Sci.* **8** (1979) 189–201.
- [12] L.G. Valiant, The complexity of enumeration and reliability problems, *SIAM J. Comput.* **8** (1979) 410–421.