2012 IEEE 53rd Annual Symposium on Foundations of Computer Science

A Permanent Approach to the Traveling Salesman Problem

Nisheeth K. Vishnoi
Microsoft Research
Bangalore, India
Email: nisheeth.vishnoi@gmail.com

***Abstract*—A randomized polynomial time algorithm is presented which, for every simple, connected, k-regular graph on n vertices, finds a tour that visits every vertex and has length at most $\left(1+\sqrt{64/\ln k}\right)n$ with high probability. The proof follows simply from results developed in the context of permanents; Egorychev's and Falikman's theorem which lower bounds the permanent of a doubly stochastic matrix and the polynomial time algorithm of Jerrum, Sinclair and Vigoda which samples a near-random, perfect matching from a bipartite graph. The techniques in this paper suggest new permanent-based approaches for TSP which could be useful in attacking other interesting cases of TSP.**

***Keywords*-Approximation Algorithms, Traveling Salesman Problem, Permanent**

I. Introduction

The Traveling Salesman Problem (TSP) is the following: Given an undirected graph $G = (V, E)$ and a cost function $c : E \mapsto \mathbb{R}_{\geq 0}$, find a subset of edges of minimum cost which connects all the vertices such that every vertex has degree 2. The most interesting case of TSP, referred to as METRICTSP, occurs when the cost function c satisfies the triangle inequality. In this case, one may assume that $E = V \times V$. METRICTSP is a central problem in optimization and computer science and has been one of the most intensely studied problems for over half a century. The problem is NP-hard and Christofides [9] provided a $3/2$ approximation to METRICTSP. Further, it was shown in [27] that it is NP-hard to approximate METRICTSP to a factor better than $220/219$. Thus remains the status of the general problem in spite of a remarkable amount of structural and algorithmic results. Further, a lot of work has been done to understand for which cost functions the factor of $3/2$ can be improved. Two important sequences of work deserve mention in this direction.

The first is the work that gives a PTAS for the case when the metric is Euclidean [5], [24]. Subsequent work includes a PTAS on planar instances [6], [16], [21] and a PTAS for low-genus metrics [10].

The second sequence of work is for instances where the cost function is the shortest path metric on an unweighted input graph. This problem, referred to as GRAPHTSP, seems to capture the difficulty associated to METRICTSP. Indeed, until recently, there was no better algorithm known for GRAPHTSP than Christofides'. The best lower bound example, achieving $4/3$, for the Held-Karp LP-relaxation [17] for METRICTSP is in fact an instance of GRAPHTSP, and this problem is known to be APX-hard [16]. In another series of work, initiated by [15], [25] and improved by [26], [30], the approximation factor for GRAPHTSP was brought down to $7/5$. This progress was preceded and perhaps inspired by work on cubic graphs [14], [1], [7].

It is believed that the limit to these approaches is $4/3$ since there exists a sub-cubic integrality gap GRAPHTSP instance achieving a lower bound of $4/3 - \delta$ for any $\delta > 0$ for the Held-Karp LP. On the other hand, it is plausible that the Held-Karp relaxation for GRAPHTSP gets better as the number of edges in the graph increases. One evidence for this is a theorem by Dirac [11] which asserts that every graph with minimum degree $n/2$ has a Hamiltonian cycle. Another series of results, starting with Karp [20] (see also [4], [22], [2], [8], [28]), show that many reasonable models of random graphs (including random regular graphs) contain a Hamiltonian cycle which can be found efficiently.

In this paper we establish that, indeed, the worst case instances for GRAPHTSP might have low degree by showing that, when the graph is k-regular, arbitrarily good tours exist and can be found in polynomial time. The following is the main result of this paper.

Theorem I.1 (Main Theorem)**.** *There is a randomized, polynomial-time algorithm which, for every k and $n \geq n_0(k)$, given a simple, connected, k-regular graph G on n vertices, finds a tour of length at most $\left(1+\sqrt{64/\ln k}\right)n$ with probability at least $1 - 1/n$.*

The proof relies on two celebrated results on permanents. 1) The results of Egorychev [12] and Falikman [13] that lower bound the permanent of a doubly stochastic matrix. 2) The result of Jerrum, Sinclair and Vigoda [19] that gives a polynomial time algorithm to sample, a nearly-random perfect matching from a bipartite graph.

Importantly, our techniques suggest new permanent-based approaches for TSP. The proof is short and we proceed directly to it. Subsequently we discuss possible extensions in Section III.

0272-5428/12 $26.00 © 2012 IEEE
DOI 10.1109/FOCS.2012.81

IEEE computer society

II. Algorithm and Proof

Overview: Consider the following algorithm: Pick a random cycle cover[1] C in G, find a minimum spanning tree T in the graph obtained from contracting all the cycles in C, and output the Eulerian tour consisting of C and two copies of each edge in T. In the GRAPHTSP case, the cost of the tree is one less than the number of cycles in C. Hence, the length of the tour, in excess of n, is at most two times the number of cycles in C. Thus, the smaller the number of cycles in a typical cycle cover in G, the better the algorithm performs provided we can sample a random cycle cover from G efficiently. This is where the Jerrum, Sinclair and Vigoda result comes in. It allows us to sample from a closely related distribution; the distribution on cycle covers that is obtained by picking a perfect matching in a natural bipartite graph associated to G. In essence, C is obtained by picking two perfect matchings in G in a correlated manner. While in general a graph may not have any cycle cover, a k-regular graph has plenty. This follows from the lower bound of $({}^{k}\!/\!{}_{e})^n$ on the permanent of the adjacency matrix of G due to [12], [13], [29]. Moreover, by a straightforward counting argument, we can show that only a small fraction of these cycle covers can have a large number of vertices that are a part of a small cycle. Roughly, this happens because, for any vertex, the number of cycles in G of length at most l that contain it is upper bounded by k^{l-1}. We now formalize this argument to show that, in any k regular graph, a typical cycle cover has $o_k(n)$ cycles to prove Theorem I.1.

Permanents: Let A denote the adjacency matrix of G. The permanent of A is $\mathrm{per}(A) \stackrel{\mathrm{def}}{=} \sum_{\sigma \in S_n} \prod_{i=1}^{n} A_{i,\sigma(i)}$. While our graph is simple, the following theorem lower bounds the permanent of a k-regular multi-graph.

Theorem II.1. *[12], [13], [29], If A is the adjacency matrix of a k-regular graph on n vertices, then $\mathrm{per}(A) \geq ({}^{k}\!/\!{}_{e})^n$.*

We state a slightly weaker version of the Jerrum, Sinclair, Vigoda result [19] which suffices for our purpose.

Theorem II.2. *There is an algorithm which, given a bipartite graph G', with D being the uniform distribution on its perfect matchings, outputs a sample from a distribution $\widetilde{D}$ such that $\|D - \widetilde{D}\|_{\mathrm{TV}} \leq \varepsilon$.[2] The running time is polynomial in the size of the graph and $\log {}^{1}\!/\!{}_{\varepsilon}$.*

Matchings and cycle covers: For a graph $G = (V, E)$, consider the (natural) bipartite graph $G' = (V_L, V_R, E')$ where $V_L \stackrel{\mathrm{def}}{=} \{v^L : v \in V\}$, $V_R \stackrel{\mathrm{def}}{=} \{v^R : v \in V\}$ and $E' \stackrel{\mathrm{def}}{=} \{(u^L, v^R), (v^L, u^R) : uv \in E\}$. For a perfect matching M' in G' let $C(M')$ be the multi-subset $\{uv \in E : (u^L, v^R) \in M'\}$. Any $C = C(M')$ is a cycle cover of G since it is precisely a collection of vertex-disjoint cycles that cover V. Note that an edge uv in $C(M')$ is repeated twice, and hence gives rise to a cycle of length 2, if and only if both $(u^L, v^R), (v^L, u^R) \in M'$. In general, however, the mapping is many-to-one since each cycle in C of length at least 3 corresponds to exactly two matchings in G'. For a cycle cover C, let $N_i(C)$ denote the number of cycles of length i in C. Then, the number of perfect matchings in G' that are mapped to C is exactly $2^{\sum_{i=3}^{n} N_i(C)}$. Since the total number of perfect matchings in G' is easily seen to be $\mathrm{per}(A)$, if we pick a perfect matching in G' uniformly at random, this induces a probability of picking a cycle cover C in G, namely $\mu(C) \stackrel{\mathrm{def}}{=} \frac{2^{\sum_{i=3}^{n} N_i(C)}}{\mathrm{per}(A)}$.

The algorithm: Given a graph $G = (V, E)$, construct the bipartite graph G' as above. Pick a near-random perfect matching M' in G' using the Jerrum, Sinclair and Vigoda algorithm with $\varepsilon \stackrel{\mathrm{def}}{=} {}^{1}\!/\!{}_{(2n^2)}$. Compute $C(M')$. Collapse all vertices belonging to the same cycle in $C(M')$ to obtain a minor H of G. Find a minimum-cost, spanning tree T in H. Output the multi-graph $(V, C(M') \cup T \cup T)$, which is Eulerian and connected. A tour is derived from this Eulerian tour by short-cutting.

Remark II.3. *Instead of a uniform distribution on the perfect matchings in G', our algorithm uses a near-uniform distribution on perfect matchings $\widetilde{D}$, which in turn induces a probability $\widetilde{\mu}$ of obtaining a cycle cover C in G. However, the mapping from perfect matchings in G' to cycle covers in G will not increase the total variation distance since $\|\mu - \widetilde{\mu}\|_{\mathrm{TV}} = {}^{1}\!/\!{}_{2} \sum_C |\mu(C) - \widetilde{\mu}(C)|$ which, by the triangle inequality, is at most ${}^{1}\!/\!{}_{2} \sum_C \sum_{M' : C(M') = C} |D(M') - \widetilde{D}(M')| = \|D - \widetilde{D}\|_{\mathrm{TV}}$.*

Counting subgraphs with small cycles: For a simple graph of degree at most k we will need an upper bound on the number of cycle covers that only use cycles of length less than l.

Lemma II.4. *For every l, $k \geq l^l$, and a simple graph F of degree at most k, there are at most $k^{(1 - 1/l)|V(F)|}$ spanning subgraphs $\widetilde{F}$ of F that satisfy the following properties:*

1) *$V(\widetilde{F}) = V(F)$.*
2) *$\widetilde{F}$ can be decomposed into vertex-disjoint edges and cycles.*
3) *Each cycle in $\widetilde{F}$ is of length at most $l - 1$.*

Proof: The proof is by strong induction on $|V(F)|$. Let $f(r)$ be an upper bound on the number of subgraphs $\widetilde{F}$ as defined above in a simple graph on r vertices with maximum degree at most k. Note that $f(0) = 1$ and, for the induction hypothesis, assume that $f(r') \leq k^{(1-1/l)r'}$ for all $r' < r$. Consider the number of cycles (or edges) of length less than l that contain vertex 1; remove these cycles and recurse. If vertex 1 belongs to a cycle of length t, then the

[1] A cycle cover in G is a collection of vertex-disjoint cycles, possibly containing length 2 cycles obtained by picking an edge twice, which cover all the vertices.

[2] The total variation distance is defined for distributions over a finite set X to be $\|F - \widetilde{F}\|_{\mathrm{TV}} \stackrel{\mathrm{def}}{=} {}^{1}\!/\!{}_{2} \sum_{x \in X} |F(x) - \widetilde{F}(x)|$.

number of choices is at most k^{t-1}. Then,

$$f(r) \le kf(r-2) + k^2 f(r-3) + \cdots + k^{l-2} f(r-l+1).$$

Then, by the induction hypothesis $f(r)$ is at most

$$\begin{aligned}
&\le k \cdot k^{(1-1/l)(r-2)} + \cdots + k^{l-2} \cdot k^{(1-1/l)(r-l+1)} \\
&\le k^{(1-1/l)r} \cdot k^{-(1-1/l)} \left(k^{1/l} + \cdots + k^{(l-2)/l}\right) \\
&\le k^{(1-1/l)r} \cdot k^{-(1-1/l)} \cdot l \cdot k^{(l-2)/l} \\
&= k^{(1-1/l)r} \cdot l \cdot k^{-1/l}.
\end{aligned}$$

The last inequality follows from the assumption that $l^l \le k$. ■

Proof of Theorem I.1: When G is connected and k-regular, G' always has a perfect matching. In fact, from Theorem II.1 $\operatorname{per}(A) \ge (k/e)^n$. Let $C = C(M')$ be the cycle cover obtained from an M' sampled as in our algorithm using the distribution $\widetilde{D}$ on near-random perfect matchings in G'. Recall that the induced distribution on cycle covers in G is $\widetilde{\mu}$. The cost of the Eulerian tour generated from C is at most $\sum_{i=2}^n iN_i(C) + 2[(\sum_{i=2}^n N_i(C)) - 1]$ since the number of vertices in the contracted graph H is $\sum_{i=2}^n N_i(C)$. For every cycle cover C we know that $\sum_{i=2}^n iN_i(C) = n$, hence, the cost is $n - 2 + 2\sum_{i=2}^n N_i(C)$.

Let $2 \le l \le n$ be an integer (depending on, but much less than, k) which will be fixed later. Start by noticing that, trivially, any cycle cover C satisfies $\sum_{i=l}^n N_i(C) \le n/l$. Thus, it suffices to upper bound the probability that $\sum_{i=2}^{l-1} N_i(C)$ is large. Let $0 < \gamma \le 1$ be a constant which will be fixed later. We will upper bound $\mathbb{P}_{\widetilde{\mu}}[\sum_{i=2}^n N_i(C) \ge \gamma n]$ by $1/n$ for large enough n. This will imply that with probability at least $1 - 1/n$, the cost of the Eulerian tour is at most $n(1 + 2/l + 2\gamma)$.

Towards this end, for an integer $0 \le t \le n$, let $\mathcal{C}_t$ be the set of cycle covers of G for which exactly t vertices are covered by cycles of length at most $l-1$. Then, $\mathbb{P}_{\widetilde{\mu}}\left[\sum_{i=2}^{l-1} N_i(C) \ge \gamma n\right] = \sum_{t=0}^n \mathbb{P}_{\widetilde{\mu}}\left[\sum_{i=2}^{l-1} N_i(C) \ge \gamma n | C \in \mathcal{C}_t\right] \mathbb{P}_{\widetilde{\mu}}[C \in \mathcal{C}_t]$.

By definition, if $C \in \mathcal{C}_t$, then $\sum_{i=2}^{l-1} N_i(C) \le t/2$ since each cycle is of length at least 2. Hence, $\mathbb{P}_{\widetilde{\mu}}\left[\sum_{i=2}^{l-1} N_i(C) \ge \gamma n | C \in \mathcal{C}_t\right] = 0$ if $t < 2\gamma n$. Hence, the only terms that survive are those for which $t \ge 2\gamma n$. Thus, $\mathbb{P}_{\widetilde{\mu}}\left[\sum_{i=2}^{l-1} N_i(C) \ge \gamma n\right] = \sum_{t=2\gamma n}^n \mathbb{P}_{\widetilde{\mu}}\left[\sum_{i=2}^{l-1} N_i(C) \ge \gamma n | C \in \mathcal{C}_t\right] \mathbb{P}_{\widetilde{\mu}}[C \in \mathcal{C}_t]$. The latter is at most

$$\sum_{t=2\gamma n}^n \mathbb{P}_{\widetilde{\mu}}[C \in \mathcal{C}_t] \le \sum_{t=2\gamma n}^n \left(\mathbb{P}_\mu[C \in \mathcal{C}_t] + 1/(2n^2)\right)$$

where we have moved from $\widetilde{\mu}$ to μ and paid in the distance between them, see Remark II.3. Thus,

$$\sum_{t=2\gamma n}^n \mathbb{P}_{\widetilde{\mu}}[C \in \mathcal{C}_t] \le 1/(2n) + \sum_{t=2\gamma n}^n \mathbb{P}_\mu[C \in \mathcal{C}_t].$$

We proceed to calculate $\mathbb{P}_\mu[C \in \mathcal{C}_t]$ for $t \ge 2\gamma n$.

Using Lemma II.4, if $k > l^l$, $\mathbb{P}_\mu[C \in \mathcal{C}_t]$ can be upper bounded by $\binom{n}{t} k^{(1-1/l)t} k^{n-t} \max_{C \in \mathcal{C}_t} \mu(C)$ by first fixing a subset of t vertices, then using the bound of $k^{(1-1/l)t}$ on the number of ways to cover the t vertices by vertex-disjoint cycles of length less than l and, finally, using the trivial upper bound of k^{n-t} on the number of ways to cover the rest of the vertices. Note that $\max_{C \in \mathcal{C}_t} \mu(C)$ is at most $\frac{2^{n/3}}{\operatorname{per}(A)}$. Since $\operatorname{per}(A) \ge (k/e)^n$, we obtain that

$$\mathbb{P}_\mu[C \in \mathcal{C}_t] \le \exp(n + t\ln(ne/t) - t(\ln k)/l + (n \ln 2)/3).$$

Thus, if $t \ge 2\gamma n$ then, provided $(\ln k)/l > \ln(e/\gamma)$, $\mathbb{P}_\mu[C \in \mathcal{C}_t] \le \exp(n(1 + \gamma\ln(e/\gamma(-\gamma(\ln k)/l + (\ln 2)/3)) \le \exp(-n\gamma[(\ln k)/l - 4/\gamma])$. Here we have used the fact that

$$1/\gamma \cdot (2 + (\ln 2)/3) + \ln 1/\gamma \le 4/\gamma$$

which holds since $\gamma \le 1$. Thus, if $(\ln k)/l - 4/\gamma > 0$ for fixed k, l and γ, the r.h.s. goes down exponentially with n. Consequently, for a fixed k, l, γ, for all large enough n, $\mathbb{P}_\mu[C \in \mathcal{C}_t]$ can be made to be smaller than $\gamma/(2n^2)$. Hence,

$$\sum_{t=2\gamma n}^n \mathbb{P}_\mu[C \in \mathcal{C}_t] \le (\gamma/(2n^2)) \cdot n \le \gamma/(2n) \le 1/(2n).$$

Thus, $\sum_{t=2\gamma n}^n \mathbb{P}_{\widetilde{\mu}}[C \in \mathcal{C}_t] \le 1/(2n) + 1/(2n) = 1/n$. Hence, then with probability at least $1 - 1/n$, the cost of the Eulerian tour derived from C is at most $n + 2\gamma n + 2n/l \le n(1 + 2\gamma + 2/l)$.

The constraints on k, l, γ are $k \ge l^l$, $(\ln k)/l > \ln(e/\gamma)$ and $(\ln k)/l > 4/\gamma$. The second constraint is strictly weaker than the third. The choice of $\gamma \stackrel{\text{def}}{=} \delta/4$, and $l \stackrel{\text{def}}{=} 4/\delta$ and $k \ge \exp(64/\delta^2)$ satisfies these constraints and gives a tour of length at most $(1 + \delta)n$ with high probability. As a function of k, δ can be chosen as $\sqrt{64/\ln k}$. This concludes the proof. ■

We have made no attempt to optimize the dependence of δ on k. The fact that our bipartite graphs are *symmetric*, i.e., there is an ordering of the vertices on both sides such that (v_i^L, v_j^R) is an edge if and only if (v_j^L, v_i^R) is an edge, may also be useful in order to tighten our estimates.

How critical is regularity?: The lower bounds used on the permanent hold only in the doubly stochastic case and, hence, the graph must be regular, perhaps with multiple edges and self loops allowed, if one intends to use these bounds. However, regularity can be mildly relaxed. There are at least two ways to do this: (1) Consider a graph with minimum degree k (which can be made large), but no upper bound on the degree. Consider the degree reduction operation which takes a high degree vertex v and replaces it with vertices $v_1, \ldots, v_t$, distributes the edges adjacent to v among $v_1, \ldots, v_t$ and puts edges among them so as to make them k-regular. An example of such an operation was used, for instance, by [25] to reduce a graph to a cubic graph. Note that regularizing the graph in this manner may increase the number of vertices significantly. The good news is that an Eulerian tour on this bigger, but regular, graph has the same cost as in the original irregular graph. Hence, if the input graph is such that it can be made regular by adding a small number of vertices, our results will hold. (2) If the graph is nearly regular and very small sets in G expand, then it can be shown via a max-flow/min-cut technique (see [23], [3]) that there is a regular subgraph in G and, hence, our results also hold. We omit the details. As a first step, it would be interesting to extend our results to graphs where all degrees are between $k/2$ and k.

III. Extensions and Discussion

Before we discuss our approach and its extensions to MetricTSP, it may be helpful to first review Christofides' algorithm and the recent approach in [15]. Henceforth, the graph will not be assumed to be k-regular unless stated explicitly.

Matchings and trees are central to algorithm design for TSP. Indeed, two efficiently computable lower bounds on the cost of the best tour are: (1) the minimum cost spanning tree in G and (2) two times the minimum cost maximum matching in G. Christofides' algorithm uses both to obtain his $3/2$ approximation algorithm: He starts with a minimum spanning tree and adds a minimum cost perfect matching on the vertices in the tree which have odd degree to make the union of the two Eulerian. Further, the cost of this Eulerian tour is no more than $3/2$ times that of the optimal tour. Once the graph is Eulerian, the metric property allows one to short circuit and obtain a tour without increasing the cost.

The algorithm in [15] brought the approximation ratio below $3/2$ for GraphTSP for the first time. Roughly, the key new idea in [15] was to start with a random spanning tree instead of a minimum spanning tree. The distribution on spanning trees is chosen such that it maximizes entropy while maintaining that an edge is present with probability $x^\star_e$, where $x^\star_e$ is obtained by solving the following Held-Karp LP relaxation to TSP: Minimize $\sum_{e\in E} c_e x_e$ subject to

$$x(\delta(v)) = 2 \quad \forall v \in V,\, x(\delta(S)) \geq 2 \quad \forall\, \emptyset \neq S \subsetneq V$$

with $x \geq 0$. Here $\delta(S)$ denotes the set of edges crossing $(S, \bar{S})$ and $x(F) \stackrel{\text{def}}{=} \sum_{e\in F} x_e$ for any $F \subseteq E$. A significant amount of effort goes into proving that this, or rather a modification of this algorithm, gives a factor strictly less than $3/2$. Notably, since random spanning trees are intimately related to the combinatorial Laplacian and the probability that an edge is present is proportional to its *effective resistance*, there is an intrinsic connection between their approach and determinants.

We now describe two approaches that arise from our work which may provide extended results. First, given a graph G, solve the Held-Karp LP mentioned above and obtain an optimal solution $x^\star$. Notice that for every vertex v in G, $x^\star(\delta(v)) = 2$. Hence, if we consider the $n \times n$ matrix A where the entry for the edge $e = uv$ is set to $x^\star_e/2$, then A is doubly stochastic. The algorithm then constructs the bipartite version of G, G' as before, albeit with the weight of an edge uv being $x^\star_{uv}/2$ and samples a random perfect matching from G'. The difference is that now a matching is picked with probability proportional to the product of the weight of its edges. This latter task can be handled by the algorithm in [19]. Now, one can pick a minimum cost tree to connect all the cycles and double it. The challenge in analyzing this approach is to handle edges which are close to 1 in the Held-Karp solution. Another issue is to relate the probability that an edge appears in a perfect matching according to the distribution above to $x^\star_e$. [3]

The second approach can be interpreted as a dual to the approach in [15]; the role of matching and trees is reversed. Here, starting with a graph G and its Held-Karp solution $x^\star$, compute non-negative weights λ_e for edges such that the probability that an edge e is contained in a random perfect matching of the bipartite graph G' is $x^\star_e$. In this case, the expected cost of the cycles produced is exactly the cost of the Held-Karp solution. Here, the difficulty seems to be to analyze the cost of connecting the cycles. Finally, it should be of interest to evaluate these permanent-based heuristics on data-sets arising in practice.

In conclusion, in this paper we show that the shortest path metrics obtained from simple k-regular graphs become easier instances of GraphTSP as k increases. This substantially extends the recent results on cubic graphs and suggests that the hard instances for GraphTSP may occur when the average degree is low.

Acknowledgment

The author would like to thank Deeparnab Chakrabarty and Mohit Singh for valuable discussions.

References

[1] Nishita Aggarwal, Naveen Garg, and Swati Gupta. A 4/3-approximation for TSP on cubic 3-edge-connected graphs. *CoRR*, abs/1101.5586, 2011.

[2] M. Ajtai, J. Komlós, and E. Szemerédi. First occurrence of Hamilton cycles in random graphs. In B.R. Alspach and C.D. Godsil, editors, *Annals of Discrete Mathematics, 27, Cycles in Graphs*, volume 115 of *North-Holland Mathematics Studies*, pages 173 – 178. North-Holland, 1985.

[3] Noga Alon, Vojtech Rödl, and Andrzej Rucinski. Perfect matchings in ε-regular graphs. *Electr. J. Comb.*, 5, 1998.

[4] Dana Angluin and Leslie G. Valiant. Fast probabilistic algorithms for Hamiltonian circuits and matchings. In *Proceedings of the ninth annual ACM symposium on Theory of computing*, STOC '77, pages 30–41, New York, NY, USA, 1977.

[3] Intriguingly, the *solution* we use, $2/k$ for every edge, for k-regular graphs *does not* satisfy the Held-Karp LP as there may be very small cuts in the input graph G which are not picked to be 1 in this solution.

[5] Sanjeev Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *J. ACM*, 45(5):753–782, 1998.

[6] Sanjeev Arora, Michelangelo Grigni, David R. Karger, Philip N. Klein, and Andrzej Woloszyn. A polynomial-time approximation scheme for weighted planar graph TSP. In *SODA*, pages 33–41, 1998.

[7] Sylvia Boyd, René Sitters, Suzanne van der Ster, and Leen Stougie. TSP on cubic and subcubic graphs. In *IPCO*, pages 65–77, 2011.

[8] Andrei Z. Broder, Alan M. Frieze, and Eli Shamir. Finding hidden Hamiltonian cycles (extended abstract). In *STOC*, pages 182–189, 1991.

[9] Nicos Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical Report 388, Graduate School of Industrial Administration, Carnegie-Mellon University, 1976.

[10] Erik D. Demaine, MohammadTaghi Hajiaghayi, and Kenichi Kawarabayashi. Contraction decomposition in H-minor-free graphs and algorithmic applications. In *STOC*, pages 441–450, 2011.

[11] G. A. Dirac. Some theorems on abstract graphs. *Proc. London Math. Soc.*, 2:69–81, 1952.

[12] G. P. Egorychev. Proof of the van der Waerden conjecture for permanents. *Siberian Mathematical Journal*, 22:854–859, 1981.

[13] D. I. Falikman. Proof of the van der Waerden conjecture regarding the permanent of a doubly stochastic matrix. *Mathematical Notes*, 29:475–479, 1981. 10.1007/BF01163285.

[14] David Gamarnik, Moshe Lewenstein, and Maxim Sviridenko. An improved upper bound for the TSP in cubic 3-edge-connected graphs. *Oper. Res. Lett.*, 33(5):467–474, 2005.

[15] Shayan Oveis Gharan, Amin Saberi, and Mohit Singh. A randomized rounding approach to the traveling salesman problem. In *FOCS*, pages 550–559, 2011.

[16] Michelangelo Grigni, Elias Koutsoupias, and Christos H. Papadimitriou. An approximation scheme for planar graph TSP. In *FOCS*, pages 640–645, 1995.

[17] Michael Held and Richard M. Karp. The traveling-salesman problem and minimum spanning trees. *Operations Research*, 18:1138–1162, 1970.

[18] Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM J. Comput.*, 18(6):1149–1178, 1989.

[19] Mark Jerrum, Alistair Sinclair, and Eric Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries. *J. ACM*, 51(4):671–697, 2004.

[20] R. M. Karp. The Probabilistic Analysis of some Combinatorial Search Algorithms. In J. F. Traub, editor, *Algorithms and Complexity: New Directions and Recent Results*, pages 1–20. Academic Press, New York, 1976.

[21] Philip N. Klein. A linear-time approximation scheme for TSP in undirected planar graphs with edge-weights. *SIAM J. Comput.*, 37(6):1926–1952, 2008.

[22] János Komlós and Endre Szemerédi. Limit distribution for the existence of Hamiltonian cycles in a random graph. *Discrete Mathematics*, 43(1):55–63, 1983.

[23] László Lovász and Michael D. Plummer. *Matching Theory*. AMS Chelsea Publishing, vol. 367, 2009.

[24] Joseph S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems. *SIAM J. Comput.*, 28(4):1298–1309, 1999.

[25] Tobias Mömke and Ola Svensson. Approximating graphic TSP by matchings. In *FOCS*, pages 560–569, 2011.

[26] Marcin Mucha. $13/9$-approximation for graphic TSP. In *STACS*, pages 30–41, 2012.

[27] Christos H. Papadimitriou and Santosh Vempala. On the approximability of the traveling salesman problem (extended abstract). *Combinatorica*, 26(1):101–120, 2006.

[28] Robert W. Robinson and Nicholas C. Wormald. Hamilton cycles containing randomly selected edges in random regular graphs. *Random Struct. Algorithms*, 19(2):128–147, 2001.

[29] Alexander Schrijver. Counting 1-factors in regular bipartite graphs. *J. Comb. Theory, Ser. B*, 72(1):122–135, 1998.

[30] András Sebö and Jens Vygen. Shorter tours by nicer ears:. *CoRR*, abs/1201.1870, 2012.