SEPARATING SUBADDITIVE EUCLIDEAN FUNCTIONALS

ALAN FRIEZE AND WESLEY PEGDEN

ABSTRACT. If we are given n random points in the hypercube $[0,1]^d$, then the minimum length of a Traveling Salesperson Tour through the points, the minimum length of a spanning tree, and the minimum length of a matching, etc., are known to be asymptotically $\beta n^{\frac{d-1}{d}}$ a.s., where β is an absolute constant in each case. We prove separation results for these constants. In particular, concerning the constants β^d_{TSP}, β^d_{MST}, β^d_{MM}, and β^d_{TF} from the asymptotic formulas for the minimum length TSP, spanning tree, matching, and 2-factor, respectively, we prove that $\beta^d_{\mathrm{MST}} < \beta^d_{\mathrm{TSP}}$, $2\beta^d_{\mathrm{MM}} < \beta^d_{\mathrm{TSP}}$, and $\beta^d_{\mathrm{TF}} < \beta^d_{\mathrm{TSP}}$ for all $d \geq 2$. Our results have some computational relevance, showing that a certain natural class of simple algorithms cannot solve the random Euclidean TSP efficiently.

1. Introduction

Beardwood, Halton, and Hammersley [3] studied the length of a Traveling Salesperson Tour through random points in Euclidean space. In particular, if $x_1, x_2, \ldots$ is a random sequence of points in $[0,1]^d$ and $\mathcal{X}_n = \{x_1, \ldots, x_n\}$, their results imply that there is an absolute constant β^d_{TSP} such that the length $\mathrm{TSP}(\mathcal{X}_n)$ of a minimum length tour through $\mathcal{X}_n$ satisfies

$$\mathrm{TSP}(\mathcal{X}_n) \sim \beta^d_{\mathrm{TSP}} n^{\frac{d-1}{d}} \qquad a.s. \tag{1}$$

This result has many extensions; for example, we know that identical asymptotic formulas hold for the the cases of the minimum length of a spanning tree $\mathrm{MST}(\mathcal{X}_n)$[3], and the minimum length of a matching $\mathrm{MM}(\mathcal{X}_n)$ [13]. Steele [14] provided a general framework which enables fast assertion of identical asymptotic formulas for these and other suitable problems. For example, we will see in Section 2 that his results imply that the length $\mathrm{TF}(\mathcal{X}_n)$ of a minimum length 2-factor admits the same asymptotic characterization.

A major problem in this area remains to obtain analytic results regarding the constants β in such formulas. In particular, the analytic bounds on such constants are generally very weak, with the best known results given for $d = 2$ in Table 1. On the other hand, there was some success as d grows large, as Bertsimas and Van

Date: January 13, 2015.
Research supported in part by NSF grant DMS-1362785.
Research supported in part by NSF grant DMS-1363136.

	lower	upper
β^2_{TSP}	.62499 [3]	.92037 [3]
β^2_{MST}	.60082 [2]	$\frac{1}{\sqrt{2}} \approx .707$ [9]
$2\beta^2_{\mathrm{MM}}$	.5 [6]	.92037

TABLE 1. Bounds on constants for $d = 2$.

Ryzen [6] showed that, asymptotically in d,

$$\beta^d_{\mathrm{MST}} \sim 2\beta^d_{\mathrm{MM}} \sim \sqrt{\frac{d}{2\pi e}}, \tag{2}$$

and conjectured that $\beta^d_{\mathrm{TSP}} \sim \sqrt{\frac{d}{2\pi e}}$ as well.

It seems that it has been overlooked that local geometric arguments are sufficient to prove the separation of constants for many natural examples of Euclidean functionals. In particular, in the present paper, we will show that $\beta^d_{\mathrm{MST}} < \beta^d_{\mathrm{TSP}}$, $\beta^d_{\mathrm{TF}} < \beta^d_{\mathrm{TSP}}$, and $2\beta^d_{\mathrm{MM}} < \beta^d_{\mathrm{TSP}}$ for all d. These are the first asymptotic separations for Euclidean functionals where the Eulidean metric is playing an essential role: the only previous separation was shown (by Bern [4]; see also [10]) for the minimum length rectilinear Steiner tree vs. the minimum rectilinear length spanning tree, which is equivalent to asymptotically distinguishing Steiner trees from trees in the L_1 norm. (The rectilinear Steiner tree is also the only case where the asymptotic *worst-case* length is known exactly [5]).

We begin by considering the degrees of vertices in the minimum spanning trees among n random points. Steele, Shepp, and Eddy [16] showed that the number $\Lambda_k(\mathcal{X}_n)$ of vertices of degree k satisfies

$$\Lambda_k(\mathcal{X}_n) \sim \alpha_{k,d} n$$

for constants $\alpha_{k,d}$, and proved that $\alpha(1,d) > 0$. Note that we must have $\alpha_{k,d} = 0$ when $k > \tau(d)$, where $\tau(d)$ is the kissing number of d dimensional space (6 in the case $d = 2$). Indeed, we must have $\alpha_{k,d} = 0$ whenever $k > \tau'(d)$, where $\tau'(d)$ denotes a *strict kissing number* of d, which we define as the maximum K such that there exists $\varepsilon > 0$ such that there is, in d dimensions, a configuration of K disjoint spheres of radius $1 + \varepsilon$ each tangent to a common unit sphere. (Note that $\tau'(d) \le \tau(d)$, and in particular, $\tau'(2) = 5$.) We prove:

Theorem 1.1. *$\alpha(k,d) > 0$ if $k \le \tau'(d)$.*

Considering Euclidean functionals $\mathrm{MST}_k(X)$ (with corresponding constants $\beta^d_{\mathrm{MST}_k}$) defined as the minimum length of a spanning tree of X whose vertices all have degree $\le k$, we will then get separation as follows:

Theorem 1.2. *We have that*

$$\beta^d_{\mathrm{TSP}} = \beta^d_{\mathrm{MST}_2} > \beta^d_{\mathrm{MST}_3} > \cdots > \beta^d_{\mathrm{MST}_{\tau'(d)}} = \beta^d_{\mathrm{MST}} \tag{3}$$

for all d.

Thus, the MST_k constants are as diverse as are allowed by the simple geometric constraint of $\tau'(d)$.

Still, there are only finitely many constants $\beta^d_{\mathrm{MST}_k}$ for each d; while we can draw trees with very large degrees, large degrees (relative to d) are not useful for minimum spanning trees in Euclidean space. In contrast to this scenario, let us recall that a *2-factor* is a disjoint set of cycles covering a given set of points. We will see in Section 2 that the length of the minimum 2-factor is indeed a subadditive Euclidean functional, and thus this length satisfies $\mathrm{TF}(\mathcal{X}_n) \sim \beta^d_{\mathrm{TF}} n^{\frac{d-1}{d}}$ for some constant β^d_{TF}. Moreover, if $\mathrm{TF}_g(X)$ is the minimum length of a 2-factor through X whose cycles all have length $\geq g$, then we will see that TF_g is also a subadditive linear functional, so that we have $\mathrm{TF}_g(\mathcal{X}_n) \sim \beta^d_{\mathrm{TF}_g} n^{\frac{d-1}{d}}$. Naturally, we must have $\beta^d_{\mathrm{TF}} = \beta^d_{\mathrm{TF}_3} \leq \beta^d_{\mathrm{TF}_4} \leq \beta^d_{\mathrm{TF}_5} \leq \cdots$. In analogy to the high-degree vertices in a tree, we can of course draw 2-factors with small cycles, but it is not clear *a priori* whether small cycles will be asymptotically essential to optimum 2-factors in random point sets. The following theorem shows that they are:

Theorem 1.3. *β_{TF_g} is a monotone increasing sequence $\beta^d_{\mathrm{TF}_3} < \beta^d_{\mathrm{TF}_4} < \beta^d_{\mathrm{TF}_5} < \cdots$.*

On the other hand, we prove that 2-factors with long (but constant) girth requirements produce close approximations to the TSP:

Theorem 1.4. $\lim_{g\to\infty} \beta^d_{\mathrm{TF}_g} = \beta^d_{\mathrm{TSP}}$.

With a bit more work, our method for proving Theorem 1.3 will also allow us to deduce the following:

Theorem 1.5. $2\beta^d_{\mathrm{MM}} < \beta^d_{\mathrm{TSP}}$.

We note in contrast to Theorem 1.3 that in the independent case where the edge lengths $X_e, e \in \binom{[n]}{2}$ are independent uniform $[0,1]$ random variables, Frieze [8] showed that that weight of the minimal 2-factor is asymptotically equivalent to the minimum length tour, with probability $1 - o(1)$.

We continue by mentioning a natural generalization of $\mathrm{MM}(\mathcal{X}_n)$. Given a fixed graph H on k vertices, an H-factor of a set of points S is a set of edges isomorphic to $\lfloor |X|/k \rfloor$ vertex disjoint copies of H. As a subadditive Euclidean functional, the minimum length $H\mathrm{F}(\mathcal{X}_n)$ of an H factor of $\mathcal{X}_n$ satisfies

$$H\mathrm{F}(\mathcal{X}_n) \sim \beta^d_H n^{\frac{d-1}{d}}.$$

We pose the following conjecture:

Conjecture 1.6. *Given H_1, H_2 and $d \geq 2$, we have that $\beta^d_{H_1} \neq \beta^d_{H_2}$ unless H_1 and H_2 are each isomorphic to a disjoint union of copies of some graph H_3. In particular, $\beta^d_{H_1} \neq \beta^d_{H_2}$ if H_1, H_2 are connected and non-isomorphic.*

We prove at least the following, showing diversity in the constants even for fixed edge density:

Theorem 1.7. *For any fixed $d \geq 2$ and rational $r \geq 1$, there are infinitely many distinct constants β_H^d over connected graphs H with edge density $\frac{|E(G)|}{|V(G)|} = r$.*

Our separation results have implications for the practical problem of solving the Euclidean TSP. *Branch and bound* algorithms are a standard approach to solving NP-hard problems, in which a bounding estimate is used to prune an exhaustive search of the solution space. There has been a great deal of success solving real-world instances of the TSP with branch-and-bound augmented with sophisticated techniques based on cutting planes for the TSP polytope (see, for example Applegate, Bixby, Chvátal and Cook [1]).

One simple and natural lower bound for the TSP is the minimum length 2-factor, and one might think that this bound would suffice to solve random instances of the Euclidean TSP with branch and bound efficiently. However, the separation of constants and the concentration of measure shows that this is not necessarily true, even if one could use 2-factors of large girth (though finding the minimum length 2-factor of girth $g \geq 4$ is known to become NP-hard for $g \geq 4$). In particular, in Section 5 we will define for absolute constants, C, δ, a (C, δ)-restricted branch and bound algorithm. This class of algorithms includes many naturally occurring variants, and we will prove:

Theorem 1.8. *Suppose that we use the 2-factor problem, with an arbitrarily large constant lower bound g on girth, to give us a lower bound for use in (C, δ)-restricted branch and bound algorithm to solve the Euclidean TSP. Then the algorithm runs in time $n^{\Omega(n^{(d-1)/d})}$, a.s.*

This gives a rigorous explanation for the observation (see [12], for example) that branch-and-bound heuristics using the Assignment Problem as a bounding estimate (even weaker than the 2-factor) perform poorly on random Euclidean instances.

2. Subadditive Euclidean Functionals

Steele defined a *Euclidean functional* as a real valued function L on finite subsets of $\mathbb{R}^d$ which is invariant under translation, and scales as $L(\alpha X) = \alpha L(X)$. It is *nearly monotone* with respect to addition of points if

$$L(X \cup Y) \geq L(X) - o(n^{\frac{d-1}{d}}) \quad \text{for } n = |X|. \tag{4}$$

It has finite variance if, fixing n, we have

$$\mathbf{Var}(L(\mathcal{X}_n)) < \infty \tag{5}$$

(in particular, if it is bounded for fixed n) and it is *subadditive* if, for $\mathcal{Y}_n$ a random set of n points from $[0, t]^d$, it satisfies

$$L(\mathcal{Y}_n) \leq \sum_{\alpha \in [m]^d} L(S_\alpha \cap \mathcal{Y}_n) + Ctm^{d-1}$$

for some absolute constant C, where here $\{S_\alpha\}$ ($\alpha \in [m]^d$) is a decomposition of $[0, t]^d$ into m^d subcubes of side length $u = t/m$.

Steele proved:

Theorem 2.1 (Steele [14])**.** *If L is a subadditive Euclidean functional on $\mathbb{R}^d$ of finite variance, $x_1, x_2, \dots$ is a random sequence of points from $[0,1]^d$, and $\mathcal{X}_n = \{x_1, x_2, \dots, x_n\}$, then there is an absolute constant β_L^d s.t.*

$$L(\mathcal{X}_n) \sim \beta_L^d n^{\frac{d-1}{d}} \qquad a.s.$$

This can thus be used to easily give the existence of the simple asymptotic formulas for the functionals $\mathrm{TF}_g(X)$, $\mathrm{MST}_k(X)$, and $H\mathrm{F}(X)$ by showing that these functionals are subadditive.

Proposition 2.2. $\mathrm{TF}_g(X)$*,* $\mathrm{MST}_k(X)$*, and* $H\mathrm{F}(X)$ *are subadditive Euclidean functionals.*

Before writing a proof, we note that for the definition of the 2-factor functionals $\mathrm{TF}_g(X)$, we can only require that the 2-factors whose length we minimize cover all the points when there are at least $\max(g,3)$ points. Similarly, the $H\mathrm{F}(X)$ functional is required just to cover at least $n - |H| + 1$ points.

Proof. We begin by noting that for each of these functionals, we can assert an upper bound $Cn^{\frac{d-1}{d}}$ for some constant C, even over worst-case arrangements of n points in $[0,1]^d$. The analogous statement for the TSP was proved by Toth [17] and by Few [7], and implies these bounds for the functionals considered here. Indeed, a tour through n points itself gives a tree of max-degree 2 (after deleting one edge), and is a 2-factor subject to any constant girth restriction. For H factors, a tour can be divided into paths of length $|H|$ (except for $< |H|$ remaining vertices) which can then be completed to instances of $H' \supseteq H$ by adding edges. Each added edge has a cost bounded by the length of the path it lies in and so this construction increases the total cost by at most a factor equal to the number of edges in H.

Subadditivity of $\mathrm{TF}_g(X)$, and $H\mathrm{F}(X)$ is now a consequence of the fact that a union of 2-factors (subject to restrictions on the cycle length, perhaps) or H-factors is again a 2-factor (subject to the same restrictions) or an H factor, respectively. In particular, the subadditive error term for these functions comes just from the fact that points may be uncovered in some of the subcubes S_α, for the exceptional reasons noted above. Since there are at most $(g-1)m^d$ or $|H|m^d$ such uncovered points, however, the error is suitably bounded by the minimum cost factor on a worst-case arrangement of the remaining points.

Subadditivity of $\mathrm{MST}_k(X)$ $(k \geq 2)$ is similar: after finding minimal spanning trees of max-degree k in each subcube S_α, we must join together these trees into a single tree. We choose 2 leaves of each subcube's tree and denote one red and the other blue. We let $\alpha_1, \alpha_2, \dots, \alpha_{m^d}$ denote a path through the decomposition $\{S_\alpha\}$, so that the subcubes S_{α_i} and $S_{\alpha_{i+1}}$ are adjacent. For each $i < m^d$, we join the red leaf of the tree in S_{α_i} to the blue leaf in the tree of S_{α_i+1}. The result is a spanning tree of the whole set of points with the same maximum degree and with extra cost at most $2\sqrt{d}um^d = 2\sqrt{d}tm^{d-1}$. □

3. Separating asymptotic constants

In the following we will use the simplest application of the Azuma-Hoeffding martingale tail inequality: It is often referred to as McDiarmid's inequality [11]. Suppose that we have a random variable $Z = Z(X_1, X_2, \ldots, X_N)$ where $X_1, X_2, \ldots, X_N$ are independent. Further, suppose that changing one X_i can only change Z by at most c in absolute value. Then for any $t > 0$,

$$\mathbf{Pr}(|Z - \mathbf{E}\, Z| \geq t) \leq 2 \exp\left\{-\frac{t^2}{c^2 N}\right\}. \tag{6}$$

We will also use the following inequality, applicable under the same conditions, when $\mathbf{E}\, Z$ is not large enough.

$$\mathbf{Pr}(Z \geq \alpha\, \mathbf{E}\, Z) \leq \left(\frac{e}{\alpha}\right)^{\alpha\, \mathbf{E}\, Z/c}. \tag{7}$$

Our method to distinguish constants is based on achieving constant factor improvements to the values of functions via local changes. Given $\varepsilon, D \in \mathbb{R}$ and a finite set of points $S \subseteq \mathbb{R}^d$ and a universe X, we say that $T \subseteq X$ is an (ε, D)-copy of S if there is a bijection f between T and a point set S' congruent to S such that $||x - f(x)|| < \varepsilon$ for all $x \in T$, and such that T is at distance $> D$ from $X \setminus T$. Here we will further assume that $||x - y|| > \varepsilon$ for $x \neq y \in S$.

For our purposes, it will be convenient notationally to work with n random points $\mathcal{Y}_n$ from $[0, t]^d$ where $t = n^{1/d}$, in place of n random points $\mathcal{X}_n$ from $[0, 1]^d$. At the end, we will scale our results by a factor $n^{-1/d}$ in order to get what is claimed above.

Observation 3.1. *Given any finite point set S, any $\varepsilon > 0$, and any D, $\mathcal{Y}_n$ a.s contains at least $C^S_{\varepsilon,D} n$ (ε, D)-copies of S, for some constant $C^S_{\varepsilon,D} > 0$.*

Proof of Observation 3.1. Let Z denote the number of (ε, D)-copies of S in $\mathcal{Y}_n$. We divide $[0, t]^d$ into $n/(3D)^d$ subcubes $C_1, C_2, \ldots,$ of side $3D$. Then let $C'_i \subseteq C_i$ be a centrally placed subcube of side D. Now choose a set S' congruent to S somewhere inside C'_1 and let $B_1, B_2, \ldots, B_s$, $s = |S|$ be the collection of balls of radius ε, centered at each point of S'. The with probability at least $\alpha = \alpha_{\varepsilon,D} > 0$, each B_i contains exactly one point of $\mathcal{Y}_n$ and there are no other points of $\mathcal{Y}_n$ in C_1. Thus $\mathbf{E}\, Z \geq \beta n$ where $\beta = \alpha/(3D)^d$. Now changing the position of one point in $\mathcal{Y}_n$ changes the number of (ε, D)-copies of S by at most two and so we can use McDiarmid's inequality [11] to show that $Z \geq \frac{1}{2}\mathbf{E}\, Z$ a.s. □

To use this to prove Theorem 1.1, we will need just a bit more.

Observation 3.2. *If $Y \subset \mathbb{R}^d$ and x lies in the interior of the convex hull of Y, then when D is sufficiently large, any point at distance $> D$ is closer to some point of Y than to x.* □

If $v_0, v_1, \ldots, v_k$ are vectors in $\mathbb{R}^d$ with pairwise negative dot-product, then $v_1, \ldots, v_k$ lie in the half-space $v_0 \cdot x < 0$, and the projections of $v_1, \ldots, v_k$ onto the hyperplane

$v_0 \cdot x = 0$ have pairwise negative dot-products. This gives the following, by induction on d:

Observation 3.3. *If $v_1, \ldots, v_{d+1} \in \mathbb{R}^d$ are vectors with negative pairwise dot-products, then 0 is a positive linear combination of the v_i's.* □

This allows us to prove:

Lemma 3.4. *If $d+1 \leq k \leq \tau'(d)$, then there exists a set of points $\bar{S}^{(k)} \subset \mathbb{R}^d$ consisting of a single point at the origin, surrounded by a set $S^{(k)}$ of k points on the unit sphere centered at the origin and separated pairwise by at least some $\varepsilon' > 0$ more than unit distance, such that $S^{(k)}$ does not lie in open half-space whose boundary passes through the origin.*

Proof. We first observe that the definition of τ' already gives us a set $S^{(k)}$ with the desired properties, except that it may all lie in some open half-space through the origin. In this case, however, we can delete a point and replace it with the point x_H on the unit sphere opposite the half-space H, and furthest away from the halfspace. We do this repeatedly and note that because the above exchange of points only happens when all points are on one side of a half-space H', x_H remains as the unique point which is in the open half-space opposite to H. Furthermore, doing this repeatedly, we can achieve either a set $S^{(k)}$ with all the desired properties, or can find after at most k steps a set $S^{(k)}$ of points on the sphere separated pairwise by at least $\varepsilon' > 0$ more than unit distance, and whose pairwise dot products as vectors in $\mathbb{R}^d$ are all negative. But then Observation 3.3 and $k \geq d+1$ implies that the points cannot all lie in the interior of some half-space whose boundary passes through the origin. □

We are now ready to prove Theorem 1.1.

Proof of Theorem 1.1. Given $k \geq 2$, we choose any $d' \leq d$ such that $d'+1 \leq k \leq \tau'(d')$.

We apply Lemma 3.4 with k, d' to get a set $S'^{(k)} \subset \mathbb{R}^d$. Observe first that the origin must lie in the convex hull X of the set $S'^{(k)}$ given by Lemma 3.4; otherwise, there would be a supporting half-space H of X not containing the origin, and $S'^{(k)}$ would lie in the open half-space through the origin which is parallel to H, a contradiction. Now we take $S^{(k)} = S'^{(k)} \times \{0\}^{d-d'}$, and the origin is still in the convex hull of $S^{(k)}$.

Now, letting Δ^d denote a unit simplex centered at the origin (with $d+1$ points), we let

$$U = \bar{S}^{(k)} \cup \bigcup_{p \in S^{(k)}} \{(1.5)p + .1 \cdot \Delta^d\}.$$

So U is a set of $1+k+(d+1)k$ points. (Figure 1 shows U for the case $d=2, k=2$; note that in this case, $d'=1$.)

FIGURE 1. A configuration for forcing degree 2 in 2-dimensions.

We now let $U_{\varepsilon,D}$ denote an (ε, D) copy of U, for sufficiently small $\varepsilon > 0$ and sufficiently large D. Observe that since the origin is in the convex hull of $S^{(k)}$, the k small copies of the d-simplex in U ensure that the origin is in the *interior* of the convex hall of U, and thus also in the interior of $U_{\varepsilon,D}$ for sufficiently small ε.

Observe that (for large D) the distance between any pair of points in an $U_{\varepsilon,D}$ is less than the minimum distance between $U_{\varepsilon,D}$ and $\mathcal{Y}_n \setminus U_{\varepsilon,D}$. In particular, if T denotes the minimum length spanning tree on $\mathcal{Y}_n$, the subgraph $T[U_{\varepsilon,D}]$ induced by the points in $U_{\varepsilon,D}$ must be connected (and so a tree), or we could exchange a long edge for a short edge. Moreover, the minimum length spanning tree on T must restrict to a minimum length spanning tree on $U_{\varepsilon,D}$, and by construction, the point of $U_{\varepsilon,D}$ corresponding to the origin point in U has degree k in the MST on U. Finally, no points in $\mathcal{Y}_n \setminus U_{\varepsilon,D}$ can be adjacent to the center of the star when D is sufficiently large, by Observation 3.2. Thus Observation 3.1 gives that $\alpha_{k,d} > 0$ for $d + 1 \leq k \leq \tau'(d)$.

Finally, $\alpha_{1,d} > 0$ is an immediate consequence of $\alpha_{3,d} > 0$. □

Indeed, Theorem 1.2 follows immediately as well:

Proof of Theorem 1.2. Suppose $2 \leq k < \tau'(d)$, and T is a minimum spanning tree of $\mathcal{Y}_n$ subject to the restriction that the maximum degree is $\leq k$. By Observation 3.1 we have that there are Cn (ε, D) copies of the set U from the previous proof, for some constant C, and from the argument above we see that each such copy S_i will induce a (connected subtree) $T[S_i]$, which will have maximum degree at most k in an instance of MST_k. Replacing each $T[S_i]$ by the optimum $(k+1)$-star produces a spanning tree of maximum degree $k + 1$, whose length is less by at least some constant $C'n$. Rescaling by t gives that the length difference is at least $C'n^{\frac{d-1}{d}}$. □

Remark 3.5. The same argument allows us to separate β^d_{MST} from $\beta^d_{Steiner}$ where the latter corresponds to the minimum length Steiner tree. We just need to use (ε, D) copies of an equilateral triangle. We remark that adding the Steiner points corresponding to the Fermat points of the copies will reduce the tree length. The details can be left to the reader.

We turn our attention now to 2-factors. We begin with two very simple geometric lemmas:

Lemma 3.6. *Suppose that points p, q, r, s satisfy*

$$||p - q||, ||r - s|| \geq \Delta \text{ and } ||r - s|| \leq \delta,$$

where $\Delta \gg \delta$ i.e Δ is sufficiently large with respect to δ.

Let $\theta(x; y, z)$ denote the angle between the line segments xy and xz. If

$$\max \{\theta(p; q, s), \theta(s; p, r)\} \geq \Delta^{-1/3}$$

then

$$||p-s|| \le ||p-q|| + ||r-s|| + \delta - \frac{\Delta^{1/3}}{4}.$$

Proof. We have

$$||p-s|| \le ||p-q|| \cos\theta(p;q,s) + \delta + ||r-s|| \cos\theta(s;p,r).$$

Now use $\cos x \le 1 - x^2/3$ for $x \le 1$. □

Lemma 3.7. *Suppose that points $p_i, q_i, r_i, s_i, i = 1, 2$ satisfy*

$$||p_i - q_i||, ||r_i - s_i|| \ge \Delta \quad \textit{for } i = 1, 2 \tag{8}$$

and also that q_1, r_1, q_2, r_2 are contained in a ball of radius δ. Then there is a matching on $\{p_1, p_2, s_1, s_2\}$ whose total length is at most

$$||p_1 - q_1|| + ||r_1 - s_1|| + ||p_2 - q_2|| + ||r_2 - s_2|| + 4\delta - \tfrac{1}{2}\Delta. \tag{9}$$

Proof. Without loss of generality we let the q_i, r_i be within distance δ of the origin, and then let $\theta(x, y)$ denote the angle between x and y via the origin that is less than or equal to π. There are three possible pairings of the points $P = \{p_1, p_2, s_1, s_2\}$, and for at least one such pairing, $\theta(x, y) < \frac{1}{2}\pi$ for one of the pairs.

Let us take $\{x, y\}$ and $\{w, z\}$ to be the pairs in such a pairing of P, with $\theta(x, y) \le \frac{1}{2}\pi$. We let T denote the triangle with vertices $x, y, 0$, let a, b, c denote the side-lengths, where a is length of the side opposite 0, and s denote the semi-perimeter $(a + b + c)/2$. Now $a \le (b^2 + c^2)^{1/2}$ and in fact

$$\begin{aligned} b + c - a \ge b + c - (b^2 + c^2)^{1/2} &= (b + c)\left(1 - \left(1 - \frac{2bc}{(b + c)^2}\right)\right) \\ &\ge \frac{bc}{b + c} \ge \frac{1}{2}\min\{b, c\} \ge \frac{1}{2}\Delta. \end{aligned}$$

Thus we find a pairing of P for which the total length is at most $||p_1|| + ||p_2|| + ||s_1|| + ||s_2|| - \frac{1}{2}\Delta$, and we will be done after applying the triangle inequality four times and using the fact that $||q_i||, ||r_i|| \le \delta$ for $i = 1, 2$. □

Proof of Theorem 1.3. Let F_{g+1} be a minimum length 2-factor in $\mathcal{Y}_n$ whose cycles all have length $\ge g + 1$. We let $U_{\varepsilon,D} \subset \mathcal{Y}_n$ denote any set of g points of radius ε and at distance D from $\mathcal{Y}_n \setminus U_{\varepsilon,D}$. Note that Lemma 3.1 implies that there are a linear number of copies of such sets. We now define $V_{\varepsilon,D,F}$ as a collection of three instances U_1, U_2, U_3 of $U_{\varepsilon,D}$, centered at the vertices of an equilateral triangle of sidelength $2D$, and lying at distance Δ from $\mathcal{Y}_n \setminus V_{\varepsilon,D,\Delta}$; we will take D large relative to ε and Δ large relative to D.

We will begin by showing how to give a constant-factor shortening of F_{g+1} to a 2-factor F, without being careful to avoid creating cycles of length shorter than g. In particular, we prove the following lemma:

Lemma 3.8. *There is an absolute constant δ such that for suitable choices of $\varepsilon \ll D \ll \Delta$, any instance of $V = V_{\varepsilon,D,\Delta}$ allows a modification F of F_{g+1} so that*

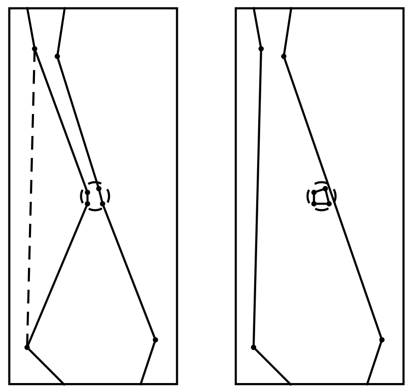

FIGURE 2. When not all pairs are nearly straight the old 2-factor (left) can be shortened to a new one (right). (The dashed circle of radius ε encloses $g+1=4$ points.)

1. *F is a 2-factor;*
2. *F has weight at least δ less than the length of F_{g+1};*
3. *Cycles of F lying entirely in V have length $\geq g$;*
4. *F is a local modification of F_{g+1}, in the sense that any edges of F_{g+1} disjoint from V are still present in F.*

Again, Lemma 3.1 implies that there are a linear number of instances of $V_{\varepsilon,D,\Delta}$ in $\mathcal{Y}_n$. In particular, this lemma would be sufficient to argue that $\beta_{\mathrm{TF}_g} < \beta_{\mathrm{TF}_{g+1}}$, except that F may not have girth g.

Proof of Lemma 3.8. For $U_i = U_{\varepsilon,D}$ in V, there are (at least 2) edges in F_{g+1} from $\mathcal{Y}_n \setminus U_i$ to U_i, since $g+1 > g = |U_i|$. We can pair these edges so that each pair lies on a common cycle of F_{g+1}, and so that the two edges in a pair are joined in F_{g+1} by a path through (possibly just 1 point of) U_i. Similarly, we can pair edges between V and $\mathcal{Y}_n \setminus V$. (Some pairs for V may also be pairs for a U_i, others may not.)

Now, by choosing D large relative to ε, we can assume that each pair of edges for a U_i is *nearly straight*, in the sense that the angle between the endpoints of the edges in $\mathcal{Y}_n$ via any point in $U_{\varepsilon,D}$ is close to π; otherwise, we can modify F_{g+1} by including all edges of some g-cycle through U_i, and shortcutting each pair of edges between $\mathcal{Y}_n \setminus U_i$ and U_i with a single edge between the endpoints in $\mathcal{Y}_n \setminus U_i$. (Figure 2.) The result has length smaller by a constant $\delta = \Omega(D^{1/3})$, see Lemma 3.6. To ensure condition (3) for F, we must now also shortcut all remaining pairs of edges between V and $\mathcal{Y}_n \setminus V$, delete any edges in $V \setminus U_i$, and then add g-cycles to the remaining U_j's. (This step adds length which can be made arbitrarily small by decreasing ε.)

We may also assume that each U_i has only a single pair of edges. Otherwise, if there are two different pairs, we delete the edges in the two pairs, use Lemma 3.7

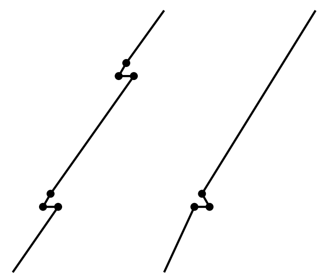

FIGURE 3. An instance of $V_{\varepsilon,D,\Delta}$ (here for $g = 2$, $d = 2$). When all pairs of edges entering/leaving U_i's are nearly straight, we must have at least 2 pairs of edges entering/leaving V, as shown here.

to find a pair of edges among the 4 outside endpoints of the pairs of total weight which is less than the total weight of the pairs by a constant, shortcut all other remaining pairs between V and $\mathcal{Y}_n$, delete all edges within V, and add g-cycles to each U_i. For sufficiently small ε, we get a constant length improvement.

Thus we may assume that each U_i in V has a single pair, and that the pair for each U_i in V is nearly straight. The crucial point is that this implies that there must be at least *two* pairs of edges joining V to $\mathcal{Y}_n \setminus V$: since, e.g., edges joining U_1 to U_2 and U_1 to U_3 would not be nearly straight. Therefore at least one of the U_i's has no edges to the other U_i's. (See Figure 3.) We conclude, as in the previous paragraph, by deleting the edges in the two pairs, using Lemma 3.7 to find a pair of edges among the 4 outside endpoints of the pairs of total weight which is less than the total weight of the pairs by a constant, shortcutting all other remaining pairs between V and $\mathcal{Y}_n$, deleting all edges within V, and adding g-cycles to each U_i. □

We must now address unintentional problems of girth (notice that, in shortcutting edges, we may have left behind short cycles). To this end, we say that $V = V_{\varepsilon,D,\Delta}$ is *ε-surrounded* if the set $\mathcal{N}_V$ of points of $\mathcal{Y}_n \setminus V$ within distance 3Δ of V has the properties that: (1) each $x \in \mathcal{N}_V$ lies within distance ε of the sphere S of radius 2Δ centered at the center of V, and (2) each $x \in S$ lies within ε of $\mathcal{N}_V$. (Essentially, $\mathcal{N}_V$ is an approximation to an ε-net on S, which surrounds V). Lemma 3.1 implies that there are a linear number of ε-surrounded V's, and additionally, a linear number of ε-surrounded sets V satisfying the requirements in the previous paragraph (each U_i has a single-pair of edges to the rest of $\mathcal{Y}_n$, etc.).

We now show that if V is ε-surrounded, then there is an constant $C_{g,\varepsilon}$, which can be made arbitrarily small by decreasing ε, such that there is a 2-factor F' such that:

(A) F' has total weight $w(F') \leq w(F_{g+1}) + C_{g,\varepsilon}$,
(B) every cycle in F' is still of length $\geq g + 1$,
(C) All edges in F' incident with V either lie in V or intersect $\mathcal{N}_V$.

To produce F' from F_{g+1}, we consider each edge $e = \{u, v\}$ from V to $\mathcal{Y}_n \setminus (\mathcal{N}_V \cup V)$ which does not intersect $\mathcal{N}_V$, and

(1) Locate a point x in $\mathcal{N}_V$ within distance ε of a point w on the edge e. Let $C = (x = x_1, x_2, \ldots x_k, x_{k+1} = x_1)$ be the cycle of F_{g+1} that contains x. If $u = x_i$ for some i, then we choose the cycle orientation so that $v = x_{i-1}$.
(2) Add the edges $\{u, x_1\}$, $\{x_k, v\}$ to the 2-factor and delete the edges e and $\{x_1, x_k\}$.

This ensures (C) and the change in cost for this one substitution is

$$\begin{aligned} &||x_1 - u|| + ||x_k - v|| - ||x_1 - x_k|| - ||v - w|| - ||u - w|| \\ \leq &||x_1 - u|| + ||x_1 - w|| - ||u - w|| \\ \leq &2||x_1 - w||. \end{aligned}$$

Thus dealing with all edges from $V_{\varepsilon,D,\Delta}$ to $\mathcal{Y}_n \setminus V_{\varepsilon,D,\Delta}$ increases the cost by at most $12g\varepsilon$, since there are $3g$ points in V and hence at most $6g$ edges from $V_{\varepsilon,D,\Delta}$ to $\mathcal{Y}_n \setminus V_{\varepsilon,D,\Delta}$.

After this, any cycle in F' but not in F_{g+1} must contain an edge added in Step (2). But either $u, v \notin \{x_1, \ldots, x_k\}$, in which case the length of this cycle is at least $k+2 \geq g+3$, or else $u = x_i, v = x_{i+1}$ and this cycle is $x_1, x_2, \ldots, x_{i-1}x_k x_{k-1} \cdots x_i x_1$ and so has length $k \geq g + 1$.

We are now prepared to find a 2-factor F_g whose weight is smaller than F_{g+1} by a constant factor. For some small constant c, we have that there are at least cn instances of ε-surrounded $V = V_{\varepsilon,D,\Delta}$'s. We take these instances as $V_1, V_2, \ldots$, in any order, and beginning with $F = F_{g+1}$ and for each $i = 1, 2, \ldots$, we

(i) Find F' for V_i as above (with weight increase $C_{g,\varepsilon}$ which we make arbitrarily small)
(ii) Apply Lemma 3.8 to shorten F' at V_i to F_0 with a constant weight improvement
(iii) At an arbitrarily small cost, modify F_0 to a 2-factor F_0' which has girth g, by merging cycles intersecting the net $\mathcal{N}_{V_i}$, and set $F = F_0'$ (explanation is below).

In particular, to carry out Step (iii), note that any cycle C of length $< g$ in F_0 includes a point x of $\mathcal{N}_V$, and we can merge C with the cycle through a point y within 2ε of x, at an additional cost of $\leq 2\varepsilon$: We join x and y, delete edges $\{x, x'\}$ and $\{y, y'\}$ incident with each in the previous 2-factor and replace them by $\{x, y\}, \{x', y'\}$ at a cost of

$$||x - y|| + ||x' - y'|| - ||x - x'|| - ||y - y'|| \leq 2||x - y||.$$

After applying Steps (i)–(iii) for each $V \in \mathcal{V}$, the result is a 2-factor $F_g = F$ of girth g, whose total weight is smaller than the total weight of F_{g+1} by a constant factor. □

The proof of the counterpoint Theorem 1.4 will be given in Section 4. For now we consider matchings:

Proof of Theorem 1.5. We define the Euclidean functional $2\mathrm{MM}(X)$ as the minimum length union of two matchings on X. Note that we make no requirement of disjointness and that we trivially have that $2\mathrm{MM}(X) = 2 \cdot \mathrm{MM}(X)$ for all X. On the other hand, a TSP through X can be viewed as a (near)-union of two matchings (alternating edges around the tour, leaving one vertex unmatched if n is odd). Our aim will be to give a constant factor improvement to the union of a pair of matchings given by the TSP, to show that $2\mathrm{MM}(\mathcal{Y}_n)$ is asymptotically less than $\mathrm{TSP}(\mathcal{Y}_n)$. To this end, we let M_1 and M_2 denote a pair of matchings derived from the minimum length TSP.

We let $U_{\varepsilon,D}$ denote a set of two points separated by distance at most ε and at distance at least D from all other points of $\mathcal{Y}_n$, and let $V_{\varepsilon,D,F}$ denote a collection of 5 instances $U_1, \ldots, U_5$ of $U_{\varepsilon,D}$, centered at the vertices of a regular pentagon of sidelength $2D$, such that all other points of $\mathcal{Y}_n$ are at distance $\geq F$ from this set. As before, Lemma 3.1 gives that there are a linear number of instances of $V_{\varepsilon,D,F}$ for any fixed F, D, and $\varepsilon > 0$. Moreover, as before, if we have a linear number of instances $U_{\varepsilon,D}$ in which a pair of edges of a matching leaves $U_{\varepsilon,D}$ and is not nearly straight, then we can make a constant improvement to the matching, by joining the two points of $U_{\varepsilon,D}$ and shortcutting the outside endpoints of the edges leaving $U_{\varepsilon,D}$ with a single edge.

Since M_1 and M_2 are disjoint, the pigeonhole principle gives that for some $s \in \{1, 2\}$ and at least three of the U_i's in any $V_{\varepsilon,D,F}$, the pair of points in U_i is omitted from M_s. In particular, we may assume without loss of generality that we have a linear number of $V_{\varepsilon,D,F}$'s for which the set I of indices i for which the points in U_i are unmatched in M_1 has cardinality $|I| \geq 3$. Moreover, from the previous paragraph, there must be a linear number of such $V_{\varepsilon,D,F}$'s which also have the property that the pair edges leaving the U_i, $i \in I$ is nearly straight. In particular, as the point sets U_i $(i \in I)$ are not nearly collinear, we must have as in the previous proof that there are (at least) 2 pairs of edges entering and leaving $V_{\varepsilon,D,F}$. We conclude by applying Lemma 3.7 (with 2ε, say) to get a constant factor improvement a linear number of times. □

We close this section by considering H-factors.

3.1. **Proof of Theorem 1.7.** It suffices to show that for fixed $r \geq 1$, there are connected graphs H with $r \cdot |V(H)|$ edges for which the constant β_H^d is arbitrarily large, which we show by demonstrating that β_T^d can be arbitrarily large even just over trees T. To this end, we let T_k be the tree on $k+1$ vertices which has k leaves.

Given any large constant $u = t/m$ for some integer m, we decompose the $[0, t]^d$ cube with m^d subcubes of side u. Now the number of points in each subcube is binomially distributed with mean u^d. Let a point in $\mathcal{Y}_n$ be *good* if the subcube S_α that it lies in has at least $(1-\varepsilon)u^d$ members of $\mathcal{Y}_n$ and the total number of points in the $\leq 3^d$ subcubes that touch S_α contain at most $(1+\varepsilon)(3u)^d$ members of $\mathcal{Y}_n$, where $\varepsilon = \frac{1}{10k}$. Assuming that u is sufficiently large, the Chernoff bounds imply that a member of $\mathcal{Y}_n$ is good with probability at least $1 - \varepsilon/2$. Thus the expected number of good points in $\mathcal{Y}_n$ is at least $(1 - \varepsilon/2)n$. Now the Chernoff bounds can be used to show that the number of members of $\mathcal{Y}_n$ in any subcube is a.s. $O(\log n)$ and therefore, changing one point only changes the number of good points by $O(\log n)$ a.s. A fairly simple modification of McDiarmid's inequality now implies that a.s. $(1 - \varepsilon)n$ of the members of $\mathcal{Y}_n$ are good.

Since $\approx n/(k + 1)$ points must have degree k in a T_k factor of $\mathcal{Y}_n$, we have that there are at least $n/(2k)$ good points which have degree k. Now let $k = 2(3u)^d$. Then a.s. a T_k factor has length at least $\frac{n}{2k} \cdot \frac{(1-\varepsilon)k}{2} \cdot u > \frac{un}{5}$.

Rescaling the $[0, t]^d$ cube by a factor of t gives that the minimum T_k factor has length at least $\frac{1}{5} un^{\frac{d-1}{d}}$, and here $\frac{u}{5}$ is an arbitrarily large constant. □

4. The 2-factor limit

Here we prove Theorem 1.4: $\lim_{g\to\infty} \beta^d_{\mathrm{TF}_g} = \beta^d_{\mathrm{TSP}}$. We will continue to work with $\mathcal{Y}_n$ as above. We divide $[0, t]^d$ into $m^d = n/L^d$ subcubes $S_\alpha, \alpha \in [m]^d$ of sidelength L, for some sufficiently large constant $L > 0$.

With each cube S_α we associate the 2^d quadrants $Q_{\alpha,j}, j = 1, 2, \dots, 2^d$, whose origin is the center s_α of S_α. We call the quadrant $Q_{a,j}$ *trivial* if the quadrant intersects $[0, t]^d$ in a unit cube (in which case S_α is one of the 2^d corner cubes in the decomposition). Then for a non-negative integer r, we let $Q_{\alpha,j,r}$ denote the cubes in $Q_{\alpha,j}$ whose centers are at distance at most rL from s_α; for convenience, we call $Q_{\alpha,j,r}$ trivial (resp. nontrivial) whenever $Q_{\alpha,j}$ is, regardless of the choice of r.

If $Q_{\alpha,j,r} \subseteq [0, t]^d$ is nontrivial and $Y_{\alpha,j,r}$ is the number of points of $\mathcal{Y}_n$ that are in $Q_{\alpha,j,r}$ then $Y_{\alpha,j,r}$ is a binomial random variable with mean

$$\alpha_d rL \le \mathbf{E}\, Y_{\alpha,j,r} \le \beta_d (rL)^d$$

for some constants $\alpha_d, \beta_d > 0$. Note that, away from the boundary cubes of the decomposition of $[0, t]^d$ we can use $(rL)^d$ in place of rL for the lower bound, but in the worst-case, we have to reduce the exponent. We can therefore write

$$\mathbf{Pr}(Y_{\alpha,j,r} = 0) \le e^{-\gamma_d rL} \tag{10}$$

for some $\gamma_d > 0$.

Next let ν_r denote the number of subcubes S_α for which there exists j, r such that $Q_{a,j}$ is nontrivial, $Q_{\alpha,j,r} \subseteq [0, t]^d$, and $Y_{\alpha,j,r} = 0$. Then

$$\mathbf{E}\, \nu_r \le n 2^d e^{-\gamma_d rL}, \tag{11}$$

for some $\gamma_d > 0$. We deduce from the above that

$$\nu_r = 0 \text{ for } r \ge r_0 = \frac{2}{L}(\gamma_d^{-1}(\log n + d \log 2). \tag{12}$$

Now ν_r is determined by n independent choices for the points in $\mathcal{Y}_n$. Changing one point changes ν_r by at most $\delta_d r^d$ for some $\delta_d > 0$. Applying McDiarmid's inequality we see that if $t > 0$ and $r < r_0$ then

$$\mathbf{Pr}(\nu_r \geq \mathbf{E}\,\nu_r + t) \leq \exp\left\{-\frac{t^2}{n\delta_d^2 r^{2d}}\right\}. \tag{13}$$

If $\mathbf{E}\,\nu_r \geq n^{2/3}$ then (11) implies that $r \leq \frac{\log n + 3d\log 2}{3\gamma_d L} \leq \frac{1}{10d}\log n$ for L sufficiently large. It follows from (13) with $t = \mathbf{E}\,\nu_r$ that

$$\nu_r \leq 2\,\mathbf{E}\,\nu_r \text{ a.s. if } \mathbf{E}\,\nu_r \geq n^{2/3}. \tag{14}$$

When $\mathbf{E}\,Z \leq n^{2/3}$ we use (7) with $\alpha = n^{3/4}/(\mathbf{E}\,\nu_r)$ and $c = \delta_d r^d = \log^{O(1)} n$ to obtain

$$\nu_r \leq n^{3/4} \text{ a.s. if } \mathbf{E}\,\nu_r \leq n^{2/3}. \tag{15}$$

Now suppose that $C_1, C_2, \ldots, C_M$ are the cycles of a minimum cost 2-factor, where $|C_i| \geq g$ for $i = 1, 2, \ldots, M$. Suppose first there exist i, j such that there exist $S_p \ni x \in C_i$ and $S_q \ni y \in C_j$ such that $||s_p - s_q|| \leq L^2$. Suppose that (x, x') is an edge of C_i and that (y, y') is an edge of C_j. Then $||x - y|| \leq (L + 2d^{1/2})L$ and $||x' - y'|| \leq ||x' - x|| + ||x - y|| + ||y' - y||$. It follows that if we delete the edges $(x, x'), (y, y')$ from C_i, C_j and add the edges $(x, y), (x', y')$ then we create a single cycle out of the vertices of $C_i \cup C_j$ at a cost of at most $2||x' - y||$. By repeating this where possible, we obtain a new set of cycles $C_1', C_2', \ldots, C_{M'}'$ such that for two distinct cycles C_i', C_j' the set of subcubes visited by C_i' have centers that are distance at least L^2 from the centers of the set of subcubes visited by C_j'. Furthermore, the increase in cost associated with this merging is at most

$$\frac{2(L + 2d^{1/2})L}{g} n. \tag{16}$$

We continue merging cycles. For $r = L + 1 \ldots, r_0$ we try to merge cycles C, C' for which there is a subcube S_i containing a point of C whose center is within rL of the center of a subcube that contains a point of C'. The cost of making these merges can be bounded by

$$2\sum_{r=L+1}^{r_0} n2^d r^d e^{-\gamma_d r L} + n^{3/4} r_0 r^d \leq 3n(2L)^d e^{-\gamma_d L^2} \tag{17}$$

for L sufficiently large.

This is because, when we merge two cycles via subcubes at distance rL we are using one of at most ν_r subcubes. Further, for each such subcube there are at most r^d other subcubes at distance r.

We argue next that after all of these merges, there can be only one cycle. Suppose that there are two cycles C, C' and let $x \in C, x' \in C'$ be as close as possible. Suppose that $x \in S_a$ and that $x' \in S_b$ where $||s_a - s_b|| > r_0$. If this happens then we can find a Q_{a,j,r_0} or a Q_{b,j,r_0} that is empty, contradiction.

It follows from this and (16), (17) that with $L = g^{2/3}$, that after scaling to $[0,1]^d$ we find that for g sufficiently large,

$$\beta^d_{\text{MST}} \le \beta^d_{\text{TF}_g} + g^{-1/4} + e^{-\gamma_d g^{2/3}/2}$$

and this completes the proof of Theorem 1.3. □

5. Branch and Bound Algorithms

In this section we prove Theorem 1.8.

We begin by defining a *branch and tree* $\mathcal{T}$. This is a rooted tree in which each vertex v has a label $L_v = (I_v, O_v, b_v)$. Here I_v, O_v are disjoint subsets of the edges $\binom{[n]}{2}$ of the complete graph K_n and $b_v \in \mathbb{R}$. Vertex v of the tree represents the problem of finding the shortest tour H through $\mathcal{X}_n$ given that H must use every edge in I_v and none of the edges in O_v. Here the weight of an edge $\{i,j\}$ is of course the Euclidean length $||x_i - x_j||$ between the corresponding points. Let Ω_v denote the set of tours that satisfy these constraints. b_v will be a lower estimate of the length of H. Here b_v will be the minimum length of a 2-factor F that satisfies the edge constraints, and has girth at least g.

We will assume that the edges in I_v induce a collection of vertex disjoint paths. Furthermore, if x is an interior point of one such path with path neighbors x_1, x_2, then no tour in Ω_v can use an edge $\{x,y\}$ where $y \ne x_1, x_2$, and such edges are not included in O_v. Moreover, for the root ρ we take $I_\rho = O_\rho = \emptyset$. Next, if vertex $p(v)$ denotes the parent of vertex v we must have $I_{p(v)} \subseteq I_v$ and $O_{p(v)} \subseteq O_v$. This means that we can assume that $b_{p(v)} \le b_v$. In addition we require that the sets Ω_v, $p(v) = w$ partition Ω_w. If v is a leaf of $\mathcal{T}$ then $\Omega_v = \{H_v\}$ i.e. it consists of a unique Hamilton cycle of K_n. Each of the Hamilton cycles of K_n appears exactly once as H_v.

At each stage of the algorithm, there will be a value B available. This will be an upper bound on the length of the minimum tour. In practice, B should decrease as the search progresses, if, in addition to branching, we simultaneously apply some heuristic to find (suboptimal) tours. In the ensuing analysis, however, we will simply assume that B is always the actual minimum tour length.

The branch-and-bound algorithm will search the tree $\mathcal{T}$ of instances for an optimum tour. It does not search the whole tree. It prunes the tree by deleting subtrees *strictly* below any vertex v for which $b_v \ge B$. This leaves the *pruned* tree $\hat{\mathcal{T}}$ and then the number of vertices of $\hat{\mathcal{T}}$ is a lower bound on the size of the branch and bound tree that is produced if we use 2-factors (with large girth) for lower bounds. One sees immediately, that a smaller B always results in fewer nodes being explored.

We now introduce our restriction on the class of branch and bound trees. We assume that there are absolute constants $C, \delta > 0$ such that if $|I_v| \le \delta n^{(d-1)/d}$, then $|O_v| \le C|I_v|$. We say that the algorithm is (C,δ)-restricted. This restriction includes many natural branching strategies, including the following one: Having solved the 2-factor problem for an instance v, we choose a cycle $C = (x_1, x_2, \dots, x_k, x_1)$ of the factor. We then create $(k-1)$ children of the current vertex w. For the ith child v_i we let

$$I_{v_i} = I_w \cup \{\{x_j, x_{j+1}\}, 1 \le j \le i\} \text{ and } O_{v_i} = O_w \cup \{\{x_{i+1}, x_{i+2}\}\}.$$

Our goal is to show that $\hat{\mathcal{T}}$ is large by showing that a.s. $b_v < B$ whenever $|I_v| + |O_v| \le \varepsilon n^{(d-1)/d}$ for some sufficiently small $\varepsilon > 0$. We do not strive for the best possible bounds; the following will suffice.

Lemma 5.1. *There is an absolute constant $c > 0$ such that,*

$$\mathbf{Pr}(TF(\mathcal{X}_n) \ge \mathbf{E}\, TF(\mathcal{X}_n) + t) \le \begin{cases} \exp\left\{-\frac{ct^2}{\log n}\right\} & d = 2, \\ \exp\left\{-\frac{ct^2}{n^{(d-2)/d}}\right\} & d \ge 3. \end{cases}$$

In particular, putting $t = \alpha n^{(d-1)/d}$, we see that

$$\mathbf{Pr}(TF(\mathcal{X}_n) \ge \mathbf{E}\, TF(\mathcal{X}_n) + \alpha n^{(d-1)/d}) \le \begin{cases} \exp\left\{-\frac{c\alpha^2 n}{\log n}\right\} & d = 2, \\ e^{-c\alpha^2 n} & d \ge 3. \end{cases} \tag{18}$$

Proof. We use a modification of the argument of Steele [15], Section 2.1. We write

$$d_i = \mathbf{E}(TF(x_1, x_2, \ldots, x_i, \ldots, x_n) - TF(x_1, x_2, \ldots, \hat{x}_i, \ldots, x_n \mid x_1, \ldots, x_i)$$

where $\{\hat{x}_i\}$ is a sequence with the same distribution as $\mathcal{X}_n$, but independent of it. We will prove that

$$|d_i| \le \begin{cases} \frac{C_{d,g}}{(n-i+1)^{1/d}} & i \le n - 100g, \\ 4d^{1/2} & i > n - 100g, \end{cases} \tag{19}$$

for some constant $C_{d,g} > 0$.

We can then use the Azuma-Hoeffding inequality

$$\mathbf{Pr}(TF \ge \mathbf{E}\, TF + t) \le \exp\left\{-\frac{t^2}{2\sum_{i=1}^n |d_i|^2}\right\}.$$

This implies the lemma, since (19) implies that $\sum_i |d_i|^t = O(\log n)$ for $d = 2$ and $O(n^{(d-2)/d})$ for $d \ge 3$.

Now fix $x_1, x_2, \ldots, x_n, \hat{x}_i$ and let

$$\Delta = |TF(x_1, x_2, \ldots, x_i, \ldots, x_n) - TF(x_1, x_2, \ldots, \hat{x}_i, \ldots, x_n)|.$$

Let F be the optimal 2-factor for $x_1, x_2, \ldots, x_n$. Suppose that the neighbors of $x = x_i$ on its cycle C in F are y, z. If $|C| = g$ then we cannot simply delete x and replace the path (y, x, z) by (y, z) as this will produce a 2-factor of girth $g - 1$. So, let a be the closest point to x that is not on C and let b be a neighbor of a on the cycle C' of F that contains a. The first thing we do now is to delete x and merge the points in $C \cup C' \setminus \{x\}$ into one cycle. We delete the edges $\{x, y\}, \{x, z\}, \{a, b\}$ and add the edges $\{y, a\}, \{z, b\}$. The change in cost is

$$\begin{aligned} &|y-a| + |z-b| - |x-y| - |x-z| - |a-b| \le \\ &\quad (|x-y| + |x-a|) + (|z-x| + |x-a| + |a-b|) - |x-y| - |x-z| - |a-b| \\ &\qquad = 2|x-a|. \end{aligned}$$

The new cycle has length at least $2g-1 \geq g$. After this we can insert $\hat{x}=\hat{x}_i$ into the cycle D,say, that contains the point c of $x_1,\dots,x_{i-1},x_{i+1},\dots,x_n$ closest to $\hat{x}$. Suppose that d is a neighbor of c on D. Then we remove the edge $\{c,d\}$ and replace it with the edges $\{\hat{x},c\},\{\hat{x},d\}$. This will not decrease the girth of the factor. The change in cost is

$$|\hat{x}-c|+|\hat{x}-d|-|c-d| \leq |\hat{x}-c|+(|\hat{x}-c|+|c-d|)-|c-d| = 2|\hat{x}-c|.$$

Thus,

$$\Delta \leq 2\min_{j\notin N_g(i)} |x_i-x_j| + 2\min_{j\neq i}|\hat{x}-x_j|,$$

where $N_g(i)$ is the set of g points in $x_1,\dots,x_{i-1},x_{i+1},\dots,x_n$ closest to x_i.

Because the definition of d_i involves conditioning on $x_1,x_2,\dots,x_i$ we will replace the above upper bound on Δ by

$$\Delta \leq 2\min_{i<j\notin N_g(i)} |x_i-x_j| + 2\min_{i<j}|\hat{x}_i-x_j|. \tag{20}$$

Now if ρ_d denotes the volume of a ball of radius one in $\mathbb{R}^d$ and

$$2^{-d}\rho_d\lambda^d(n-i) \geq 2g \text{ or } \lambda \geq \lambda_0 = A_{d,g}(n-i)^{-1/d}$$

for some constant $A_{d,g}$, then

$$\begin{aligned}\mathbf{Pr}\left(\min_{i<j\notin N_g(i)}|x_i-x_j|\geq\lambda\right) &\leq \mathbf{Pr}(\mathrm{Bin}(n-i,2^{-d}\rho_d\lambda^d)\leq g)\\ &\leq 2\binom{n-i}{g}(2^{-d}\rho_d\lambda^d)^g\exp\left\{-(n-i-g)2^{-d}\rho_d\lambda^d\right\}.\end{aligned}$$

The 2^{-d} factor accounts for the possibility that x_i is close to a corner of $[0,1]^d$. Also, in this probability estimate, x_i is fixed and x_j is chosen uniformly from $[0,1]^d$. So, for some constant $B_{d,g}$ and $i \leq n-100g$,

$$\begin{aligned}&\mathbf{E}\min_{i<j\notin N_g(i)}|x_i-x_j|\\ &\leq \frac{A_{d,g}}{(n-i)^{1/d}} + B_{d,g}\int_{\lambda_0}^{\infty}((n-i-g)\lambda^d)^g\exp\left\{-(n-i-g)2^{-d}\rho_d\lambda^d\right\}d\lambda\\ &= \frac{A_{d,g}}{(n-i)^{1/d}} + \frac{B_{d,g}}{d(n-i-g)^{1/d}}\int_0^{\infty}\mu^g e^{-2^{-d}\rho_d\mu}d\mu\\ &\leq \frac{K_{d,g}}{(n-i)^{1/d}}.\end{aligned}$$

Going back to (20) we see that this is good enough to prove the case $n-i\geq 100g$ in (19). The case $n-i\leq 100g$ is trivial, because the diameter of $[0,1]^d$ is $d^{1/2}$. □

So, we fix v,I_v,O_v and estimate the probability that $b_v \geq (\beta^{(d)}_{\mathrm{TF}_g}+\varepsilon)n^{(d-1)/d}$ for a small ε. Indeed we can bound b_v by $2(|I_v|+|O_v|)d^{1/2}$ plus the minimum cost of a 2-factor on the vertices that are not involved in edges defined by I_v,O_v. So,

$$\mathbf{Pr}(b_v \geq 2(|I_v|+|O_v|)d^{1/2} + (\beta_{TF}+\varepsilon)n^{(d-1)/d}+t) \leq RHS[(18)]. \tag{21}$$

To cover all possible choices for I_v,O_v, we need only inflate the above probability upper bound by $\binom{n}{2(|I_v|+|O_v|)} = e^{o(n/\log n)}$.

Because $\beta_{TSP} > \beta_{TF}$ we see that a.s. for ε sufficiently small,

$$b_v < B \text{ for all } v \text{ such that } |I_v| + |O_v| \leq \varepsilon n^{(d-1)/d}. \tag{22}$$

Now let L be the set of leaves of $\hat{\mathcal{T}}$. We must have

$$\frac{(n-1)!}{2} = \sum_{v \in L} |\Omega_v| \tag{23}$$

Suppose now that for some vertex $v \in T$ the set of edges I_v induces a_v paths. Then the number of Hamilton cycles in K_n that contain the edges of I_v is

$$2^{a_v - 1}(n - |I_v| - 1)! \leq 2^{|I_v|-1}(n - |I_v| - 1)!.$$

Going back to (23) we see that

$$\frac{(n-1)!}{2} \leq \sum_{v \in L} 2^{|I_v|-1}(n - |I_v| - 1)!$$

or

$$\sum_{v \in L} \frac{2^{|I_v|}}{(n-1)_{|I_v|}} \geq 1, \tag{24}$$

where $(x)_a = x(x-1)\cdots(x-a+1)$.

Now if $v \in L$ then $b_v \geq B$ and so we can assume from (18) and (21) that $|I_v|+|O_v| > \varepsilon n^{(d-1)/d}$ for all $v \in T$. Furthermore, if $|I_v| \leq \delta n^{(d-1)/d}$ then the (C, δ) restriction implies that $\delta(C+1)|I_v| > \varepsilon n^{(d-1)/d}$. So we have

$$v \in L \text{ implies that } |I_v| > \Lambda = c_1 n^{(d-1)/d}$$

where $c_1 = \min\left\{\delta, \frac{\varepsilon}{\delta(C+1)}\right\}$.

It follows from (24) that

$$|L| \geq \frac{(n-1)_\Lambda}{2^\Lambda}.$$

This completes the proof of Theorem 1.8. □

6. Final Remarks

Our results lead to many natural directions of inquiry, and here we mention just a few. Apart from simply increasing the list of separated pairs of constants, the following seems like a very good challenge:

1. What is the relationship between β^d_{MST}, β^d_{TF}, and $2\beta^d_{\text{MM}}$?

In connection with Theorem 1.4:

2. The minimum length of covering of $\mathcal{X}_n$ by paths of lengths $\geq k$ is a Euclidean functional; let $\beta_{P,k}$ denote the constant in its asymptotic formula. Is it true that $\lim\limits_{k\to\infty} \beta_{P,k} = \beta_{\text{TSP}}$?

Short of a full confirmation of Conjecture 1.6, one could warm up with some special cases:

3. Pick an integer k, and then prove or disprove that distinct unlabeled trees T on k vertices have distinct asymptotic constants β_T^d.

Though the condition in Theorem 1.8 is not so restrictive, it would be nice to remove it:

4. Does Theorem 1.8 remain true if we remove the (C,δ)-restriction?

Finally, we note that as our methods for separating constants give only very small differences, we have not attempted to calculate lower bounds on, say, $\beta_{\mathrm{TSP}}^d - \beta_{\mathrm{TF}}^d$, or optimize our techniques for this purpose, though this project could be pursued.

References

[1] D. Applegate, R. Bixby, V. Chvátal and W. Cook, The Traveling Salesman Problem: A Computational Study, Princeton University Press, 2007.

[2] F. Avram and D. Bertsimas, The minimum spanning tree constant in geometrical probability and under the independent model: a unified approach, in *The Annals of Applied Probability* 2 (1992) 113–130.

[3] J. Beardwood, J. H. Halton and J. M. Hammersley, The shortest path through many points, *Mathematical Proceedings of the Cambridge Philosophical Society* 55 (1959) 299-327.

[4] M. Bern, Two probabilistic results on rectilinear Steiner trees, in *Algorithmica* 3 (1988) 191-204.

[5] F.R.K. Chung and R.L. Graham, On Steiner trees for bounded point sets, in *Geometriae Dedicata* 11 (1981) 353–361.

[6] D. J. Bertsimas and G. Van Ryzin, An asymptotic determination of the minimum spanning tree and minimum matching constants in geometrical probability, *Operations Research Letters* **9** (1990) 223–231.

[7] L. Few, The shortest path and the shortest road through n points, *Mathematika* 2 (1955) 141-144.

[8] A.M. Frieze, On random symmetric travelling salesman problems, *Mathematics of Operations Research* 29 (2004) 878-890.

[9] E. N. Gilbert, Random minimal tress, *Journal of the Society for Industrial and Applied Mathematics* 13 (1965) 376–387.

[10] F. K. Hwang and Y. C. Yao, Comments on Bern's probabilistic results on rectilinear Steiner trees, in *Algorithmica* 5 (1990) 591–598.

[11] C. McDiarmid, On the Method of Bounded Differences, *Surveys in Combinatorics* 141 (1989) 148188.

[12] D. Miller and J. Pekny, Exact Solution of Large Asymmetric Traveling Salesman Problems, in *Science* 251 754–761.

[13] C. H. Papadimitriou, The probabilistic analysis of matching heuristics, in *Proc. 15th Conference Comm. Contr. Computing.* Univ. Illinois.

[14] J. Michael Steele, Subadditive Euclidean functionals and nonlinear growth in geometric probability, *The Annals of Probability* 9 (1981) 365-376.

[15] J. Michael Steele, Probability Theory and Combinatorial Optimization, SIAM CBMS series, 1996.

[16] J. M. Steele, L. A. Shepp and W. F. Eddy, On the Number of Leaves of a Euclidean Minimal Spanning Tree, *Journal of Applied Probability* 24 (1987) 809–826.

[17] L. Tóth, Über einen geometrischen Satz, *Mathematische Zeitschrift* 46 (in German) (1940) 8385.

Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, U.S.A.

E-mail address, Alan Frieze: alan@random.math.cmu.edu

E-mail address, Wesley Pegden: wes@math.cmu.edu