

SEPARATING SUBADDITIVE EUCLIDEAN FUNCTIONALS

ALAN FRIEZE AND WESLEY PEGDEN

ABSTRACT. The classical Beardwood-Halton-Hammersly theorem (1959) asserts the existence of an asymptotic formula of the form $\beta\sqrt{n}$ for the minimum length of a Traveling Salesperson Tour through n random points in the unit square, and in the decades since it was proved, the existence of such formulas has been shown for other such *Euclidean functionals* on random points in the unit square as well. Despite more than 50 years of attention, however, it remained unknown whether the minimum length TSP through n random points in $[0, 1]^2$ was asymptotically distinct from its natural lower bounds, such as the minimum length spanning tree, the minimum length 2-factor, or, as raised by Goemans and Bertsimas, from its linear programming relaxation.

We prove that the TSP on random points in Euclidean space is indeed asymptotically distinct from these and other natural lower bounds, and show that this separation implies that branch-and-bound algorithms based on these natural lower bounds must take nearly exponential ($e^{\tilde{\Omega}(n)}$) time to solve the TSP to optimality, even in average case. This is the first average-case superpolynomial lower bound for these branch-and-bound algorithms (a lower bound as strong as $e^{\tilde{\Omega}(n)}$ was not even been known in worst-case analysis).

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, \dots is a random sequence of points in $[0, 1]^d$ and $\mathcal{X}_n = \{x_1, \dots, x_n\}$, their results imply that there is an absolute constant β_{TSP}^d such that the length $\text{TSP}(\mathcal{X}_n)$ of a minimum length tour through \mathcal{X}_n satisfies

$$(1) \quad \text{TSP}(\mathcal{X}_n) \sim \beta_{\text{TSP}}^d n^{\frac{d-1}{d}} \quad a.s.$$

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 $\text{MST}(\mathcal{X}_n)$ [3], and the minimum length of a matching $\text{MM}(\mathcal{X}_n)$ [24]. Steele [28] provided a general framework which enables fast assertion of identical asymptotic formulas for these and other suitable *Euclidean functionals*; we recall his definitions and theorem in Appendix A.

A major remaining problem in this area is to obtain analytic results regarding the constants β in such formulas. In particular, the best rigorous bounds on such

Date: April 29, 2024.

Research supported in part by NSF grant DMS-1362785.

Research supported in part by NSF grant DMS-1363136.

	lower	upper
β_{TSP}^2	.62866 [32]	.92116... [3, 32]
β_{MST}^2	.60082 [2]	$\frac{1}{\sqrt{2}} \approx .707$ [12]
$2\beta_{\text{MM}}^2$.5 [5]	.92116

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

constants are generally very weak, with known results for $d = 2$ given in Table 1. In particular, the bounds on β_{TSP}^2 were not improved since 1959, until the paper [32] of Steinerberger improved the lower bound by $\frac{19}{5184} \approx .0036\dots$ and the upper bound by $\approx 10^{-6}$. On the other hand, there was some success as d grows large, as Bertsimas and Van Ryzen [5] showed that, asymptotically in d ,

$$(2) \quad \beta_{\text{MST}}^d \sim 2\beta_{\text{MM}}^d \sim \sqrt{\frac{d}{2\pi e}},$$

and conjectured that $\beta_{\text{TSP}}^d \sim \sqrt{\frac{d}{2\pi e}}$ as well. This conjecture was subsequently verified by Rhee [25].

In the present paper, we will asymptotically separate the TSP from its natural lower bounds; in particular, we will prove that $\beta_{\text{MST}}^d < \beta_{\text{TSP}}^d$, $\beta_{\text{TF}}^d < \beta_{\text{TSP}}^d$, and $2\beta_{\text{MM}}^d < \beta_{\text{TSP}}^d$ for all d . Here β_{TF}^d denotes the minimum length of a 2-factor, see below for more details. These are the first asymptotic separations for Euclidean functionals where the Euclidean metric is playing an essential role: the only previous separation was shown (by Bern [4]; see also [18]) 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 [7].) Finally, we will also asymptotically separate the TSP from its linear programming relaxation.

We begin by considering the degrees of vertices in the minimum spanning trees among n random points. Steele, Shepp, and Eddy [31] 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 $\tau'(d) \leq \tau(d)$, and in particular, $\tau'(2) = 5$.) We prove:

Theorem 1.1. $\alpha_{k,d} > 0$ if and only if $k \leq \tau'(d)$.

We now consider the Euclidean functionals $\text{MST}_k(\mathcal{X})$ defined as the minimum length of a spanning tree of \mathcal{X} whose vertices all have degree $\leq k$. It was shown

in [27] that $\text{MST}_k(\mathcal{X}) \sim \beta_{\text{MST}_k}^d n^{\frac{d-1}{d}}$ for each k , and we prove separation of these asymptotic formulas as follows:

Theorem 1.2. *We have that*

$$(3) \quad \beta_{\text{TSP}}^d = \beta_{\text{MST}_2}^d > \beta_{\text{MST}_3}^d > \cdots > \beta_{\text{MST}_{\tau'(d)}}^d = \beta_{\text{MST}}^d$$

for all d .

Thus, the MST_k constants are as diverse as are allowed by the simple geometric constraint of $\tau'(d)$. Note that Theorem 1.2 has consequences for classical TSP approximation schemes. Christophides' approximation for the Metric TSP works by finding a minimum spanning tree and a minimum matching on the odd-degree vertices of the spanning tree, allowing an Euler circuit among the edge set, which can be shortcutted to a TSP assuming the triangle inequality. Christophedes' algorithm always produces a tour within a factor of $\frac{3}{2}$ of the optimum length (the weight of the tree plus the weight of the matching), which remains the best worst-case approximation ratio for an efficient algorithm for the metric TSP. Theorem 1.2 implies that the approximation ratio is strictly better than $\frac{3}{2}$ for random Euclidean point sets, w.h.p, since the tree used in the construction will be asymptotically $\beta_{\text{MST}}^d n^{\frac{d-1}{d}} < \text{TSP}(\mathcal{X}_n)$ w.h.p. (Similarly, Theorem 1.2 implies that the double-tree heuristic has approximation ratio < 2 for random Euclidean configurations, w.h.p.)

Another natural lower bound for the TSP is given by the minimum length 2-factor; let us recall that a *2-factor* is a disjoint set of cycles covering a given set of points. We verify in Appendix A that the length of the minimum 2-factor has an asymptotic formula $\text{TF}(\mathcal{X}_n) \sim \beta_{\text{TF}}^d n^{\frac{d-1}{d}}$ for some constant β_{TF}^d . Moreover, if $\text{TF}_g(X)$ is the minimum length of a 2-factor through X whose cycles all have length $\geq g$, then $\text{TF}_g(\mathcal{X}_n) \sim \beta_{\text{TF}_g}^d n^{\frac{d-1}{d}}$ for constants $\beta_{\text{TF}_g}^d$. Naturally, we must have $\beta_{\text{TF}}^d = \beta_{\text{TF}_3}^d \leq \beta_{\text{TF}_4}^d \leq \beta_{\text{TF}_5}^d \leq \cdots$. It is not clear *a priori* whether small cycles are asymptotically essential to optimum 2-factors in random point sets. The following theorem shows that they are:

Theorem 1.3. *$\beta_{\text{TF}_g}^d$ is a monotone increasing sequence $\beta_{\text{TF}_3}^d < \beta_{\text{TF}_4}^d < \beta_{\text{TF}_5}^d < \cdots$. In particular, $\beta_{\text{TF}_g}^d < \beta_{\text{TSP}}^d$ for all g .*

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

Theorem 1.4. $\lim_{g \rightarrow \infty} \beta_{\text{TF}_g}^d = \beta_{\text{TSP}}^d$.

Theorem 1.3 stands in contrast to the independent case where the edge lengths $X_e, e \in \binom{[n]}{2}$ are independent uniform $[0, 1]$ random variables. In that setting, Frieze [11] showed that with probability $1 - o(1)$, the weight of the minimal 2-factor is asymptotically equal to the minimum length of a tour.

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

Theorem 1.5. $2\beta_{\text{MM}}^d < \beta_{\text{TSP}}^d$.

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

$$\text{HF}(\mathcal{X}_n) \sim \beta_H^d n^{\frac{d-1}{d}}.$$

We pose the following conjecture:

Conjecture 1.6. *Given H_1, H_2 and $d \geq 2$, we have that $\beta_{H_1}^d \neq \beta_{H_2}^d$ unless H_1 and H_2 are each isomorphic to a disjoint union of copies of some graph H_3 . In particular, $\beta_{H_1}^d \neq \beta_{H_2}^d$ 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 final separation result concerns the linear programming relaxation of the TSP. The TSP through a set of points can be given as the following integer program, on variables $x_{\{ij\}}$ ($i, j \in V, i < j$) indicating the presence of an edge between vertices i and j in the tour, where $c_{\{ij\}}$ gives the cost of the edge $\{i, j\}$:

$$\begin{aligned} & \min \sum_{\{i,j\} \subseteq V} c_{\{ij\}} x_{\{ij\}} \\ & \text{subject to} \\ & \hline & (\forall i) \quad \sum_{j \neq i} x_{\{ij\}} = 2 \\ & (\forall \emptyset \neq S \subsetneq V) \quad \sum_{\{i,j\} \subseteq S} x_{\{ij\}} \leq |S| - 1 \\ & (\forall i < j \in V) \quad x_{\{ij\}} \in \{0, 1\} \end{aligned}$$

The linear programming relaxation of the above integer program replaces the final constraint with the requirements that $0 \leq x_{\{ij\}} \leq 1$ for each $x_{\{ij\}}$. It is often referred to as the Held-Karp relaxation of the TSP, but its origins go back to the paper of Dantzig, Fulkerson and Johnson [9]. It should be noted that, although this LP has an exponential number of constraints, it can be solved in polynomial time, e.g., using the ellipsoid algorithm [14].

We denote by $\text{HK}(\mathcal{X}_n)$ the value of a solution to the Held-Karp relaxation on \mathcal{X}_n . Of course, $\text{HK}(\mathcal{X}_n) \leq \text{TSP}(\mathcal{X}_n)$. The Held-Karp bound on the TSP is generally considered to be a good bound which is algorithmically useful on ‘typical instances’ (see [6, 15, 19, 26, 34]), including as pruning bound for branch-and-bound algorithms,

discussed below. With this motivation, Goemans and Bertsimas [13] showed that $\text{HK}(\mathcal{X}_n) \sim \beta_{\text{HK}} n^{\frac{d-1}{d}}$ for some constant β_{HK} , and asked whether $\beta_{\text{HK}} = \beta_{\text{TSP}}$. Experimental evidence has suggested that any gap between the constants would be less than 1%. [20]

Theorem 1.8. *For all $d \geq 2$, $\beta_{\text{HK}}^d < \beta_{\text{TSP}}^d$.*

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]).

Our results show, however, that several natural approximators are asymptotically distinct from the TSP. This will be algorithmically relevant as follows:

Theorem 1.9. *Suppose that we use branch and bound to solve the TSP on \mathcal{X}_n , using any branching strategy, any tour heuristic, and using TF_g or HK as a lower bound. Then, w.h.p, the algorithm runs in time $e^{\Omega(n/\log^6 n)}$.*

In particular, this gives a rigorous explanation for the observation (see [22], for example) that branch-and-bound heuristics using the Assignment Problem as a bounding estimate (even weaker than the 2-factor) perform poorly on Euclidean instances, and indicates that the success of the Held-Karp bound in branch and bound algorithms will be limited for sufficiently large Euclidean point sets. We emphasize that this is the first average-case hardness result (stronger than worst-case hardness) for the Euclidean TSP of which we are aware.

Remark 1.10. Asymptotic formulas are available for subadditive Euclidean functionals in more general settings. If $x_1, x_2, \dots \in \mathbb{R}^d$ are independent identically distributed random variables with bounded support, then the length $L(x_1, \dots, x_n)$ of the functional on the points \mathcal{X}_n satisfies

$$L(x_1, \dots, x_n)/n^{\frac{d-1}{d}} \rightarrow \beta_L^d \int_{\mathbb{R}^d} f(x)^{\frac{d-1}{d}} dx,$$

where f is the absolutely continuous part of the distribution of the x_i 's and β_L^d is a constant depending only on d and L (see [3, 28]). Note that this gives an asymptotic formula for $L(x_1, \dots, x_n)$ unless the right hand side is zero. The latter case will happen if the x_i 's lie exclusively on some m -dimensional manifold embedded in \mathbb{R}^d where $m < d$, but the BHH theorem also has a suitable extension to this setting [8], allowing asymptotic formulas involving $n^{\frac{m-1}{m}}$. Our results are all immediately valid in these more general settings, however: as the constants β_L^d depend only on L and d or m (in particular, not on the distribution or, in the second case, the particular manifold), it is enough to study the constants in the case of points which are uniformly distributed in the hypercube.

2. 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 [21]. Suppose that we have a random variable $Z = Z(X_1, X_2, \dots, X_N)$ where X_1, X_2, \dots, 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 $t > 0$,

$$(4) \quad \Pr(|Z - \mathbf{E} Z| \geq t) \leq 2 \exp \left\{ -\frac{t^2}{c^2 N} \right\}.$$

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 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.

The underpinning of our separation results is the fact that for any fixed configuration, a linear number of approximations to that configuration appear in a random point-set with high probability, at the appropriate scale. The intuition for this is perhaps particularly strong in the case of a Poisson cloud of points, but it holds just as well for \mathcal{Y}_n . We capture this essential fact as follows:

Observation 2.1. *Given any finite point set S , any $\varepsilon > 0$, and any D , \mathcal{Y}_n w.h.p. contains at least $C_{\varepsilon, D}^S n$ (ε, D) -copies of S , for some constant $C_{\varepsilon, D}^S > 0$.*

Proof of Observation 2.1. 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, \dots , of side $3D$. Then let $C'_1 \subseteq C_1$ 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, \dots, B_s , $s = |S|$ be the collection of balls of radius ε , centered at each point of S' . Then 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 [21] to show that $Z \geq \frac{1}{2} \mathbf{E} Z$ w.h.p. \square

Observation 2.1 captures all of the probability theory needed for our separation results—in fact, the statement $\mathbf{E} Z \geq \beta n$ from the proof suffices for our separation proofs. The real work of the paper is now discrete geometry—of varying levels of difficulty for the various separations we aim to prove.

For Theorem 1.1, Observation 2.1 is almost all we need; let us now complete the argument.

Observation 2.2. *If $Y \subseteq \mathbb{R}^d$ and x lies in the interior of the convex hull of Y , then when R is sufficiently large, any point at distance greater than R from Y is closer to some point of Y than to x .* \square

If v_0, v_1, \dots, v_k are vectors in \mathbb{R}^d with pairwise negative dot-product, then v_1, \dots, v_k lie in the half-space $v_0 \cdot x < 0$, and the projections of v_1, \dots, v_k onto the hyperplane $v_0 \cdot x = 0$ have pairwise negative dot-products. This gives the following, by induction on d :

Observation 2.3. *If $v_1, \dots, v_{d+1} \in \mathbb{R}^d$ are vectors with negative pairwise dot-products, then 0 is a positive linear combination of the v_i 's.* \square

This allows us to prove:

Lemma 2.4. *If $d + 1 \leq k \leq \tau'(d)$, then there exists a set of points $\bar{S}^{(k)} \subseteq \mathbb{R}^d$ 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 an 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 2.3 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. \square

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 2.4 with k, d' to get a set $S'^{(k)} \subseteq \mathbb{R}^d$. Observe first that the origin must lie in the convex hull X of the set $S'^{(k)}$ given by Lemma 2.4; 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 $S^{(k)}$.



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

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

$$U = \{0\} \cup S^{(k)} \cup \bigcup_{p \in S^{(k)}} \{(1.5)p + .1 \cdot \Delta^d\}.$$

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, $d' = 1$.)

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 hull 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 2.2. Thus Observation 2.1 gives that $\alpha_{k, d} > 0$ for $d + 1 \leq k \leq \tau'(d)$.

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

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 2.1 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 $C'n^{\frac{d-1}{d}}$. \square

Remark 2.5. The same argument allows us to separate β_{MST}^d from β_{Steiner}^d 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 2.6. *Suppose that points p, q, r, s satisfy*

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

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

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

then

$$\|p - s\| \leq \|p - q\| + \|r - s\| + \delta - \frac{\Delta^{1/3}}{12}.$$

Proof. We have

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

Now use $\cos x \leq 1 - x^2/12$ for $x \leq \pi$. □

Lemma 2.7. *Suppose that points $p_i, q_i, r_i, s_i, i = 1, 2$ satisfy*

$$(5) \quad \|p_i - q_i\|, \|r_i - s_i\| \geq \Delta \text{ for } i = 1, 2$$

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

$$(6) \quad \|p_1 - q_1\| + \|r_1 - s_1\| + \|p_2 - q_2\| + \|r_2 - s_2\| + 4\delta - \Delta.$$

Proof. Without loss of generality we let the $q_i, r_i, i = 1, 2$ be within distance δ of the origin. Then points x, y we 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, $\theta(x, y) \leq \frac{1}{2}\pi$ for one of the pairs.

Let us take $\{x, y\}$ and $\{w, z\}$ to be the pairs in such a pairing of P , with $\theta(x, y) \leq \frac{1}{2}\pi$. 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 $a \leq (b^2 + c^2)^{1/2}$ and in fact

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

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 $\|q_i\|, \|r_i\| \leq \delta$ 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 $\geq g + 1$. 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 2.1 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




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.)

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 call a multiset of edges a *2-matching* if every vertex is incident with exactly 2 (not necessarily distinct) edges in the multiset. This is the same as a 2-factor, except that it can contain “2-cycles”, which consist of a single edge included twice.

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

Lemma 2.8. *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*

- (1) F is a 2-matching;
- (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 2.1 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 $\beta_{\text{TF}_g} < \beta_{\text{TF}_{g+1}}$, except that F may not have girth g .

Proof of Lemma 2.8. 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.)




FIGURE 3. An instance of $V_{\epsilon, D, \Delta}$ (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.

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 $U_{\epsilon, D}$ 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 2.6. 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 2.7 to add 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 (note that we may have created a 2-cycle if one of these edges was already present, which is why F is only a 2-*matching*), 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 2.7 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 . \square

We must now address unintentional problems of girth. (Notice that, in shortcutting edges, we may have left behind short cycles—in particular, any 2-cycles must be eliminated.) To this end, we say that $V = V_{\epsilon, D, \Delta}$ 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 2.1 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 a 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 $C = (x = x_1, x_2, \dots, x_k, x_{k+1} = x_1)$ 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 $u, v \notin \{x_1, \dots, x_k\}$, 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 $x_1, x_2, \dots, x_{i-1}x_kx_{k-1} \cdots x_ix_1$ 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, \dots , in any order, and beginning with $F = F_{g+1}$ and for each $i = 1, 2, \dots$, we

- (i) Find F' for V_i as above (with weight increase $C_{g,\varepsilon}$ which we make arbitrarily small)

- (ii) Apply Lemma 2.8 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. \square

The proof of the counterpoint Theorem 1.4 will be given in Section 3. For now we consider matchings. In fact, Theorem 1.5 can be viewed as a consequence of Theorem 1.8, via Proposition 5 of [13]. However, we also give a short self-contained proof.

Proof of Theorem 1.5. We define the Euclidean functional $2\text{MM}(X)$ as the minimum length union of two matchings on X . Note that we make no requirement of disjointness and that we trivially have that $2\text{MM}(X) = 2 \cdot \text{MM}(X)$ 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 $2\text{MM}(\mathcal{Y}_n)$ is asymptotically less than $\text{TSP}(\mathcal{Y}_n)$. 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, \dots, 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 2.1 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 2.7 (with 2ε , say) to get a constant factor improvement a linear number of times. \square

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. \square

Proof of Theorem 1.8. We begin with some general observations regarding the shortest TSP through Euclidean point sets:

Observation 2.9. *Suppose that $S_{\varepsilon,D}$ is an (ε, D) copy of a fixed set S for fixed ε and sufficiently large D , and that at least 2 pairs of edges of a shortest TSP tour \mathcal{L} join $S_{\varepsilon,D}$ to $V \setminus S_{\varepsilon,D}$. Then the pairs are nearly straight (i.e., the angle for each pair is arbitrarily close to π as $\varepsilon \rightarrow 0$, and $k, D \rightarrow \infty$).*

(Here the edges of \mathcal{L} joining $S_{\varepsilon,D}$ to $V \setminus S_{\varepsilon,D}$ are *paired* such that the endpoints of the edges in a pair which are inside $S_{\varepsilon,D}$ are joined by a portion of \mathcal{L} lying entirely in $S_{\varepsilon,D}$.)

Proof. Otherwise, we shortcut the edge pair which is not nearly straight to obtain a constant improvement (which, for a fixed angle, can be made large by increasing

D). The tour portion between one of the other edge pairs is modified to cover any vertices of $S_{\varepsilon,D}$ which are now missed by the tour, at an increased cost which does not depend on D . \square

Observation 2.10. *Suppose that $S_{\varepsilon,D}$ is an (ε, D) copy of any fixed set S for fixed ε and sufficiently large D . Then there are at most 2 pairs of edges in a shortest TSP tour which join $S_{\varepsilon,D}$ to $V \setminus S_{\varepsilon,D}$.*

Proof. Let \mathcal{L} denote a shortest TSP tour, and suppose there are three pairs (e_1, e_2) (f_1, f_2) and (g_1, g_2) of edges in \mathcal{L} between $S_{\varepsilon,D}$ and $V' = V \setminus S_{\varepsilon,D}$. We let $x_1, x_2, y_1, y_2, z_1, z_2$ denote the endpoints of the edges $e_1, e_2, f_1, f_2, g_1, g_2$, respectively, which lie in V' , and we suppose, without loss of generality, that the pairs x_1, y_1, y_2, z_2 , and z_1, x_2 , respectively, are joined by paths in $V'' \cap \mathcal{L}$, for $V'' = V \setminus \{x_1, x_2, y_1, y_2, z_1, z_2\}$.

We now modify \mathcal{L} as follows.

- (1) We remove the edges $e_1, e_2, f_1, f_2, g_1, g_2$.
- (2) We add new edges between the pairs (x_2, y_2) and (y_1, z_1) .
- (3) We add a path which travels from x_1 to the set $S_{\varepsilon,D}$ (visiting every vertex of the set) and then to z_2 .

Observe that the result is a new TSP tour; it remains to estimate the change in cost. The path P added in part 3 has Euclidean length $\ell(P)$ at most $\ell(e_1 - g_2) + C_{S,\varepsilon,D}$, where we are viewing the edges as vectors from their point in V'' to their point in $S_{\varepsilon,D}$, and where $C_{S,\varepsilon}$ is a constant depending only on S, ε . (For example, we can take $C_{S,\varepsilon} = \text{TSP}(S_{\varepsilon,D}) + \text{diam}(S)$.)

Similarly, the edge (x_2, y_2) has length at most $\ell(e_2 - f_2) + C_{S,\varepsilon,D}$, and the edge (y_1, z_1) has length at most $\ell(f_1 - g_1) + C_{S,\varepsilon,D}$. Applying the triangle inequality to the three lengths immediately gives that our new tour has length at most $\text{TSP}(\mathcal{X}_n) + 3C_{S,\varepsilon,D}$. In fact, we should be hoping to do better: If $\ell(e_2 - f_2)$ is within a constant of $\ell(e_2) + \ell(f_2)$ as D grows large, then the points x_2, y_2 are constrained to be nearly antipodal about center c of the set $S_{\varepsilon,D}$ (with angle tending to π as $D \rightarrow \infty$). Similarly, we have that $\ell(f_1 - g_1)$ is far from $\ell(f_1) + \ell(g_1)$ unless y_1, z_1 are nearly antipodal.

Thus if we have not achieved a contradiction by shortening the tour, then, x_2, y_2 are nearly antipodal, and y_1, z_1 are nearly antipodal. Observation 2.9 means that the pairs $\{x_1, x_2\}$, $\{y_1, y_2\}$, and $\{z_1, z_2\}$ are nearly antipodal as well. Thus, in particular, x_1 and z_1 are *not* nearly antipodal, and so we can produce a shortening of the original TSP tour by instead:

- (1) Removing the edges $e_1, e_2, f_1, f_2, g_1, g_2$;
- (2) Adding new edges between the pairs (x_1, z_1) and (x_2, z_2) ;
- (3) Adding a path which travels from y_1 to the set $S_{\varepsilon,D}$ (visiting every vertex of the set) and then to y_2 . \square




FIGURE 4. The set S^{12} . Such sets are good for the HK-bound, but bad for the tour. (Thick lines have weight 1, while thin lines have weight $\frac{1}{2}$.)

We now consider the $d = 2$ case of Theorem 1.8. We let S^k be a set consisting of k equally spaced points on a unit circle centered at the origin, $2k$ equally spaced points at the radius 4 circle centered at the origin, and the points $(2, 0)$ and $(-2, 0)$ (See Figure 4. The particular ratios $2k : k$ and $4 : 2 : 1$ are chosen just to make a clean figure.). We will argue that if ε is sufficiently small and D, k are sufficiently large, each instance of an (ε, D) copy of S^k allows us to locally modify an instance of the TSP so that it is still a solution to the HK linear program, but is shorter by some additive constant.

To this end, let \mathcal{L} be some shortest tour, and $S_{\varepsilon, D}^k$ be an (ε, D) -copy of S^k . Observation 2.10 implies that \mathcal{L} enters (and leaves) $S_{\varepsilon, D}^k$ either once or twice, for sufficiently large D .

Case 1: The tour enters $S_{\varepsilon, D}^k$ once. We let x_1, x_2 denote the vertices in $V' = V \setminus S_{\varepsilon, D}^k$ which are adjacent to vertices in $S_{\varepsilon, D}^k$. We let \mathcal{L}° be the portion of \mathcal{L} between x_1 and x_2 . Then we have that

$$(7) \quad \ell(\mathcal{L}^\circ) \geq \text{dist}(x_1, S_{\varepsilon, D}^k) + \text{dist}(S_{\varepsilon, D}^k, x_2) + 10\pi + 6 + K - o(1),$$

where we are using $o(1)$ to denote a function which tends to 0 as $\varepsilon \rightarrow 0$, $k \rightarrow \infty$, and $D \rightarrow \infty$ simultaneously, and K is an absolute constant (in fact, K can be 2). To see this lower bound, observe that the tour must cover the outer circle ($\approx 8\pi$), the inner circle ($\approx 2\pi$), and must cross the gap between the inner and outer circles twice ($2 \cdot 3 = 6$). Finally, the tour must also spend more (bounded below by some constant K) to cover one of the two “gap” vertices in $S_{\varepsilon, D}^k$. To see that the tour can not cover both gap vertices in gap crossings while crossing the gap only twice, let a_1, a_2 denote the first and last vertices of \mathcal{L}° on the inner circle. Either a_1, a_2 lie at the two ends of the inner circle close to the gap vertices, in which case the tour spends an additive constant K more than $2\pi - o(1)$ to cover the entire inner circle, or, say, a_1 lies at distance $1 + K$ from the gap vertex to which it is joined, incurring an extra cost K again.

We now modify the portion \mathcal{L}° so that the result is still a solution to the Held-Karp LP, but is shorter by some additive constant. We let y_1, y_2 and z_1, z_2 be




FIGURE 5. An example of the set S^{12} with two passes of a tour. (Thick lines have weight 1, while thin lines have weight $\frac{1}{2}$.)

pairs of points on the outer circle which are closest to the gap vertices g_1 and g_2 , respectively, and similarly let a_1, a_2 and b_1, b_2 be points of the inner circle which are closest to the first and second gap vertex, respectively. We join all pairs among each of the triples $y_1, y_2, g_1, a_1, a_2, g_1, z_1, z_2, g_2, b_1, b_2, g_2$ by edges of weight $\frac{1}{2}$ (Figure 4). Next, we let $\alpha_1, \alpha_2, \alpha_3$ be a consecutive triple of points on the outer circle which is as close as possible to x_1 (but disjoint from the set $\{y_1, y_2, z_1, z_2\}$), and $\beta_1, \beta_2, \beta_3$ be a consecutive triple of points on the outer circle which is as close as possible to x_2 (but disjoint from the set $\{y_1, y_2, z_1, z_2, \alpha_1, \alpha_2, \alpha_3\}$). We join all pairs among each of the triples $\alpha_1, \alpha_2, \alpha_3$ and $\beta_1, \beta_2, \beta_3$ with edges of weight $\frac{1}{2}$. We join α_2 and β_2 to x_1 and x_2 , respectively, by edges of weight 1. Finally, using edges of weight 1, we join all consecutive pairs of points on each circle which were not already joined (by edges of weight $\frac{1}{2}$). The result for S^{12} is shown on the right-hand side of Figure 4. As k, D grow large and ε grows small, the total cost of this is

$$\text{dist}(x_1, S_{\varepsilon, D}^k) + \text{dist}(S_{\varepsilon, D}^k, x_2) + 10\pi + 6 + o(1),$$

thus it suffices to check that what we have given is indeed a valid instance of the Held-Karp LP. We carry out this verification below, after describing the second case of the construction.

Case 2: The tour enters $S_{\varepsilon, D}^k$ twice. We let x_1^1, x_2^1 and x_1^2, x_2^2 denote the two pairs of entry/exit points in $V' = V \setminus S_{\varepsilon, D}^k$. We let \mathcal{L}^1 be the portion of \mathcal{L} between x_1^1 and x_2^2 , \mathcal{L}^2 be the portion between x_1^2 and x_2^1 , and let $\mathcal{L}^\circ = \mathcal{L}^1 \cup \mathcal{L}^2$. Then we have that

$$(8) \quad \ell(\mathcal{L}^\circ) \geq \sum_{\substack{i=1,2 \\ j=1,2}} \text{dist}(x_j^i, S_{\varepsilon, D}^k) + 10\pi + 6 + K - o(1),$$

where K is again an absolute constant (and where again, in fact, K can be 2). Again, $10\pi + 6$ is needed to cover both circles, and transition to the inner circle and back. If both \mathcal{L}^1 and \mathcal{L}^2 visit the inner circle this gives an extra cost of ≈ 6 for the transitions, so we assume that \mathcal{L}^1 is the only portion to visit the inner circle. But the argument from the previous case shows that \mathcal{L}^1 cannot cover the entire inner circle and visit both gap vertices without incurring an additive constant extra cost.

We now modify the portion \mathcal{L}° so that the result is still a solution to the Held-Karp LP, as follows. We still let y_1, y_2 and z_1, z_2 be pairs of points on the outer circle which are closest to the gap vertices g_1 and g_2 , respectively, and similarly let a_1, a_2 and b_1, b_2 be points of the inner circle which are closest to the first and second gap vertex, respectively. We join all pairs among each of the triples y_1, y_2, g_1 , a_1, a_2, g_1 , z_1, z_2, g_2 , b_1, b_2, g_2 by edges of weight $\frac{1}{2}$. Next, we let $\alpha_1^i, \alpha_2^i, \alpha_3^i$ be $\beta_1^i, \beta_2^i, \beta_3^i$ be consecutive triples of points on the outer circle which are close to the points x_1^i and x_2^i , respectively. (All named points must be distinct) We join all pairs among each of the triples $\alpha_1^i, \alpha_2^i, \alpha_3^i$ and $\beta_1^i, \beta_2^i, \beta_3^i$ with edges of weight $\frac{1}{2}$ for $i = 1, 2$. We join α_2^i and β_2^i to x_1^i and x_2^i , respectively, for $i = 1, 2$, by edges of weight 1. Finally, using edges of weight 1, we join all consecutive pairs of points on each circle which were not already joined (by edges of weight $\frac{1}{2}$). The result for S^{12} is shown on the right-hand side of Figure 5. As k, D grow large and ε grows small, the total cost of this is

$$\sum_{\substack{i=1,2 \\ j=1,2}} \text{dist}(x_j^i, S_{\varepsilon, D}^k) + 10\pi + 6 + o(1),$$

and so we have improved the length by an additive constant.

Feasibility of the solutions. We now check that making many local modifications according to the cases above does not disturb the property that \mathcal{L} is a feasible instance of the Held-Karp LP. It is immediate that all degree-weights $\sum_{j \neq i} x_{\{ij\}}$ are 2; it remains to check the condition that

$$(9) \quad (\forall \emptyset \subsetneq S \subsetneq V) \quad \sum_{\{i,j\} \subseteq S} x_{\{ij\}} \leq |S| - 1.$$

In other words, the total weight of edges in any proper nonempty subset is at least one less than the number of vertices.

Since the degree-weights are 2 at every vertex, we can show (9) by showing that any proper nonempty subset S has the property that edges of total weight at least 2 leave the subset; i.e., that

$$(10) \quad (\forall \emptyset \subsetneq S \subsetneq V) \quad \sum_{\substack{i \in S \\ j \notin S}} x_{\{ij\}} \geq 2.$$

If this fails, there is a cut in the graph of weight < 2 . First we consider the possibility that the cut includes an edge of weight $\frac{1}{2}$. Such edges only occur in triangles, and a minimum cut in a graph cannot contain exactly one edge of any triangle. But this already implies that no cut of weight < 2 in our graph can include any edges of weight $\frac{1}{2}$, since deleting even all the edges of any one triangle does not disconnect our graph, and since 4 such edges, or 2 such edges plus an edge of weight 1, already gives weight 2.

But now we are done: since no minimum cut of weight < 2 can include an edge of weight $\frac{1}{2}$ (and all other edges have weight 1), it suffices to note that there is no single cut edge in our modified graph. This completes the proof for the case $d = 2$.

The case $d > 2$. We cannot use exactly the same point set in higher dimensions. The only trouble with the previous argument is the lower bound on the tour length. In particular, if an edge enters a set S^k at a sharp angle to the 2-D plane containing S^k , then it may join to an inner circle vertex or gap vertex at negligible (or zero) extra cost over entering at the nearest possible point, and this is not accounted for in (7) or (8). Our goal now is to create a set out of many copies of S^k , in a way that allows us to be certain that some copy of S^k must be incident only with edges nearly parallel to its containing plane, in any optimal tour.

We will let $S_{\varepsilon,D,R}^{k,\ell}$ be an (ε, D) copy of a certain elaborate set $S_R^{k,\ell}$. The set $S_R^{k,\ell}$ consists of ℓ copies of S^k . Each copy of S^k lies in a 2-dimensional hyperplane normal to some unit vector v , and we orient the copies of S^k so that these hyperplanes are tangent to the $(d-1)$ -sphere S_R of radius R . Moreover, we ensure that the centers c_1, \dots, c_ℓ lie on S_R , and are roughly “evenly spaced” in the sense that as ℓ grows large, we can ensure that for any given unit vector $x \in S_1$, we have that x lies arbitrarily close to $\frac{1}{R}c_i$ for one of the c_i ’s.

We now appeal to the following fact:

Lemma 2.11. *If X is a subset of the unit $(d-1)$ -sphere S_1 in \mathbb{R}^d of cardinality n , then $\text{TSP}(X) = o(n)$.*

Proof. We can rescale the sphere and embed it in the hypercube $[0, 1]^d$, and then appeal to the fact (see e.g., [10, 29]) that the *worst-case* length of a tour through n points in $[0, 1]^d$ is $\leq C_d n^{\frac{d-1}{d}}$ for some constant C_d . (It is also easy to obtain a bound $C'_d n^{\frac{d-2}{d-1}}$ for the sphere, but this is unnecessary for us.) \square

In particular, this lemma implies that if we take $R \gg \ell \gg k$, then typical edges joining pairs of instances of (approximate) S^k ’s in the set $S_{\varepsilon,D,R}^{k,\ell}$ are of length $o(R)$ (where the asymptotics are as $R \rightarrow \infty$). Furthermore, Observation 2.10 implies that the tour enters $S_{\varepsilon,D,R}^{k,\ell}$ at most twice, and also that it enters each S^k at most twice. In particular, for sufficient choices $R \gg \ell \gg k$, there will necessarily be at least one set X , which is an approximate S^k in the set $S_{\varepsilon,D,R}^{k,\ell}$ such that

- (1) All edges entering/leaving X go to other points of $S_{\varepsilon,D,R}^{k,\ell}$, and
- (2) All such edges have length $o(R)$.

In particular, the edges entering and leaving X will be nearly parallel to the 2-dimensional hyperplane in which X (approximately) lies. This is because

- (a) the portion of the sphere within $o(R)$ of X will lie within a small angular distance of X , and moreover,
- (b) the portion of the ε -neighborhood of the sphere within distance $o(R)$ of X but also at distance at least some large constant from X is also within a small angular distance of X . (Observe that points in $S_{\varepsilon,D,R}^{k,\ell} \setminus X$ are a large constant distance from X by our choice of $R \gg \ell$.)

But this is now sufficient to ensure that we can modify the tour in X (according to Case 1 or 2 from above) to obtain a constant additive improvement. \square

3. THE 2-FACTOR LIMIT

Here we prove Theorem 1.4: $\lim_{g \rightarrow \infty} \beta_{TF_g}^d = \beta_{TSP}^d$. Let F_g denote a minimum length 2-factor of girth g covering \mathcal{X}_n . Let the cycles of F_g be C_1, C_2, \dots, C_k where $k \leq n/g$ and let $L_1 \sim \beta_{TF_g}^d n^{(d-1)/d}$ denote the length of F_g . Now choose a vertex $x_i \in C_i$ for $i = 1, 2, \dots, k$ and let C_0 be a minimum length tour through $X = \{x_1, x_2, \dots, x_k\}$. It follows from Few [10] and Tóth[33] that C_0 has length at most $L_0 = C(n/g)^{(d-1)/d}$ for some absolute constant $C > 0$.

The graph Γ induced by the edges of C_0, C_1, \dots, C_k is connected and has even degrees i.e. it is Eulerian. We obtain a tour through \mathcal{X}_n of length at most $L_0 + L_1$ by traversing an Euler tour through g and short-cutting when a vertex is re-visited. This gives a tour of length at most $(1+o(1))\beta_{TF_g}^d n^{(d-1)/d} + C(n/g)^{(d-1)/d}$. Allowing $g \rightarrow \infty$ proves the theorem. \square

4. BRANCH AND BOUND ALGORITHMS

In this section we prove Theorem 1.9. Branch-and-bound is a pruning process, which can be used to search for an optimum TSP tour. Branch-and-bound as we consider here depends on 3 choices:

- (1) A choice of heuristic to find (not always optimal) TSP tours;
- (2) A choice of lower bound for the TSP (such as the 2-factor, HK bound, etc.);
- (3) A branching strategy (giving a *branch-and-bound* tree).

As an example, we will consider the case where we use some heuristic for the TSP, the 2-factor as a lower bound, and use a branching strategy based on the 2-factor as well.

Given our point-set \mathcal{X}_n , we begin by letting B be the value of the tour found by our TSP heuristic. We let b_x be the length of the shortest 2-factor in \mathcal{X}_n . Here x represents the root of the branch-and-bound tree, which we will construct iteratively; Λ_x is the set of all TSP tours in \mathcal{X}_n . Unless $b_x \geq B$, we do not know that B is an optimal tour, so we *branch* in the following way: we choose some cycle C in the 2-factor we have found, and, for the edges e_1, \dots, e_k of C , generate k children x_1, \dots, x_k of x , letting

$$\begin{aligned} I_{x_1} &= \emptyset, I_{x_2} = \{e_1\}, I_{x_3} = \{e_1, e_2\}, \dots \\ O_{x_1} &= \{e_1\}, O_{x_2} = \{e_2\}, O_{x_3} = \{e_3\}, \dots \end{aligned}$$

These are sets of required *inclusions* and *exclusions*, respectively. In particular, for any vertex v of our tree, Λ_v is the set of tours containing all edges in I_v and avoiding all edges in O_v . (For the root, we had $I_x = O_x = \emptyset$.) Thus, in this example, the vertex x_2 corresponds to the set of TSP tours which do contain e_1

but do not contain e_2 . For each x_i , we use our TSP heuristic to find a tour, with the additional constraints that the tour includes all edges in I_{x_i} and excludes those in O_{x_i} . Whenever we find a TSP tour shorter than the current value of B , we update B . We also, for each x_i , let b_{x_i} be the minimum-length 2-factor subject to the constraints I_{x_i}, O_{x_i} . For any b_{x_i} for which b_{x_i} is at least B , we know that no shorter tour than B exists subject to I_{x_i}, O_{x_i} , and the tree is *pruned* below x_i , so that x_i becomes a leaf of the pruned branch-and-bound tree. For other vertices, we continue to branch as above, by adding further constraints to kill some other cycle of the minimum 2-factors found.

This process terminates when the set L of leaves of the pruned branch-and-bound tree satisfies $v \in L \implies b_v \geq B$; such a tree corresponds to a certificate that the best TSP tour found so far by our heuristic is indeed optimum.

In general, as a *branching strategy*, we allow any method to produce, given an input weighted graph, a rooted tree (the branch-and-bound tree) labeled with sets I_v, O_v of edges from the graph such that:

- (1) When v is a child of u , $I_v \supseteq I_u$ and $O_v \supseteq O_u$.
- (2) If the children of u are v_1, \dots, v_k , then we have $\Lambda_u = \bigcup_{i=1}^k \Lambda_{v_i}$.
- (3) The leaves of the (unpruned) branch-and-bound tree satisfy $|\Lambda_v| = 1$.

Following any such branching strategy and pruning when $b_v \geq B$ will eventually lead to a proof that an optimum tour has been found (assuming a reasonable TSP heuristic and lower bound), since, in the worst case, the heuristic and lower bound will match on the leaf v for which Λ_v contains just the optimum tour. For branch-and-bound to be efficient, we would hope that all but polynomially many vertices of the branch-and-bound tree can be pruned because of inequalities $b_v \geq B$.

We can restate Theorem 1.9 more precisely as follows

Theorem 1.9 (restated). *For any TSP heuristic, any branching strategy, and a lower bound heuristic which is TF_g or HK, the pruned branch-and-bound tree will have $e^{\Omega(n/\log^6 n)}$ leaves a.s.*

In particular, in our proof of Theorem 1.9, we will make the most optimistic assumption regarding the TSP heuristic: we will simply assume that it always returns an optimum tour (B will always be the true value of the minimum TSP). Theorem 1.9 asserts that even in this case, there can be no polynomially-sized branch-and-bound tree which certifies optimality w.h.p.

One natural strategy to prove that L is large would be to show that each Λ_v ($v \in L$) is small; then, $\Lambda = \bigcup_{v \in L} \Lambda_v$ would give that L must be large. To show that Λ_v is small, one can hope to argue that to have $\text{LB}(\mathcal{X}_n | I_v, O_v) \geq \text{TSP}(\mathcal{X}_n)$, either I_v or O_v must be large, severely restricting the number of tours in Λ_v . The problem is that while large I_v does restrict the size of Λ_v considerably, having a large O_v can be a rather weak restriction.

We will thus modify this basic approach by paying attention to a special set of tours $\bar{\Lambda}$. Given the point set \mathcal{X}_n , we will consider the division of $[0, 1]^d$ into $s = \frac{n}{K \log n}$ boxes of sidelength $\left(\frac{K \log n}{n}\right)^{\frac{1}{d}}$ where K is at least some sufficiently large constant. B_1, B_2, \dots, B_s denote these boxes, taken in some order such that consecutive terms are adjacent (i.e., sharing a $(d-1)$ -dimensional face). Note that

$$|x - y| \leq \sqrt{d} \left(\frac{K \log n}{n}\right)^{\frac{1}{d}} \text{ if } x, y \text{ lie in the same box.}$$

We consider $\mathcal{X}_n = \{x_1, x_2, \dots, x_n\}$, and, for for each $2 \leq j \leq s-1$, we let $x_j^1, x_j^2, x_j^3, x_j^4$ denote the four points $x_i \in \mathcal{X}_n \cap B_j$ of smallest index (this particular choice is arbitrary, and is just for definiteness). We also choose points $x_1^3, x_1^4 \in \mathcal{X}_n \cap B_1$ and $x_s^1, x_s^2 \in \mathcal{X}_n \cap B_s$, again by simply choosing points of minimum index. Letting

$$(11) \quad \mathcal{I} = \{x_i^1, x_i^2 \mid 1 < i \leq s\} \cup \{x_i^3, x_i^4 \mid 1 \leq i < s\},$$

the points in \mathcal{I} can be viewed as preselected ‘‘interface points’’ between the boxes B_j . In particular, we let $\bar{\Lambda}$ denote the set of TSP tours in \mathcal{X}_n with the properties that, in the tour,

- (1) x_1^4 is joined to x_1^3 by a path lying entirely in B_1 ;
- (2) for $1 \leq j \leq s-1$, x_j^3 and x_{j+1}^1 are adjacent;
- (3) for $2 \leq j \leq s-1$, x_j^1 is joined to x_j^3 by a path lying entirely in B_j ;
- (4) x_s^1 is joined to x_s^2 by a path lying entirely in B_s ;
- (5) for $s \geq j \geq 2$, x_j^2 and x_{j-1}^4 are adjacent; and
- (6) for $s-1 \geq j \geq 2$, x_j^2 is joined to x_j^4 by a path lying entirely in B_j .

Note that tours in $\bar{\Lambda}$ use only edges of length $\leq 2\sqrt{d} \left(\frac{K \log n}{n}\right)^{\frac{1}{d}}$. Instead of using the full strength of the condition $\Lambda = \bigcup_{v \in L} \Lambda_v$, we will use the weaker (but apparently more useful) condition

$$\bar{\Lambda} = \bigcup_{v \in L} \bar{\Lambda}_v$$

where $\bar{\Lambda}_v = \bar{\Lambda} \cap \Lambda_v$. Intuitively, we are focusing our attention on a restricted set of tours (chosen such that the value of LB at relevant leaves $v \in L$ with nonempty $\bar{\Lambda}_v$ will be close to its typical length), and this restricted set $\bar{\Lambda}$ has the property that the allowable set of edges at each vertex is now small enough that having a large excluded set O_v really will force $\bar{\Lambda}_v = \bar{\Lambda} \cap \Lambda_v$ to be small.

Our proof will require us to analyze the performance of LB conditioned on the exclusions I_v and exclusions O_v . Thus we begin by proving that some simple operations preserve the property of being an feasible instance of the Held-Karp LP.

Lemma 4.1. *Suppose we are given a feasible instance of the Held-Karp LP on a set X of cardinality n . Then the following operations will all result in feasible instances of the Held-Karp LP:*

- (1) Downweight all edges by a factor of $(1 - \frac{1}{n})$, add a new pair of vertices y, z joined by an edge of weight 1, and join each of y, z to all vertices in X by edges of weight $\frac{1}{n}$.
- (2) For $k < \frac{n}{2}$, downweight all edges by a factor of $(1 - \frac{k}{n})$, add k new vertices y_1, \dots, y_k , and join each y_i to all vertices in X by edges of weight $\frac{2}{n}$.

Proof. It is immediate for both operations that the (weighted) degree of every vertex is 2; we need to check that

$$(12) \quad \forall \emptyset \subsetneq S \subsetneq V \quad \mathcal{S}(S) := \sum_{e \subseteq S} x_e \leq |S| - 1.$$

For Operation 1, observe first that if $\emptyset \subsetneq S \subsetneq X$, we have $\mathcal{S}(S) \leq (1 - \frac{1}{|X|})(|S| - 1)$ and $\mathcal{S}(X) = |X| - 1$. For $\emptyset \subsetneq S \subsetneq X$, we have

$$\mathcal{S}(S \cup \{y\}) \leq (1 - \frac{1}{|X|})(|S| - 1) + |S| \frac{1}{|X|} = |S| + \frac{1}{|X|} - 1 \leq |S \cup \{y\}| - 1.$$

and

$$\mathcal{S}(S \cup \{y, z\}) \leq (1 - \frac{1}{|X|})(|S| - 1) + |S| \frac{2}{|X|} + 1 = |S| + \frac{|S|+1}{|X|} \leq |S \cup \{y, z\}| - 1.$$

And $\mathcal{S}(X \cup \{y\}) = \mathcal{S}(X \cup \{z\}) = |X|$.

Finally, for Operation 2, we have that $\mathcal{S}(S) \leq (1 - \frac{k}{|X|})(|S| - 1)$ if $\emptyset \subsetneq S \subsetneq X$, and $\mathcal{S}(X) = (1 - \frac{k}{n})(|X|) = |X| - k$. If $Y \subseteq \{y_1, y_2, \dots, y_k\}$ and $\emptyset \subsetneq S \subseteq X$ then

$$(13) \quad \begin{aligned} \mathcal{S}(S \cup Y) &\leq (1 - \frac{k}{|X|})(|S| - \mathbf{1}_{S \neq X}) + |S| \cdot |Y| \cdot \frac{2}{|X|} \\ &= |S| + \frac{(2|Y| - k)|S|}{|X|} - \mathbf{1}_{S \neq X} \left(1 - \frac{k}{|X|}\right) \\ &\leq |S| + (|Y| - \mathbf{1}_{|Y| < k}) \left(1 - \frac{\mathbf{1}_{S \neq X}}{|X|}\right) - \mathbf{1}_{S \neq X} \left(1 - \frac{k}{|X|}\right), \end{aligned}$$

and we see that $(Y \subsetneq \{y_1, \dots, y_k\} \vee S \neq X) \implies \mathcal{S}(S \cup Y) \leq |S| + |Y| - 1$. \square

To know that the lower bound $\text{LB} = \text{HK}$ performs well at leaves $v \in L$ such that $\bar{\Lambda}_v \neq \emptyset$, we will also want to patch several smaller solutions to the Held-Karp LP into a single global solution using the same edges which tours in $\bar{\Lambda}$ use to cross the B_i 's:

Lemma 4.2. *Suppose we are given feasible solutions to the Held-Karp LP on disjoint sets X_1, X_2, \dots, X_s of cardinalities n_1, n_2, \dots, n_s . Write $X_i = \{x_i^1, \dots, x_i^{n_i}\}$, and suppose that the edge $\{x_i^1, x_i^2\}$ has weight 1 for all $2 \leq i \leq s$ and that the edge $\{x_i^3, x_i^4\}$ has weight 1 for all $1 \leq i \leq s - 1$. Then we can patch these solutions into a feasible solution on $\bigcup_{i=1}^s X_i$ by*

- (a) Deleting all the edges $\{x_i^1, x_i^2\}$ ($2 \leq i \leq s$) and $\{x_i^3, x_i^4\}$ ($1 \leq i \leq s - 1$)
(b) Joining x_i^3 to x_{i+1}^1 and x_i^4 to x_{i+1}^2 with edges of weight 1, for all $1 \leq i \leq s - 1$

Proof. By induction on s , it suffices to handle the case $s = 2$. (For $s > 2$, first patch together the X_1, \dots, X_{s-1} , to obtain a Held-Karp solution containing $\{x_1^1, x_1^2\}$ and $\{x_{s-1}^3, x_{s-1}^4\}$ with weight 1, then patch this with X_s .)

It is apparent that after the patching operation, the degree of every vertex is still 2. Suppose now that $\emptyset \subsetneq S \subsetneq X_1 \cup X_2$ and write $S_i = S \cap X_i$ for $i = 1, 2$. Let $0 \leq \rho \leq 2$ denote the number of edges from among $\{x_1^3, x_2^1\}, \{x_1^4, x_2^2\}$ whose endpoints are both in the set S . Note that

$$\mathcal{S}(S) = \mathcal{S}(S_1) + \mathcal{S}(S_2) + \rho.$$

We let $\bar{\mathcal{S}}(S_i)$ denote the weight of the HK instance on S_i before the patching operation (of course, $\mathcal{S}(S_i) \leq \bar{\mathcal{S}}(S_i)$) and consider cases:

Case 1: $\rho = 0$

We have $S_i \neq X_i$ for some $i \in \{1, 2\}$. For at least one i , $S_i \neq \emptyset$, giving

$$\mathcal{S}(S) \leq \bar{\mathcal{S}}(S_1) + \bar{\mathcal{S}}(S_2) \leq |S_1| + |S_2| - 1 = |S| - 1.$$

Case 2: $\rho = 1$

Now $\rho < 2$ implies that $S_1 \neq X_1$ or $S_2 \neq X_2$, while $\rho > 0$ implies that $S_1 \neq \emptyset$ and $S_2 \neq \emptyset$. If in fact both $S_1 \neq X_1$ and $S_2 \neq X_2$, then we have

$$\mathcal{S}(S) \leq \bar{\mathcal{S}}(S_1) + \bar{\mathcal{S}}(S_2) + 1 \leq |S_1| - 1 + |S_2| - 1 + 1 = |S| - 1,$$

while if (without loss of generality) $S_1 = X_1$ and $S_2 \neq X_2$, then the deletion step (a) implies that $\mathcal{S}(S_1) = \bar{\mathcal{S}}(S_1) - 1$, and so we have

$$\mathcal{S}(S) \leq \bar{\mathcal{S}}(S_1) - 1 + \bar{\mathcal{S}}(S_2) + 1 \leq |S_1| - 1 + |S_2| - 1 + 1 = |S| - 1.$$

Case 3: $\rho = 2$

In this case, we have that the deletion step (a) implies that $\mathcal{S}(S_i) \leq \bar{\mathcal{S}}(S_i) - 1$. In particular, we have

$$\mathcal{S}(S) \leq \bar{\mathcal{S}}(S_1) - 1 + \bar{\mathcal{S}}(S_2) - 1 + 2 = \bar{\mathcal{S}}(S_1) + \bar{\mathcal{S}}(S_2)$$

and thus $\mathcal{S}(S) \leq |S| - 1$ if $\emptyset \neq S \neq X_1 \cup X_2$, since we must have either $\emptyset \neq S_1 \neq X_1$ or $\emptyset \neq S_2 \neq X_2$ ($S_i = \emptyset$ is not possible since $\rho > 0$). \square

Next we prove a concentration lemma, which is a simple modification of what appears in [30]. This will allow us to argue that modest conditioning does not significantly alter the value of LB in the leaves of interest.

Let $J \subseteq \{1, \dots, n\}$ be any fixed set of indices, and let $X_J \subseteq \mathcal{X}_n = \{x_1, x_2, \dots, x_n\}$ denote the random set $\{x_j \mid j \in J\}$. Recalling \mathcal{I} from (11), we let

$$\text{LB}_i(J) = \text{LB}((B_i \cap X_J) \setminus \mathcal{I}).$$

In particular, $\text{LB}_i(J)$ is the value of $\text{LB}(X_J)$ restricted to the box B_i , after throwing away the (two or four) special points x_i^j .

Lemma 4.3. *For $\widehat{\text{LB}}(J) = \sum_i \text{LB}_i(J)$, where $\text{LB} = \text{TF}_g$ or HK , and $|J| = \eta$, we have*

$$(a) \quad \mathbf{E} \widehat{\text{LB}}(J) \leq \beta_{\text{LB}} \eta^{(d-1)/d} + o(\eta^{(d-1)/d}).$$

(b) *There is an absolute constant $c > 0$ such that,*

$$\Pr(\widehat{\text{LB}}(J) \geq \mathbf{E} \widehat{\text{LB}}(J) + t) \leq \exp \left\{ -\frac{ct^2}{\eta^{(d-2)/d} \log^{2/d} \eta} \right\}.$$

Proof.

(a) Suppose that $X_J \cap B_i$ contains n_i points for $i = 1, 2, \dots, s$. Then, in particular, $X_J \cap B_i \setminus \mathcal{I}$ is $n_i - 2$ or $n_i - 4$ for each i , and we have

$$\mathbf{E} \widehat{\text{LB}}(J) \leq (1 + o(1)) \sum_{j=1}^s \frac{\beta_{\text{LB}} n_j^{(d-1)/d}}{s^{1/d}} \leq (1 + o(1)) \frac{\beta_{\text{LB}}}{s^{1/d}} \times s \left(\frac{\sum_{j=1}^s n_j}{s} \right)^{(d-1)/d},$$

where we have used Jensen's inequality.

(b) Assume first that $\text{LB} = \text{TF}_g$. Let $J = \{j_1, j_2, \dots, j_\eta\}$. Let

$$d_i = \max_{j_i, \hat{j}_i} |\widehat{\text{LB}}(\{j_1, j_2, \dots, j_i, \dots, j_\eta\}) - \widehat{\text{LB}}(\{j_1, j_2, \dots, \hat{j}_i, \dots, j_\eta\})|.$$

The Azuma-Hoeffding inequality implies that

$$(14) \quad \Pr(\widehat{\text{LB}} \geq \mathbf{E} \widehat{\text{LB}} + t) \leq \exp \left\{ -\frac{t^2}{2 \sum_{i=1}^n |d_i|^2} \right\}.$$

Now fix $j_1, j_2, \dots, j_\eta, \hat{j}_i$ and suppose that $j_i \in B_k$ and $\hat{j}_i \in B_l$. Then let

$$\Delta = |\widehat{\text{LB}}(j_1, j_2, \dots, j_i, \dots, j_\eta) - \widehat{\text{LB}}(j_1, j_2, \dots, \hat{j}_i, \dots, j_\eta)|.$$

Let $F = F_1 \cup \dots \cup F_s$ where F_i is the optimal 2-factor for $(X_J \cap B_j) \setminus \mathcal{I}$, for $i = 1, 2, \dots, s$. Suppose that the neighbors of $x = x_{j_i}$ 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 in B_k 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 at most $2d^{1/2} \left(\frac{K \log n}{n} \right)^{\frac{1}{d}}$. 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, of F_l that contains the point c of $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n$ closest to \hat{x} . Thus,

$$(15) \quad d_i \leq 4d^{1/2} \left(\frac{K \log n}{n} \right)^{\frac{1}{d}}.$$

This proves the lemma for $\text{LB} = \text{TF}_g$.

Assume now that $\text{LB} = \text{HK}$. Here we can use the results of Goemans and Bertsimas [13]. Proposition 3 and Lemma 11 of the same paper shows that $d_i \leq 2 \min_{j \neq i} |\hat{x}_i - x_j|$ and the proof goes through as before. \square

We are ready to proceed with the proof of Theorem 1.9. We consider the sizes of $\bar{\Lambda}$ and $\bar{\Lambda}_v = \bar{\Lambda} \cap \Lambda_v$. We let $\beta_j = |\mathcal{X}_n \cap B_j|$. Then we have that

$$|\bar{\Lambda}| = (\beta_1 - 2)! \left(\prod_{j=2}^{s-1} (\beta_j - 3)! \right) (\beta_s - 2)!.$$

Now, given I_v , we let $I_v^j \subseteq I_v$ denote those edges in I_v whose endpoints both lie in B_j , and $I'_v \subseteq I_v$ denote those edges in I_v of the form $\{x_j^3, x_{j+1}^1\}$ or $\{x_j^2, x_{j+1}^4\}$.

Given O_v , we let O_v^j denote the set of those edges in O_v whose endpoints both lie in B_j and are not equal to the edges $\{x_j^1, x_j^2\}$, $\{x_j^3, x_j^4\}$, $\{x_j^1, x_j^4\}$, or $\{x_j^2, x_j^3\}$. Observe now that $\bar{\Lambda}_v = \emptyset$ unless $I_v = I_v^s \cup \bigcup_{j=1}^s I_v^j$ and $O_v = \bigcup_{j=1}^s O_v^j$. We now observe that

$$(16) \quad |\bar{\Lambda}_v| \leq (\beta_1 - 2 - |I_v^1|)! 2^{|I_v^1|} \left(\prod_{j=2}^{s-1} (\beta_j - 3 - |I_v^j|)! 2^{|I_v^j|} \right) (\beta_s - 2 - |I_v^s|)! 2^{|I_v^s|}$$

and, letting $\delta_A = 1$ when $|A| \geq 1$, and 0 otherwise:

$$(17) \quad |\bar{\Lambda}_v| \leq (\beta_1 - 2 - \delta_{O_v^1})(\beta_1 - 3)! \left(\prod_{j=2}^{s-1} (\beta_j - 3 - \delta_{O_v^j})(\beta_j - 4)! \right) (\beta_s - 2 - \delta_{O_v^s})(\beta_s - 3)!$$

since, e.g., the number of ways of covering $K_{[\beta]}$ with paths from 1 to 3 and 2 to 4, respectively, while avoiding an edge e which is not $\{1, 2\}$, $\{3, 4\}$, $\{1, 4\}$, or $\{2, 3\}$ is exactly either $(\beta - 3)! - (\beta - 4)!$ or $(\beta - 3)! - 2(\beta - 4)!$, depending, respectively on whether or not e is incident with a vertex in $\{1, 2, 3, 4\}$.

Observe now that the Chernoff bounds give that w.h.p all β_j 's satisfy $\beta_j < 2K \log n$. In particular, there must be at least $|O_v| \left(\frac{1}{2K \log n} \right)^2$ j 's such $|O_v^j| \geq 1$, so that (17) gives

$$(18) \quad |\bar{\Lambda}_v| \leq |\bar{\Lambda}| \cdot \left(1 - \frac{1}{2K \log n} \right)^{|\bar{O}_v| \left(\frac{1}{2K \log n} \right)^2} \leq e^{-|O_v|/(2K \log n)^3}$$

where $\bar{O}_v = \bigcup_{j=1}^s O_v^j$. Also, for $\bar{I}_v = \bigcup_{j=1}^s I_v^j$, (16) gives (very crudely) that, say,

$$(19) \quad |\bar{\Lambda}_v| \leq |\bar{\Lambda}| \cdot \left(\frac{1}{3} \right)^{-|I_v|} \leq \bar{\Lambda} \cdot e^{-|I_v|}$$

Now (18) and (19) establish that large \bar{I}_v or \bar{O}_v forces $\bar{\Lambda}_v$ to be small. (Note that this part of the argument would have failed if we were working with Λ_v 's in place of $\bar{\Lambda}_v$'s.) Thus, defining $\bar{L} = \{v \in L \mid \bar{\Lambda}_v \neq \emptyset\}$, we have that $\bar{\Lambda} = \bigcup_{v \in \bar{L}} \bar{\Lambda}_v$. In particular, since $\bar{\Lambda}$ is large, we can show that the set of leaves L of the branch and bound tree must be large (in fact, that $\bar{L} \subseteq L$ is large) by showing that $v \in \bar{L}$ implies that either \bar{I}_v or \bar{O}_v is large. This is where we use separation of constants. Indeed, we will prove:

Lemma 4.4. *Let LB be TF_g or HK. We have w.h.p. that for all $v \in \bar{L}$, either $|\bar{I}_v| + |\bar{O}_v| > \frac{n}{\log^3 n}$, or else*

$$\text{LB}(\mathcal{X}_n | I_v, O_v) \leq \beta_{\text{LB}} n^{\frac{d-1}{d}} + C(|\bar{I}_v| + |\bar{O}_v|) \left(\frac{\log n}{n} \right)^{\frac{1}{d}} + o(n^{\frac{d-1}{d}})$$

for some constant C .

To use the lemma to complete the proof of Theorem 1.9, we observe that $v \in \bar{L}$ implies that $\text{LB}(\mathcal{X}_n) \geq \text{TSP}(\mathcal{X}_n)$ and thus, from Lemma 4.4, that either $|\bar{I}_v| + |\bar{O}_v| > \frac{n}{\log^3 n}$, or else that

$$|\bar{I}_v| + |\bar{O}_v| \geq (\beta_{\text{TSP}} - \beta_{\text{LB}}) \frac{n}{\log n} - o\left(\frac{n}{\log n}\right)$$

In particular, in the latter case, the separation $\beta_{\text{TSP}} > \beta_{\text{LB}}$ gives that

$$|\bar{I}_v| + |\bar{O}_v| = \Omega\left(\frac{n}{\log n}\right)$$

In either case, (18) and (19) will then give that

$$|\bar{\Lambda}_v| \leq |\bar{\Lambda}| e^{-\Omega(n/\log^6 n)}$$

and thus that

$$\bar{L} \geq e^{\Omega(n/\log^6 n)},$$

completing the proof of Theorem 1.9. \square

Proof of Lemma 4.4. We have that $v \in \bar{L}$ implies that $I_v = \bar{I}_v \cup I'_v$ and $O_v = \bar{O}_v$. Now let

$$\mathcal{Y}_n^v = \mathcal{X}_n \setminus (V(\bar{I}_v) \setminus V(\bar{O}_v)).$$

Letting $J(X)$ for $X \subseteq \mathcal{X}_n$ denote the set such that $X = \{x_j \mid j \in J(X)\}$, Lemma 4.3 gives that

$$(20) \quad \Pr \left(\exists S \subseteq \mathcal{X}_n, |S| \leq \frac{n}{\log^3 n} : \widehat{\text{LB}}(J(\mathcal{X}_n \setminus S)) \geq (\beta_{\text{LB}} + \varepsilon)n^{(d-1)/d} \right) \\ \leq \binom{n}{n/\log^3 n} \exp \left\{ -\frac{c\varepsilon^2 n^{2(d-1)/d}}{4n^{(d-2)/d} \log^{2/d}} \right\} \leq \exp \left\{ \frac{n}{\log^2 n} - \frac{c\varepsilon^2 n}{4 \log^{2/d} n} \right\} = o(1).$$

Thus, taking, e.g., $\varepsilon = \frac{1}{\sqrt{\log n}}$ we have that w.h.p all leaves $v \in \bar{L}$ satisfy

$$\widehat{\text{LB}}(J(\mathcal{Y}_n^v)) < \beta_{\text{LB}} n^{(d-1)/d} + o(n^{(d-1)/d}).$$

Recall that an instance of $\widehat{\text{LB}}$ consists of independent instances H_j of LB in each $B_j \cap \mathcal{Y}_n \setminus \mathcal{I}$. Now, for each j , we patch into H_j the edges in $\bar{I}_v \cap B_j$ (using Operation 1 of Lemma 4.1, if LB = HK), the endpoints of the edges in $\bar{O}_v \cap B_j$ (using Operation 2 of Lemma 4.1, if LB = HK), and the (one or two, depending in j) edges

$$\begin{aligned} &\{x_j^1, x_j^2\} \text{ for } 2 \leq j \leq s \\ &\{x_j^3, x_j^4\} \text{ for } 1 \leq j \leq s-1 \end{aligned}$$

(again using Operation 1, if LB = HK). In the case where LB = TF_g the patching is simply accomplished by rerouting cycles through the edges and points. Note that, as the squares B_j have diameter $\sqrt{d} \left(\frac{K \log n}{n} \right)^{1/d}$, the total increase in cost due to this patching is $\leq C_d (|\bar{I}_v| + |\bar{O}_v|) \left(\frac{K \log n}{n} \right)^{1/d}$ for some constant C_d .

Next, we patch the resulting solutions together at the points x_i^j to a global instance of the LB (using Lemma 4.2, if LB = HK). It is important to note that by beginning with an instance of $\widehat{\text{LB}}$ and then patching to this particular instance of LB, we are guaranteed this instance includes any edges in I'_v . Since there are $O\left(\frac{n}{\log n}\right)$ squares of diameter $O\left(\sqrt{d} \left(\frac{\log n}{n}\right)^{\frac{1}{d}}\right)$, the total increased cost from this patching is $o(n^{(d-1)/d})$, and so we have shown that

$$(21) \quad \text{LB}(\mathcal{X}_n | \bar{I}_v \cup I'_v, O_v) < \widehat{\text{LB}}(\mathcal{Y}_n^v) + C_d (|\bar{I}_v| + |\bar{O}_v|) \left(\frac{\log n}{n} \right)^{1/d} + o(n^{(d-1)/d}) \\ < \beta_{\text{LB}} n^{(d-1)/d} + C_d (|\bar{I}_v| + |\bar{O}_v|) \left(\frac{\log n}{n} \right)^{1/d} + o(n^{(d-1)/d}),$$

as desired. □

5. 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_{MST}^d , β_{TF}^d , $2\beta_{\text{MM}}^d$?

In connection with Theorem 1.4:

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

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 .

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, $\beta_{\text{TSP}}^d - \beta_{\text{HK}}^d$, or optimize our techniques for this purpose, though this project could be pursued.

Acknowledgement: The proof of Theorem 1.4 was provided by a referee. It is much simpler than our original proof. We are happy to acknowledge this.

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] 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.
- [6] N. Christofides, The Traveling Salesman Problem, in *Combinatorial Optimization*, N. Christofides, P. Toth, A. Mingozzi and C. Sandi, Eds., John Wiley and Sons, New York, 1971.
- [7] F.R.K. Chung and R.L. Graham, On Steiner trees for bounded point sets, in *Geometriae Dedicata* 11 (1981) 353–361.
- [8] J. A. Costa and A. O. Hero, Geodesic entropic graphs for dimension and entropy estimation in manifold learning, in *Signal Processing, IEEE Transactions on* 52 (2004) 2210–2221.

- [9] G.B. Dantzig, D.R. Fulkerson and S.M. Johnson, Solution of a large scale traveling salesman problem, *Operations Research* 2 (1954) 393-410.
- [10] L. Few, The shortest path and the shortest road through n points, *Mathematika* 2 (1955) 141-144.
- [11] A.M. Frieze, On random symmetric travelling salesman problems, *Mathematics of Operations Research* 29 (2004) 878-890.
- [12] E. N. Gilbert, Random minimal tress, *Journal of the Society for Industrial and Applied Mathematics* 13 (1965) 376-387.
- [13] M.X. Goemans and D.J. Bertsimas, Probabilistic analysis of the Held and Karp lower bound for the Euclidean traveling salesman problem, *Mathematics of Operations Research* 16 (1991) 72-89.
- [14] M. Grötschel, L. Lovász and L. Schrijver, The ellipsoid method and its consequences in combinatorial optimization, *Combinatorica* 1 (1981) 169-197.
- [15] K. Helbig Hansen and J. Krarup, Improvements of the Held-Karp Algorithm for the symmetric Traveling Salesman Problem, *Mathematical Programming* 7 (1974) 87-96.
- [16] M. Held and R.M. Karp, The Traveling Salesman Problem and Minimum Spanning Trees, *Operations Research* 18 (1970) 1138-1162.
- [17] M. Held and R.M. Karp, The Traveling Salesman Problem and Minimum Spanning Trees Part II, *Mathematical Programming* 1 (1971) 6-25.
- [18] F. K. Hwang and Y. C. Yao, Comments on Bern's probabilistic results on rectilinear Steiner trees, in *Algorithmica* 5 (1990) 591-598.
- [19] D. Johnson, Talk presented at the Mathematical Programming Symposium, Tokyo, 1988.
- [20] D. Johnson, L. Mcgeoch and E. Rothberg, Asymptotic Experimental Analysis of the Held-Karp Traveling Salesman Bound, *Proceedings of the 7th Annual ACM Symposium on Discrete Algorithms* (1991) 341-350.
- [21] C. McDiarmid, On the Method of Bounded Differences, *Surveys in Combinatorics* 141 (1989) 148-188.
- [22] D. Miller and J. Pekny, Exact Solution of Large Asymmetric Traveling Salesman Problems, in *Science* 251 754-761.
- [23] C. H. Papadimitriou, The Euclidean travelling salesman problem is NP-complete, in *Theoretical Computer Science* 4 (1977) 237-244.
- [24] C. H. Papadimitriou, The probabilistic analysis of matching heuristics, in *Proc. 15th Conference Comm. Contr. Computing*. Univ. Illinois.
- [25] W.T. Rhee, On the travelling salesman problem in many dimensions, *Random Structures and Algorithms* 3 (1992) 227-233.
- [26] T. Smith and G. Thompson, A LIFO Implicit Enumeration Search Algorithm for the Symmetric Traveling Salesman problem Using Held and Karp's 1-Tree Relaxation, *Annals of Discrete Mathematics* 1 (1977) 479-493.
- [27] A. Srivastav and S. Werth, Probabilistic analysis of the degree bounded minimum spanning tree problem, in *FSTTCS 2007: Foundations of Software Technology and Theoretical Computer Science* 497-507.
- [28] J. Michael Steele, Subadditive Euclidean functionals and nonlinear growth in geometric probability, *The Annals of Probability* 9 (1981) 365-376.
- [29] J. M. Steele, T. L. Snyder. Worst-case growth rates of some classical problems of combinatorial optimization, in *SIAM Journal on Computing* 18 (1989) 278-287.
- [30] J. Michael Steele, Probability Theory and Combinatorial Optimization, SIAM CBMS series, 1996.
- [31] 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.
- [32] S. Steinerberger, New Bounds for the Traveling Salesman Constant, *Advances in Applied Probability* 47 (2013)
- [33] L. Tóth, Über einen geometrischen Satz, *Mathematische Zeitschrift* 46 (in German) (1940) 83-85.
- [34] T. Volgenant and R. Jonker, A Branch and Bound Algorithm for the Symmetric Traveling Salesman Problem based on the 1-tree Relaxation, *European Journal of Operations Research* 9 (1982) 83-89.

APPENDIX A. SUBADDITIVE EUCLIDEAN FUNCTIONALS

In this section, we establish that the functionals $\text{TF}_g(X)$ and $\text{HF}(X)$ have asymptotic formulas of the form $\sim \beta n^{\frac{d-1}{d}}$, as proofs for these two cases seem to be absent from the literature.

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

$$(22) \quad L(X \cup Y) \geq L(X) - o(n^{\frac{d-1}{d}}) \quad \text{for } n = |X|.$$

It has finite variance if, fixing n , we have

$$(23) \quad \mathbf{Var}(L(\mathcal{X}_n)) < \infty$$

(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 A.1 (Steele [28]). *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.*

$$(24) \quad L(\mathcal{X}_n)/n^{\frac{d-1}{d}} \rightarrow \beta_L^d \quad \text{a.s.}$$

In particular, we have a.s. that either $L(\mathcal{X}_n) = o(n^{\frac{d-1}{d}})$ (if $\beta_L^d = 0$) or that

$$(25) \quad L(\mathcal{X}_n) \sim \beta_L^d n^{\frac{d-1}{d}}.$$

This can thus be used to easily give the existence of the simple asymptotic formulas for many natural structures for Euclidean points. In the present section, we will carry out the argument for the functionals $\text{TF}_g(X)$ and $\text{HF}(X)$. We note that the constant β_L^d cannot be zero in either of these cases and thus that (24) implies (25), since for any $1 \leq i \leq n$,

$$\mathbf{E}(\min_{\substack{x \in \mathcal{X}_n \\ x \neq x_i}} \text{dist}(x_i, x)) = \Omega(n^{-\frac{1}{d}}),$$

giving a lower bound of $\Omega(n^{\frac{d-1}{d}})$ for each functional by Linearity of Expectation.

Proposition A.2. *$\text{TF}_g(X)$ and $\text{HF}(X)$ are subadditive Euclidean functionals.*

Before writing a proof, we note that for the definition of the 2-factor functionals $\text{TF}_g(X)$, 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 $\text{HF}(X)$ 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 [33] and by Few [10], 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 $\text{TF}_g(X)$, and $\text{HF}(X)$ is now a consequence of the fact that a union of disjoint 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.

□

DEPARTMENT OF MATHEMATICAL SCIENCES, CARNEGIE MELLON UNIVERSITY, PITTSBURGH, PA 15213, U.S.A.

Email address, Alan Frieze: `alan@random.math.cmu.edu`

Email address, Wesley Pegden: `wes@math.cmu.edu`