Improved Bounds for the Traveling Salesman Problem with Neighborhoods on Uniform Disks

Ioana Oriana Bercea*

Abstract

Given a set of n disks of radius R in the Euclidean plane, the Traveling Salesman Problem With Neighborhoods (TSPN) on uniform disks asks for the shortest tour that visits all of the disks. The problem is a generalization of the classical Traveling Salesman Problem(TSP) on points and has been widely studied in the literature. For the case of disjoint uniform disks of radius R, Dumitrescu and Mitchell [14] give a PTAS and also show that the optimal TSP tour on the centers of the disks is a 3.547-approximation to the TSPN version. The core of the latter analysis is based on bounding the detour that the optimal TSPN tour has to make in order to visit the centers of each disk and shows that it is at most $2Rn$ in the worst case. Häme, Hyytiä and Hakula [21] asked whether this bound is tight when R is small and conjectured that it is at most $\sqrt{3}Rn$.

We further investigate this question and derive structural properties of the optimal TSPN tour to describe the cases in which the bound is smaller than $2Rn$. Specifically, we show that if the optimal TSPN tour is not a straight line, at least one of the following is guaranteed to be true: the bound is smaller than $1.999Rn$ or the TSP on the centers is a 2-approximation. The latter bound of 2 is the best that we can get in general. Our framework is based on using the optimality of the TSPN tour to identify local structures for which the detour is large and then using their geometry to derive better lower bounds on the length of the TSPN tour. This leads to an improved approximation factor of 3.53 for disjoint uniform disks and 6.728 for the general case. We further show that the Häme, Hyytiä and Hakula conjecture is true for the case of three disks and discuss the method used to obtain it.

1 Introduction

We study the Traveling Salesman Problem with Neighborhoods (TSPN) when each neighborhood is a disk of fixed radius R. The problem is a generalization of the classical Euclidean Traveling Salesman Problem (TSP), when each point to be visited is replaced with a region (interchangeably, a neighborhood) and the objective is to compute a tour of minimum length that visits at least one point from each of these regions. While it is known that Euclidean TSP admits a Polynomial Time Approximation Scheme (PTAS) due to the celebrated results of Arora [3] and Mitchell [25], Euclidean TSPN has been shown in fact to be APX-hard [29, 13] even for line segments of comparable length [17]. The geometric version of TSPN was first studied by Arkin and Hassin [2] who gave constant factor approximations for a variety of cases. Since then, there has been a wide ranging study of TSPN for different types of regions. In the case of connected regions, there is a series of $O(\log n)$ approximations [24, 17, 20]. Better approximations are known for cases that consider various restrictions on the regions such as comparable sizes (i.e. diameter), fatness (ratio between the smallest circumscribing radius and largest inscribing radius, or how well can a disk approximate the region) and pairwise disjointness or limited intersection [14, 13, 17, 27, 8, 15, 6, 26, 9]. We refer the reader to [19] for a comprehensive survey of the results.

We study the disk version which models the situation in which each customer is willing to travel a distance R to meet the salesperson. This is considered an important special case of the general TSPN [15] and is especially relevant since it has found applicability in other areas such as path planning algorithms for coverage with a circular field of view [1, 18] and most recently for data collection in wireless sensor networks [33, 23, 12]. Various heuristics [33, 7, 10, 22] and variations [4, 23, 30] have been considered, all of which have as their basis the TSPN on uniform disks problem.

Dumitrescu and Mitchell [14] were the first to specifically address the case of uniform disks in 2001. They showed a PTAS for disjoint unit disks and simpler constant factor approximations for the disjoint and overlapping cases. The specific factor of 3.547 for disjoint disks is relative to using a routine for TSP on points (i.e. the actual constant depends on the subroutine used). The PTAS and the 3.547-approximation are the best known factors for the disjoint case. Later, Dumitrescu and Tóth [15] improved the constant factor in [14] for the overlapping case from 7.62 to 6.75 and extended it to unit balls in $\mathbb{R}^d$, giving a $O(7.73^d)$-approximation. When the balls are disjoint, Elbassioni et al. [17] showed a $O(2^d/\sqrt{d})$-approximation. Most recently, Dumitrescu and Tóth [16] gave a randomized constant factor for (po-

*Department of Computer Science, University of Maryland, College Park, ioana@cs.umd.edu

tentially overlapping) disks of arbitrary radii (the actual constant is not mentioned and seems large). As noted by the authors in [15], while the complexity of the disk case is well understood generally, the question of obtaining practical and better constant factor approximations remains of high interest.

In this paper, we aim for an improved constant factor algorithm for the case of uniform radius disks and note that the algorithm proposed by Dumitrescu and Mitchell [14] for the disjoint case outputs an approximate TSP tour on the *centers* of each disk. The core of their analysis is a bound that compares the length of the optimal TSP on the centers of each disk($|TSP^*|$) with the length of the optimal TSPN on the disks ($|TSPN^*|$) and says that

$$|TSP^*| \leq |TSPN^*| + 2Rn, \qquad (1)$$

where $n \geq 2$ is the number of disks in the instance. In addition, the authors use a packing argument to lower bound the length of the optimal TSPN tour in terms of R and n and get that $\frac{\pi}{4}Rn - \pi R \leq |TSPN^*|$. Overall, this gives a 3.547- approximation and in addition, the authors show that the algorithm cannot give better than a factor 2 approximation. While other methods for choosing representative points can be employed [16, 2, 14], this approach is appealing both in its elegance and because it does not depend on R. Moreover, other existing constant factors approximations for TSPN often hide large constants [13, 17, 16] that are incurred as a consequence of using general bounds on the length of the optimal tour that do not directly exploit the structure of the regions or of the optimal TSPN tour (bounding rectangle argument and Combination Lemma in [2]). In order to improve on them, the challenge then becomes to develop bounds that exploit the difference in behavior between a TSP tour (on points) and the TSPN tour on the regions and furthermore, avoid using general purpose techniques that add on to the overall approximation factor.

In this context, one way to improve the approximation factor for disjoint disks is to better understand the relationship between the optimal TSP tour on the centers and the optimal TSPN on the disks. Specifically, is the $2Rn$ term in (1) tight or can it be improved by using specific structural properties of the optimal TSPN? A similar question was asked in 2011 by Häme, Hyytiä and Hakula [21] for the case when R is very small (and hence, TSP^* and $TSPN^*$ respect the same order and the disks are pairwise disjoint). They conjectured that the true detour term should be $\sqrt{3}nR$ and constructed arbitrarily large instances of disjoint disks that converge to this case. We refer to this as the **Häme, Hyytiä and Hakula conjecture.** Subsequent experiments by Müller [28], however, suggest that this might be true only for tours up to five disks and higher otherwise. No further progress has been made towards the conjecture since then.

Contributions. We make the first progress on the conjecture and develop a twofold method that either improves the bound in (1) or shows that the TSP on the centers is a good approximation for the TSPN on the disks. Formally, we get that

Theorem 1 *For any $n \geq 4$ disjoint disks of radius R at least one of the following is true:*

- *$TSPN^*$ is supported by a straight line,*
- *$|TSP^*| \leq |TSPN^*| + 1.999Rn$,*
- *$|TSP^*| \leq 2 \cdot |TSPN^*|$.*

Our framework also gives an overall 3.53*-approximation for the case of uniform disjoint disks and a* 6.728*-approximation for the overlapping case.*

While the improvement in the overall approximation factor is small, our framework strives to provide new insight into the problem that can be explored further. Specifically, the 2-approximation (optimal with respect to the method of computing a TSP on the centers [14]) comes from the case in which the TSPN tour takes a lot of sharp turns. Furthermore, it is based on a lower bound that does not rely on packing arguments. To the best of our knowledge, this is the first such bound specifically for TSPN out of all arguments for general fat regions [27]. As such, it might be of independent interest and it could, for example, lead to improved approximation factors for balls in $\mathbb{R}^d$ that do not depend exponentially on the dimension. Moreover, the fatness of the disks is used in showing that short sharp turns lead to a disk being visited multiple times and can conceivably be used to show similar properties for other fat regions.

We start by fixing an order σ and comparing the TSPN tour that visits the disks in that order to the TSP tour that visits their respective centers in the same order. The $2Rn$ term in (1) comes from considering the points at which the TSPN touches the boundary of each disk and charging each such vertex with a $2R$ detour for going to its respective center and coming back. In this view, the $2Rn$ term cannot be improved since the charge on each vertex will always be $2R$. Instead, we reinterpret the bound as charging the *edges* of the TSPN tour instead of its vertices and notice that the charge for each edge can now be anywhere between $-2R$ and $2R$, depending on how close the tour is to (locally) visiting pairs of disks optimally. In this context, we define a "bad" edge to be one that incurs a large charge (i.e. $> (2-\epsilon)R$ for some $\epsilon > 0$). We show that such bad edges lead to the TSPN tour exhibiting sharp turns (i.e. with small interior angle). When the edges of the sharp turn are long, we use that to derive a better lower bound on the overall TSPN tour length. On the other hand, when one of them is short, we show that the tour must then visit a disk twice (i.e. visit it once, then touch another disk and return back to it). The crux of the argument is in understanding how these short sharp turns that visit a disk multiple times influence the global detour term.

When a tour visits a disk more than once, two scenarios follow naturally from the classical TSP case of just visiting points: either the order σ is not optimal or the tour must follow a straight line. Surprisingly, we show that a third alternative scenario is also possible, whose local structure we call a β-triad. The main technical contribution of the paper is in describing structural properties of such β-triads and showing that they actually have a low *average* detour. Specifically, we construct an additional order σ' and use an averaging argument to show that β-triads have low detour when compared to the TSP tours that visit the centers in the order σ and σ'. This then allows us to conclude that they have a low detour with respect to the optimal TSP on the centers.

Along the way, we also show that the Häme, Hyytiä and Hakula conjecture is true for $n = 3$ and use it to bound the average detour of β-triads. We include a discussion of the method used to derive it, involving Fermat-Weber points, which might be useful for the case of $n \geq 4$. We also discuss how our approach can be used within the framework of Dumitrescu and Tóth [15] to yield improved approximation factors for the overlapping disks case.

Preliminaries. We consider $n \geq 3$ disjoint disks of radius R in the Euclidean plane. We denote an optimal TSP tour on the centers of the disks as TSP^*. Similarly, $TSPN^*$ will denote an optimal TSPN tour on the disks. Our results will be with respect to a fixed TSPN tour (which we call simply $TSPN$) described by a sequence of ordered points $P_1, P_2, \ldots, P_n$ on the boundary of the disks such that the tour is a polygonal cycle with edges (P_i, P_{i+1}). Furthermore, we have that for each of the input disks, there exists some $i \in [1, n]$ such that point P_i is on the boundary of the disk.

Notice that the points P_i induce a natural order σ on the disks with centers $O_1, O_2, \ldots O_n$, i.e. σ corresponds to the identity permutation on $P_1, \ldots, P_n$. For the majority of our theorems, we will assume that TSP always refers to a tour *on the centers* and in the order σ on the disks. When we need to make a difference, we will further use $TSP(\sigma')$ to be the tour which visits the centers in the order given by the permutation σ'. Given two such permutations σ and σ', we say that $\sigma \cap \sigma'$ refers to the maximal set of points on which σ and σ' agree. In this context, $TSPN(\sigma \cap \sigma')$ refers to the collection of paths we get from visiting the points P_i according to $\sigma \cap \sigma'$. Similarly, $TSP(\sigma \cap \sigma')$ corresponds to the collection of paths that we get from visiting the points O_i according to $\sigma \cap \sigma'$.

Finally, we denote the length of a tour $\mathcal{T}$ as $|\mathcal{T}|$. When $\mathcal{T}$ is a collection of paths, we have that $|\mathcal{T}|$ represents the total length of each of the paths. When A and B are points, we have that $|AB|$ denotes the length of the segment AB. We therefore have that $|TSPN| = \sum_{i=1}^{n} |P_i P_{i+1}|$ and $|TSP| = \sum_{i=1}^{n} |O_i O_{i+1}|$, where $P_{n+1} = P_1$ and $O_{n+1} = O_1$.

2 β-Triads and a Structural Theorem

Before we formally define what a "bad" edge is, we will describe how to interpret the $2Rn$ *detour bound* from [14] as charging edges instead of vertices. We fix an order σ and consider the points P_i and O_i as previously defined. The argument in [14] then says that we must have:

$$\textstyle\sum_i |O_i O_{i+1}| \leq \sum_i |P_i P_{i+1}| + 2Rn.$$

In this context, the term $\sum_i P_i P_{i+1} + 2Rn$ is the length of a tour that follows the TSPN tour and additionally, at each point P_i, takes a detour of $2R$ to visit the center O_i and come back. Choosing σ to be the optimal order in which $TSPN^*$ visits the disks gives us (1). In this view, the detour term $2Rn$ is obtained by charging $2R$ to each point P_i of the TSPN tour. Instead, we can also think of it as coming from charging each *edge* $P_i P_{i+1}$ of the tour with a *local detour* of $2R$ in the following sense:

$$|O_i O_{i+1}| \leq |P_i P_{i+1}| + 2R.$$

This new perspective is quite natural since it captures the observation that the shortest edge which visits the disks centered at O_i and O_{i+1} has length exactly $|O_i O_{i+1}| - 2R$ and hence the $TSPN$ tour has to pay at least that for each pair of consecutive disks it visits. In this sense, we decompose the global detour term of $2Rn$ into n local detour terms $|O_i O_{i+1}| - |P_i P_{i+1}|$ that essentially quantify how efficient the TSPN on the disks is locally.

In this context, saying a TSPN edge has a high local detour is equivalent to saying that it is close to being locally optimal or shortest possible: when the edge is exactly of length $|O_i O_{i+1}| - 2R$, its local detour is $2R$ (the maximum). If, on the other hand, we know that the edge is bounded away from $|O_i O_{i+1}| - 2R$, i.e. $|P_i P_{i+1}| > |O_i O_{i+1}| - 2R + \epsilon R$, for some $\epsilon > 0$, this translates into a local detour of at most $(2 - \epsilon)R$. Intuitively, such an edge is "good" for us because it allows us to lower the overall detour term. In contrast, "bad" $P_i P_{i+1}$ edges are the ones for which the local detour term is large and consequently, their length is closer to

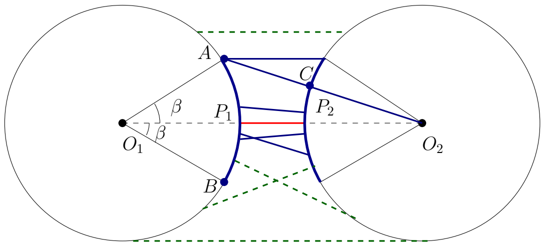

Figure 1: Bad edges are guaranteed to have both endpoints in the blue arcs. Furthermore, if $|P_1P_2| \leq |AC|$, then P_1P_2 is bad. In contrast, the dashed edges are guaranteed to be good edges.

$|O_iO_{i+1}| - 2R$. Our technique is motivated by trying to describe the behavior of such bad edges.

Formally, we consider a fixed angle parameter $\beta \in [0, \pi/12]$ that we instantiate later when we derive the overall bounds. We define the function:

$$f(O_1O_2, \beta) = \sqrt{|O_1O_2|^2 + R^2 - 2R|O_1O_2|\cos\beta},$$

which is $|O_1O_2| - R$ when $\beta = 0$ and $|O_1O_2| + R$ when $\beta = \pi$. Intuitively, the quantity $f(O_1O_2, \beta) - R$ will control how close we are to $|O_1O_2| - 2R$. We then say that the edge P_1P_2 is **bad** if $|P_1P_2| \leq f(O_1O_2, \beta) - R$ and **good** otherwise (we abstract away the dependency on β for simplicity). Bad edges are close to $|O_1O_2| - 2R$ and will incur a large local detour. In contrast, using straightforward algebra, one can show that a good edge P_1P_2 is guaranteed to have a small detour: $|O_1O_2| \leq |P_1P_2| + (1 + \cos\beta)R$.

Consecutive bad edges. The idea behind defining bad edges in terms of $f(O_1O_2, \beta) - R$ is that it allows us to restrict the location of P_1 and P_2 on the boundary of their respective disks as seen in Figure 1. Specifically, there are exactly two points A and B on the boundary of the first disk with the property that the shortest distance from A or B to the boundary of the second disk is exactly $f(O_1O_2, \beta) - R$. Not coincidentally, they form an angle of β with $O_1O_2 : \angle AO_1O_2 = \angle BO_1O_2 = \beta$. In general, P_1 (and in a similar fashion P_2) is guaranteed to lie in the short arc between A and B whenever P_1P_2 is upper bounded by $f(O_1O_2, \beta) - R$ (Lemma 5).

When a second bad edge P_2P_3 is considered, we can conclude that the angle $O_1O_2O_3$ has to be at most 2β and hence the TSP on the centers must make a sharp turn after it visits O_2 (Corollary 6). If that happens and the disks are close to each other, we have that one of the edges of the TSP must actually intersect a disk twice. Specifically, if $|O_1O_2| \leq R/\sin(2\beta)$, then the support line for O_2O_3 must pass through the disk centered at O_1. Theorem 7 shows that if this happens, then the corresponding TSPN edge P_2P_3 must also cross this disk (Appendix A).

The fact that the disk centered at O_1 is crossed by both P_1P_2 and P_2P_3 suggests that the TSPN might not be optimal because it could be shortcut. Our structural theorem identifies when that is the case and isolates the remainder as having a specialized local structure which we call a β-triad. Formally, we say that a specific TSPN subpath $P_n - P_1 - P_2 - P_3$ is a **β-triad** if it satisfies all of the following properties (Figure 2):

- P_1P_2 and P_2P_3 are bad edges and $O_1O_2 \leq R/\sin(2\beta)$,
- P_1, P_2, P_3 are not collinear but P_n, P_1, P_2 are collinear with P_1 between P_n and P_2.

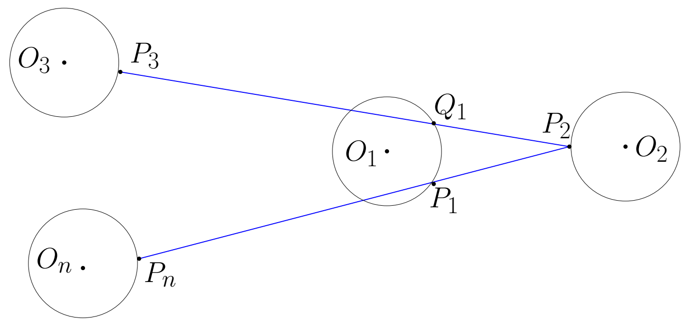

Figure 2: The path $P_n - P_1 - P_2 - P_3$ forms a β-triad.

We state the structural theorem here and refer the reader to Appendix A for a complete argument. The case in which the TSPN tour follows a straight line that stabs all the disks is discussed separately in Appendix C and is of separate interest.

Theorem 2 *For $n \geq 4$, if P_1P_2 and P_2P_3 are bad edges and $O_1O_2 \leq R/\sin(2\beta)$ then at least one of the following is true: the TSPN tour is not optimal, the TSPN tour is supported by a straight line or the path $P_n - P_1 - P_2 - P_3$ forms a β-triad.*

β-triads. The case of β-triads is interesting because it arises naturally as a consequence of dealing with regions instead of points. Given any optimal tour that exhibits internal angles $\leq \pi/6$, we can always add an extra disk at each sharp turn that will maintain optimality, pairwise disjointness and be intersected twice by this tour, giving rise to a β-triad. It is therefore important that we understand their behaviour.

Because of the fact that P_1P_2 and P_2P_3 are bad edges, the β-triad is likely to have a high detour with respect to $TSP(\sigma)$. Nevertheless, we show that there exists an alternate ordering σ' such that the average detour of the three edges in the β-triad with respect to $(|TSP(\sigma)| + |TSP(\sigma')|)/2$ is $3\sqrt{3}R$. The order σ' takes advantage of the fact that the disk centered at O_1 is crossed twice and inverts the order in which it is visited without changing the cost of the underlying TSPN tour. The $3\sqrt{3}R$ bound comes from proving the Häme, Hyytiä and Hakula conjecture for $n = 3$ (Appendix F). In order to be able to construct σ' consistently across multiple β-triads, we also show that β-triads are isolated events and specifically that they are edge-disjoint (Lemma 8 in Appendix B).

Theorem 3 *If the TSPN in the order σ has k β-triads that together cover a set of edges of total length L_T, then we can construct another order σ' that agrees with σ on everything except the order inside the β-triads such that:*

$$\frac{|TSP(\sigma)| + |TSP(\sigma')|}{2} \leq |TSP(\sigma \cap \sigma')| + L_T + 3\sqrt{3}Rk.$$

Proof. We discuss the case for $k = 1$ and show how to modify the argument for $k > 1$. Suppose that $P_n - P_1 - P_2 - P_3$ form a β-triad. We know that $TSP(\sigma)$ visits the centers of each disk in the order $O_1, O_2, O_3, \ldots O_n$. We consider an additional order σ' such that $TSP(\sigma')$ visits the centers in the order $O_2, O_1, O_3, \ldots O_n$. Notice that both σ ad σ' agree on the order $O_3, O_4, \ldots, O_n$ and that they differ in the fact that σ visits O_1 after O_n and before O_2 and σ' visits O_1 after O_2 and before O_3. Therefore, for $T' = TSP(\sigma \cap \sigma')$, we have that $|T'| = |O_3O_4| + \ldots + |O_{n-1}O_n|$, $|TSP(\sigma)| = |T'| + |O_nO_1| + |O_1O_2| + |O_2O_3|$ and $|TSP(\sigma')| = |T'| + |O_nO_2| + |O_2O_1| + |O_1O_3|$.

On the other hand, the length of the TSPN with respect to the orders σ and σ' stays the same. The local cost of visiting $P_n - P_1 - P_2 - P_3$ is $L_T = |P_nP_1| + |P_1P_2| + |P_2P_3| = |P_nP_2| + |P_2P_3|$, since P_n, P_1 and P_2 are collinear and P_1 is between P_n and P_2. We also know that P_2P_3 intersects the disk centered at O_1 at some point Q_1 that is different from P_1 (Theorem 7). In other words, the TSPN that visits the points $P_n - P_1 - P_2 - P_3$ can be reimagined as visiting the points $P_n - P_2 - Q_1 - P_3$ and therefore respecting the order σ'. The local cost of crossing these edges is the same as before: $|P_nP_2| + |P_2Q_1| + |Q_1P_3| = |P_nP_2| + |P_2P_3| = L_T$.

We now apply Theorem 10 (the $3\sqrt{3}R$ bound for $n = 3$) on the TSP tour $O_n - O_1 - O_2$ with the TSPN tour $P_n - P_1 - P_2$ and get that:

$$\begin{aligned} |O_nO_1| + |O_1O_2| + |O_nO_2| &\leq |P_nP_1| + |P_1P_2| + |P_nP_2| + \\ &+ 3\sqrt{3}R \\ &\leq 2|P_2P_n| + 3\sqrt{3}R. \end{aligned}$$

On the other hand, if we consider the tour $O_1 - O_2 - O_3$ with the TSPN tour $P_2 - Q_1 - P_3$, we get that:

$$\begin{aligned} |O_1O_2| + |O_2O_3| + |O_1O_3| &\leq |Q_1P_2| + |P_2P_3| + |P_3Q_1| + \\ &+ 3\sqrt{3}R \\ &\leq 2|P_2P_3| + 3\sqrt{3}R. \end{aligned}$$

Combining the two inequalities and rearranging some terms gives us that:

$$\begin{aligned} |TSP(\sigma)| + |TSP(\sigma')| =& 2|T'| + |O_nO_1| + |O_1O_2| + |O_2O_3| + \\ &+ |O_nO_2| + |O_2O_1| + |O_1O_3| \\ =& 2|T'| + |O_nO_1| + |O_1O_2| + |O_nO_2| + \\ &+ |O_1O_2| + |O_1O_3| + |O_2O_3| \\ \leq& 2|T'| + 2|P_2P_n| + 2|P_2P_3| + 6\sqrt{3}R. \end{aligned}$$

Since $L_T = |P_2P_n| + |P_2P_3|$, we get our conclusion.

When $k > 1$, we construct the order σ' by switching the order in which we visit the centers in each β-triad in the same way as before. Since all the β-triads are edge disjoint (Lemma 8), we can construct σ' without any conflicts because any reordering that happens in one β-triad will not affect another β-triad. □

3 Improved Bounds on the TSPN Tour

Our main strategy will be a careful balancing of good and bad edges, in which the detour of good edges will be upper bounded by $(1+\cos\beta)R$ and that of bad edges by $2R$. While the bad edges will have the highest detour possible, we will use the fact that they must also be large in order to lower bound the TSPN tour more efficiently than Lemma 4 from [14] and [15], which we quote here for completeness.

Lemma 4 *[14, 15] For n disjoint disks of radius R, we have that any TSPN tour $\mathcal{T}$ on them satisfies:*

$$\tfrac{\pi}{4}Rn - \pi R \leq |\mathcal{T}|.$$

We will now show the proof of Theorem 1.

Proof. Assume the $TSPN^*$ is not a straight line. We start by singling out the β-triads and considering the two orderings σ and σ' from Theorem 3. If there are k_1 β-triads $\mathcal{T}_1, \ldots, \mathcal{T}_{k_1}$ spanning edges of total length L_T, we get that:

$$\begin{aligned} |TSP^*| &\leq \frac{|TSP(\sigma)| + |TSP(\sigma')|}{2} \\ &\leq |TSP(\sigma \cap \sigma')| + L_T + 3\sqrt{3}R \cdot k_1. \end{aligned}$$

Observe that $TSPN(\sigma \cap \sigma')$ is a collection of disjoint paths. From all of these paths, we further extract each from these a total of k_2 subpaths $\mathcal{G}_1, \ldots, \mathcal{G}_{k_2}$ consisting of good edges. Notice that the remaining subpaths left in $\sigma \cap \sigma'$ consist of bad edges which do not form a β-triad. Suppose we obtain l such remaining subpaths paths $\mathcal{B}_1, \ldots, \mathcal{B}_l$. In other words, we have decomposed the TSPN into three categories of subpaths:

- k_1 β-triads $\mathcal{T}_1, \ldots, \mathcal{T}_{k_1}$,
- k_2 paths $\mathcal{G}_1, \ldots, \mathcal{G}_{k_2}$ that cover the remaining good edges, and

- l paths $\mathcal{B}_1, \ldots, \mathcal{B}_l$ that consist only of bad edges which do not form β-triads.

We are now ready to evaluate the detour that each of these paths takes. For each $i \in [1, k_2]$ let ψ_i the natural order on the disks associated with $\mathcal{G}_i$ and let n_i be the number of edges in $\mathcal{G}_i$. We have that:

$$|TSP(\psi_i)| \le |TSPN(\psi_i)| + (1 + \cos\beta)R \cdot n_i.$$

When it comes to the paths $\mathcal{B}_j$, with $j \in [1, l]$, let σ_j be their natural associated orders and let m_j be the number of edges it contains. We have that $|TSP(\sigma_j)| \le |TSPN(\sigma_j)| + 2R \cdot m_j$.

Let $N = \sum_{i=1}^{k_2} n_i$ be the total number of edges in $\mathcal{G}_1, \ldots, \mathcal{G}_{k_2}$ and $M = \sum_{j=1}^{l} m_j$ the total number of edges in $\mathcal{B}_1, \ldots, \mathcal{B}_l$. By construction, we decomposed $TSPN(\sigma \cap \sigma')$ into these two groups of edge disjoint paths and we therefore get that:

$$\begin{aligned}
|TSP(\sigma \cap \sigma')| &= \sum_{i=1}^{k_2} |TSP(\psi_i)| + \sum_{j=1}^{l} |TSP(\sigma_j)| \\
&\le \sum_{i=1}^{k_2} \Big(|TSPN(\psi_i)| + (1 + \cos\beta)R \cdot n_i\Big) \\
&+ \sum_{j=1}^{l} \Big(|TSPN(\sigma_j)| + 2R \cdot m_j\Big) \\
&\le |TSPN(\sigma \cap \sigma')| + (1 + \cos\beta)RN + 2RM.
\end{aligned}$$

Including the β-triads back into our bound, we get that:

$$\begin{aligned}
|TSP^*| &\le \frac{|TSP(\sigma)| + |TSP(\sigma')|}{2} \\
&\le |TSP(\sigma \cap \sigma')| + L_T + 3\sqrt{3}R \cdot k_1 \\
&\le |TSPN(\sigma \cap \sigma')| + L_T + 3\sqrt{3}R \cdot k_1 + \\
&+ (1 + \cos\beta)RN + 2RM \\
&\le |TSPN| + 3\sqrt{3}R \cdot k_1 + (1 + \cos\beta)RN + 2RM.
\end{aligned}$$

In other words, we've expressed the total detour of the $TSPN$ according to edges that participate in β-triads, edges in $\mathcal{G}_1, \ldots, \mathcal{G}_{k_2}$ and edges in $\mathcal{B}_1, \ldots, \mathcal{B}_l$. By construction, none of these paths share edges and so $3k_1 + N + M = n$. Let $K = 3k_1 + N$ be the total number of edges either in a β-triad or in $\mathcal{G}_1, \ldots, \mathcal{G}_{k_2}$ and since $\sqrt{3} \le 1 + \cos\beta$, we have that:

$$|TSP^*| \le |TSPN| + (1 + \cos\beta)R \cdot K + 2R \cdot (n - K).$$

Case 1: when $K \ge \frac{n}{2}$. In this situation, we have that:

$$|TSP^*| \le |TSPN^*| + \frac{3 + \cos\beta}{2} \cdot R \cdot n.$$

The average detour per edge $\frac{3+\cos\beta}{2}$ is better than the $2R$ bound, but it is constrained by the choice of $\beta \in [0, \pi/12]$, which means that the best we could hope for is an average detour of $\frac{1}{2}(3 + \cos\frac{\pi}{12})R < 1.983R$. We note that the average detour in the Häme, Hyytiä and Hakula conjecture is $\sqrt{3}R \approx 1.732R$. Using Lemma 4 gives us that

$$\begin{aligned}
|TSP^*| \le &\left(1 + \frac{2}{\pi} \cdot (3 + \cos\beta)\right) \cdot |TSPN^*| + \\
&+ 2 \cdot (3 + \cos\beta)R \qquad .
\end{aligned}$$

For large n, the $1 + \frac{2}{\pi} \cdot (3 + \cos\beta)$ term will dominate our approximation factor and is at most 3.525 when $\beta = \pi/12$.

Case 2: when $K < \frac{n}{2}$. In this situation, even the overall detour might be large, we will show that in fact, in this case, TSP^* is a 2-approximation and therefore, the best that it can be in general. We know that each path $\mathcal{B}_j$ consists of bad edges which do not form any β-triads. In other words, if P_1P_2 is an edge in it, then we know that $|O_1O2| > R/\sin(2\beta)$ which in turn means that $|P_1P_2| > (1/\sin(2\beta) - 2) \cdot R$. Overall we have that:

$$\begin{aligned}
|TSPN| &\ge \sum_{j=1}^{l} |\mathcal{B}_j| \ge \Big(\frac{1}{\sin(2\beta)} - 2\Big)R \cdot (n - K) \\
&\ge \Big(\frac{1}{2\sin(2\beta)} - 1\Big)R \cdot n.
\end{aligned}$$

Since the total detour could be at most $2R$ per edge, we get that:

$$|TSP^*| \le \left(1 + \frac{2}{\frac{1}{2\sin(2\beta)} - 1}\right) \cdot |TSPN|.$$

When $\beta = \frac{1}{2}\arcsin\frac{1}{6}$, the detour from Case 1 becomes $\frac{3+\cos\beta}{2} \approx 1.998$ and the approximation factor from Case 2 becomes exactly 2. We note that the machinery described can be used to obtain more nuanced results. In particular, lower choices for β will drive the approximation factor in Case 2 even lower than 2, at the expense of a higher detour bound for Case 1.

□

In Appendix D, we perform a similar analysis for a more general threshold of $\frac{N}{\alpha}$ for some $\alpha > 1$ (rather than $\frac{N}{2}$). We give there the specific values of α and β that give us a factor 2.53-approximation. The case for overlapping disks follows from the same analysis when we use the algorithm of Dumitrescu and Tóth [15]. We include the details for it in Appendix E.

Acknowledgements. The author would like to thank Prof. Samir Khuller for suggesting the problem and for inspiring conversations on the topic, as well as the reviewers for helpful comments on the paper.

References

[1] E. M. Arkin, S. P. Fekete, and J. S. Mitchell. Approximation algorithms for lawn mowing and milling. *Computational Geometry*, 17(1-2):25–50, 2000.

[2] E. M. Arkin and R. Hassin. Approximation algorithms for the geometric covering salesman problem. *Discrete Applied Mathematics*, 55(3):197–218, 1994.

[3] S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM (JACM)*, 45(5):753–782, 1998.

[4] D. Bhadauria, O. Tekdas, and V. Isler. Robotic data mules for collecting data over sparse sensor fields. *Journal of Field Robotics*, 28(3):388–404, 2011.

[5] B. Bhattacharya, J. Czyzowicz, P. Egyed, G. Toussaint, I. Stojmenovic, and J. Urrutia. Computing shortest transversals of sets. *International Journal of Computational Geometry & Applications*, 2(04):417–435, 1992.

[6] H. L. Bodlaender, C. Feremans, A. Grigoriev, E. Penninkx, R. Sitters, and T. Wolle. On the minimum corridor connection problem and other generalized geometric problems. *Computational Geometry*, 42(9):939–951, 2009.

[7] F. Carrabs, C. Cerrone, R. Cerulli, and M. Gaudioso. A novel discretization scheme for the close enough traveling salesman problem. *Computers & Operations Research*, 78:163–171, 2017.

[8] T.-H. H. Chan and K. Elbassioni. A QPTAS for TSP with fat weakly disjoint neighborhoods in doubling metrics. *Discrete & Computational Geometry*, 46(4):704–723, 2011.

[9] T.-H. H. Chan and S. H.-C. Jiang. Reducing curse of dimensionality: Improved PTAS for TSP (with neighborhoods) in doubling metrics. *ACM Transactions on Algorithms (TALG)*, 14(1):9, 2018.

[10] W.-L. Chang, D. Zeng, R.-C. Chen, and S. Guo. An artificial bee colony algorithm for data collection path planning in sparse wireless sensor networks. *International Journal of Machine Learning and Cybernetics*, 6(3):375–383, 2015.

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

[12] G. Citovsky, J. Gao, J. S. Mitchell, and J. Zeng. Exact and approximation algorithms for data mule scheduling in a sensor network. In *International Symposium on Algorithms and Experiments for Sensor Systems, Wireless Networks and Distributed Robotics*, pages 57–70. Springer, 2015.

[13] M. de Berg, J. Gudmundsson, M. J. Katz, C. Levcopoulos, M. H. Overmars, and A. F. van der Stappen. Tsp with neighborhoods of varying size. *Journal of Algorithms*, 57(1):22–36, 2005.

[14] A. Dumitrescu and J. S. Mitchell. Approximation algorithms for TSP with neighborhoods in the plane. *Journal of Algorithms*, 48(1):135–159, 2003.

[15] A. Dumitrescu and C. D. Tóth. The traveling salesman problem for lines, balls, and planes. *ACM Transactions on Algorithms (TALG)*, 12(3):43, 2016.

[16] A. Dumitrescu and C. D. Tóth. Constant-factor approximation for TSP with disks. In *A Journey Through Discrete Mathematics*, pages 375–390. Springer, 2017.

[17] K. Elbassioni, A. V. Fishkin, and R. Sitters. Approximation algorithms for the Euclidean traveling salesman problem with discrete and continuous neighborhoods. *International Journal of Computational Geometry & Applications*, 19(02):173–193, 2009.

[18] E. Galceran and M. Carreras. A survey on coverage path planning for robotics. *Robotics and Autonomous systems*, 61(12):1258–1276, 2013.

[19] J. E. Goodman, J. O'Rourke, and C. D. Tóth, editors. *Handbook of Discrete and Computational Geometry, Third Edition*. CRC Press, 2017.

[20] J. Gudmundsson and C. Levcopoulos. A fast approximation algorithm for TSP with neighborhoods. *Nord. J. Comput.*, 6(4):469, 1999.

[21] L. Häme, E. Hyytiä, and H. Hakula. The Traveling Salesman Problem with Differential Neighborhoods. In *European Workshop on Computational Geometry (EuroCG)*, Morschach, Switzerland, Mar. 2011.

[22] J.-S. Liu, S.-Y. Wu, and K.-M. Chiu. Path planning of a data mule in wireless sensor network using an improved implementation of clustering-based genetic algorithm. In *Computational Intelligence in Control and Automation (CICA), 2013 IEEE Symposium on*, pages 30–37. IEEE, 2013.

[23] M. Ma, Y. Yang, and M. Zhao. Tour planning for mobile data-gathering mechanisms in wireless sensor networks. *IEEE Transactions on Vehicular Technology*, 62(4):1472–1483, 2013.

[24] C. S. Mata and J. S. Mitchell. Approximation algorithms for geometric tour and network design problems. In *Proceedings of the eleventh annual symposium on Computational geometry*, pages 360–369. ACM, 1995.

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

[26] J. S. Mitchell. A PTAS for TSP with neighborhoods among fat regions in the plane. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 11–18. Society for Industrial and Applied Mathematics, 2007.

[27] J. S. Mitchell. A constant-factor approximation algorithm for TSP with pairwise-disjoint connected neighborhoods in the plane. In *Proceedings of the twenty-sixth annual symposium on Computational geometry*, pages 183–191. ACM, 2010.

[28] C. L. Müller. Finding maximizing Euclidean TSP tours for the Häme-Hyytiä-Hakula conjecture. Technical report, Technical Report CGL-TR-13, ETHZ, 2011.

[29] S. Safra and O. Schwartz. On the complexity of approximating TSP with neighborhoods and related problems. *Computational Complexity*, 14(4):281–307, 2006.

[30] O. Tekdas, D. Bhadauria, and V. Isler. Efficient data collection from wireless nodes under the two-ring communication model. *The International Journal of Robotics Research*, 31(6):774–784, 2012.

[31] A. Weber. *Ueber den standort der industrien*, volume 2. 1909.

[32] G. O. Wesolowsky. The Weber problem: History and perspectives. *Computers & Operations Research*, 1993.

[33] B. Yuan, M. Orlowska, and S. Sadiq. On the optimal robot routing problem in wireless sensor networks. *IEEE Transactions on Knowledge and Data Engineering*, 19(9):1252–1261, 2007.

Appendix

A Proof of Theorem 2 : Introducing β-triads

We begin by getting some intuition as to where on the boundary must the TSPN hit the disks in order to visit them within a short distance (depending on the parameter β). Specifically, we show that if a segment P_1P_2 is bad, then the possible locations for P_1 and P_2 are limited to a small interval on the boundary.

Lemma 5 *If P_1P_2 is bad, then the angles $\angle O_1O_2P_2$ and $\angle O_2O_1P_1$ are $\leq \beta$.*

Proof. Let $\gamma = \angle O_1O_2P_2 \leq \pi$ and notice that $O_1P_2 = f(O_1O_2, \gamma)$. Consider the point Q where O_1P_2 intersects the first disk and note that the shortest distance from P_2 to the first disk is exactly $P_2Q = f(O_1O_2, \gamma) - R$. We therefore get that $P_1P_2 \geq P_2Q$. Now notice that, if $\gamma > \beta$, then $f(O_1O_2, \gamma) > f(O_1O_2, \beta)$ and so $P_2Q > f(O_1O_2, \beta) - R$, which would lead to a contradiction. The same argument can be applied for P_2 and we get our conclusion. □

We now consider the scenario in which there is a second bad edge P_2P_3 and further explore the local structure of the associated TSP on the centers. Specifically, let O_3 be the center of the disk visited next at P_3 and assume that the edge P_2P_3 is also bad. Notice that the angle $\angle O_1O_2O_3$ formed by the TSP is either $\angle O_1O_2P_2 + \angle P_2O_2O_3$ or $|\angle O_1O_2P_2 - \angle P_2O_2O_3|$. Regardless, we have that $\angle O_1O_2O_3 \leq \angle O_1O_2P_2 + \angle P_2O_2O_3$ and get the following corollary:

Corollary 6 *If both P_1P_2 and P_2P_3 are bad edges, then the angle $\angle O_1O_2O_3$ is $\leq 2\beta$.*

We now have that if P_1P_2 and P_2P_3 are bad edges and O_1O_2 is small, then the TSP edge O_2O_3 edge must itersect the disk centered at O_1. In general, it is not true that if O_2O_3 intersects the first disk, we immediately get that the associated TSPN edge P_2P_3 must also intersect it. In our case, however, we have that the slope of P_2P_3 is very close to the one of O_2O_3 due to the fact that it is a bad edge. We use this information to show that if O_2O_3 does not intersect the first disk, then P_2P_3 cannot be a bad edge.

Theorem 7 *If P_1P_2 and P_2P_3 are bad edges and $O_1O_2 \leq R/\sin(2\beta)$, then the segment P_2P_3 intersects the disk centered at O_1.*

Proof. We consider the case in which $\angle O_1O_2O_3 = \angle O_1O_2P_2 + \angle P_2O_2O_3$ and note that all the other cases are similar. We denote the two lines originating at O_2 that are tangent to the first circle as ℓ_1 and ℓ_2 such that the line O_2O_3 is in between ℓ_1 and O_2O_1. Note that this is possible because the angle that ℓ_1 forms with O_2O_1 is at least 2β (since $O_1O_2 \leq R/\sin(2\beta)$) but the angle that O_2O_3 forms with O_2O_1 is at most 2β (Corollary 6).

Our strategy will be to first show that the segment P_2P_3 is contained in the wedge defined by ℓ_1 and ℓ_2 (Figure 3). Notice that, since the wedge defines a convex space, it is enough to show that P_2 and P_3 are contained in it.

We first show that the point P_2 has to be in the wedge. Let S_1 and S_2 be the points in which the segment O_1O_2 intersects the first and second disk. Similarly, let T_2 and T_3 be the points in which O_2O_3 intersects the second and third disk. We then have that P_2 is between T_2 and S_2.

Now we only need to show that P_3 is in between ℓ_1 and ℓ_2. We will do that by arguing that any choice of P_3 outside of the wedge will contradict the fact that P_2P_3 is a bad edge.

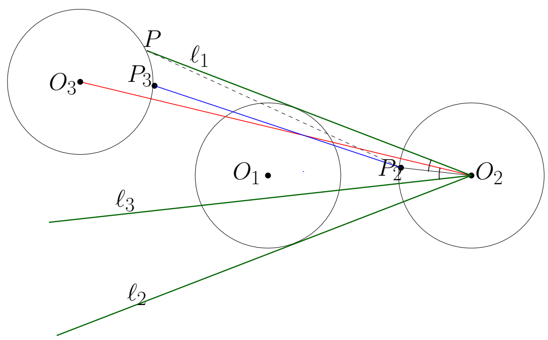

Figure 3: When O_2O_3 crosses the disk centered at O_1, we must also have that the segment P_2P_3 also crosses it. We show this by arguing that P_2P_3 is contained between the two lines ℓ_1 and ℓ_3 and that P_2 and P_3 are on separate sides of the first disk.

Let $\alpha_1 = \angle O_1O_2P_2$ and $\alpha_2 = \angle P_2O_2O_3$, and so $\alpha_1, \alpha_2 \leq \beta$ (Lemma 5). First notice that if neither ℓ_1 nor ℓ_2 intersect the third disk, then we are done because we have that the entire boundary is contained in the convex space (since O_3 is already in between ℓ_1 and ℓ_2). Assume then that ℓ_1 intersects the third disk at a point P above the line O_2O_3, since $\angle O_3O_2O_1 \leq 2\beta \leq \angle PO_2O_1$. Moreover, since $\angle PO_2O_1 \geq 2\beta$, we have that $\angle PO_2P_2 \geq 2\beta - \alpha_1 \geq \beta$ and so $PP_2 \geq f(PO_2, \beta)$ (because $|P_2O_2| = R$). Since $|PO_2| \geq |T_3O_2|$ and $|T_3O_2| \geq R$, this implies that $|PP_2| \geq f(T_3O_2, \beta)$. Using the fact that $f(x, \beta) \geq x - R\cos\beta$ for any x and $\beta \neq 0$, one can verify that:

$$\begin{aligned} f(T_3O_2, \beta) &= f(O_2O_3 - R, \beta) \\ &= \sqrt{(|O_2O_3| - R)^2 + R^2 - 2R(|O_2O_3| - R)\cos\beta} \\ &> \sqrt{|O_2O_3|^2 + R^2 - 2R|O_2O_3|\cos\beta} - R \\ &> f(O_2O_3, \beta) - R. \end{aligned}$$

This means that P cannot be a possible position for P_3 because then P_2P_3 would be too big. Moreover, any point Q "above" P (i.e. such that $\angle O_3O_2Q > \angle O_3O_2P$) would also not work as a possible position for the same reason. In other words, P_3 has to be underneath the line $PO_2 = \ell_1$.

In order to prove that P_3 is also above the line ℓ_2, we will consider an additional line ℓ_3 originating at O_2 that makes an angle of β with O_2P_2 and is underneath it. This new line makes an angle of $\beta + \alpha_2$ with O_2O_3 and since ℓ_2 makes an angle of $\geq 2\beta + \alpha_1 + \alpha_2$ with O_2O_3, we get that ℓ_3 is in between O_2O_3 and ℓ_2. In other words, if we show that P_3 is above ℓ_3, then we also get that P_3 is above ℓ_2. If ℓ_3 does not intersect the third disk, then we are done as before, so assume that it intersects it at a point Q on the boundary. Similarly as before, we have that $P_2Q = f(O_2Q, \beta) \geq f(O_2T_3, \beta) > f(O_2O_3, \beta) - R$. This in turn implies that P_3 has to be above Q, otherwise P_2P_3 would be too big. Therefore P_3 must be above the line ℓ_3.

At this point, we have that the segment P_2P_3 is contained in the wedge defined by ℓ_1 and ℓ_2. We know that the first disk is tangent on both sides to ℓ_1 and ℓ_2 but this does not directly imply that P_2P_3 must actually intersect it. In order to have that, we must also ensure that P_2 and P_3 lie on different sides of the first disk. We argue this by showing that O_3 itself must be on the other side of the first disk as O_2. Since the disks do not intersect, this implies that P_3 is on a different side from P_2. In order to show this, notice that we can assume, without loss of generality, that $O_1O_2 \leq O_2O_3$. Let T be the point on O_2O_3 such that $O_1T \perp O_2O_3$. Since $\angle O_1O_2O_3 \leq 2\beta$ and $O_1O_2 \leq R/\sin(2\beta)$, this means that T is contained in the first disk. Suppose that O_3 is on the segment O_2T (effectively in between O_1 and O_2). Then $O_2O3 < O_2T$ but, since $O_2T = O_1O_2\cos(\angle O_1O_2T) \leq O_1O_2$, this would lead to a contradiction. We therefore get that O_2 and O_3 are on different sides of the first disk and that the same is true for P_2 and P_3. This shows that the segment P_2P_3 must intersect the first disk. □

We are now ready to prove **Theorem 2**: *For $n \geq 4$, if P_1P_2 and P_2P_3 are bad edges and $O_1O_2 \leq R/\sin(2\beta)$ then at least one of the following is true:*

- *the TSPN tour is not optimal,*
- *the TSPN tour is supported by a straight line or*
- *the path $P_n - P_1 - P_2 - P_3$ forms a β-triad.*

Proof. We distinguish between the case in which P_2P_3 intersects the first disk at P_1 and otherwise. In the first case, we will show that either the TSPN is not optimal or all the disks are stabbed by it. The second case is more involved and reduces to describing what the local structure of the TSPN must be such that it does not necessarily fall in the previous two cases.

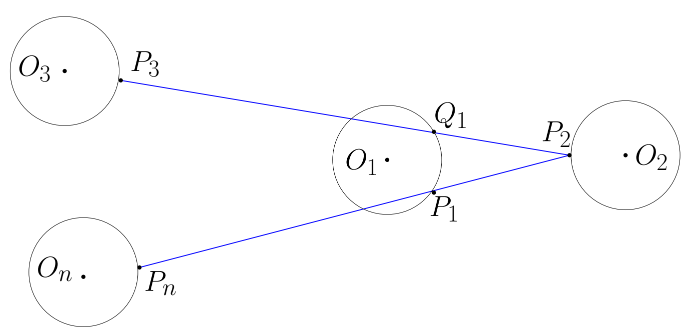

Figure 4: The path $P_n - P_1 - P_2 - P_3$ forms a β-triad.

Case 1: P_1, P_2, P_3 are collinear. Then consider the point P_n that connects to P_1. The cost that the TSPN pays for visiting the four disks is $|P_nP_1| + |P_1P_2| + |P_2P_3|$ but by triangle inequality, we know that $|P_nP_2| \leq |P_nP_1| + |P_1P_2|$, so the TSPN would visit P_2 directly and pass through P_1 on its way to P_3. If the inequality is strict, then this directly implies that the TSPN is not optimal. When we have equality, however, this implies that P_n, P_1 and P_2 are now also collinear and furthermore, that P_1 lies between P_2 and P_n. In other words, we have that on the line from P_2 to P_n, we have both P_3 and P_n to the left of P_1. Now look at how point P_4 connects to P_3 and notice that the portion of TSPN for the five disks is now $|P_4P_3| + |P_3P_2| + |P_1P_2| + |P_1P_n|$ and again, we can ask the question of why wouldn't the TSPN go straight to P_2 instead and visit P_3 along the line P_2P_3. Specifically, we have $|P_4P_2| \leq |P_4P_3| + |P_3P_2|$ with the TSPN not being optimal whenever this inequality is strict. We therefore consider the case in which $|P_4P_2| = |P_4P_3| + |P_3P_2|$ and get that now P_4 has to also be collinear with the other points and furthermore, P_3 has to be between P_4 and P_2. Continuing this process, we get that all the TSPN points would have to be collinear and in the order $P_2, P_1, P_3, P_4, \ldots, P_{n-1}$ with P_n potentially being anywhere past P_1. In this case, we have that the TSPN is a supported by a straight line that stabs all of the disks.

□

Case 2: P_1, P_2, P_3 are not collinear. Let the line P_1P_2 intersect the first disk for the first time at Q_1. By the argument from before, we know that if $|P_nP_2| < |P_nP_1| + |P_1P_2|$, then the TSPN cannot be optimal since another tour could go from P_n straight to visiting P_2 and then visit P_1 on the way to P_3, at a lesser cost. When P_n, P_1 and P_2 are collinear, in that order, we say that $P_n - P_1 - P_2 - P_3$ form a β-triad.

B Proof of Lemma 8 : β-triads are Edge Disjoint

As a reminder, we defined a bad triad to be a subpath of the tour $P_n - P_1 - P_2 - P_3$ that visits the disks centered at O_n, O_1, O_2 and O_3 and has all the following properties:

- P_1P_2 is a bad edge, i.e. $|P_1P_2| \leq f(O_1, O_2, \beta) - R$,
- P_2P_3 is a bad edge, i.e. $|P_2P_3| \leq f(O_2, O_3, \beta) - R$,
- O_1O_2 is short, i.e. $|O_1O_2| \leq R/\sin(2\beta)$,
- P_1, P_2, P_3 are not collinear and
- P_n, P_1, P_2 are collinear with P_1 between P_n and P_2.

Theorem 2 says that if $|TSPN^*|$ is not a straight line, then the triad has a local detour of at most $3\sqrt{3}R$. Lemma 8 further states that all the bad triads are also edge disjoint. In order to prove that, we go back to the proof of Theorem 2. Note that we distinguished between the case in which P_1, P_2 and P_3 are collinear (Case 1) and when they are not(Case 2). The first case leads to the TSPN being a straight line, which is ruled out by our assumptions. In the second case, the optimality of TSN^* implies that P_n, P_1 and P_2 are also collinear, with P_1 between P_2 and P_n.

Lemma 8 *All the β-triads in a given TSPN tour are edge disjoint.*

Proof. Assume there is another bad triad that shares edges with $P_n - P_1 - P_2 - P_3$. We distinguish four cases, based on the type of edges they have in common.

Case 1: $P_{n-1} - P_n - P_1 - P_2$ is a bad triad. This case cannot happen since P_n, P_1, P_2 are collinear.

Case 2: $P_{n-2} - P_{n-1} - P_n - P_1$ is a bad triad. Then, by definition, we must have that P_nP_1 is also a bad edge. We will show, however, than this cannot be. For this, we will use an additional lemma:

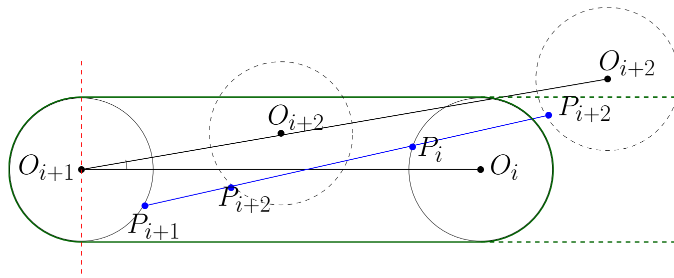

Figure 5: A potential TSPN path is drawn in blue. Because the angle $\angle O_iO_{i+1}O_{i+2} \leq \pi/2$, we have that O_i and O_{i+1} are on the same side of the hyperplane described by the red line. That, in turns, give two options for the disk centered at O_{i+2} to intersect the (extended) convex hulls of the other two diskss, drawn in green. In each case, the points are visited in the wrong order.

Lemma 9 *If P_iP_{i+2} is a straight line that passes through point P_{i+1} such that P_{i+1} is between P_i and P_{i+2}, then it cannot be that both P_iP_{i+1} and $P_{i+2}P_{i+1}$ are bad edges.*

Proof. Assume that both P_iP_{i+1} and $P_{i+2}P_{i+1}$ are bad edges. Then Corollary 6 implies that the angle $\angle O_iO_{i+1}O_{i+2} \leq 2\beta$. Now consider the convex hull of the two disks centered at O_i and O_{i+1} (Figure 5). If the disk centered at O_{i+2} intersects the convex hull, then P_{i+2} must be contained in that convex hull, otherwise the line $P_i - P_{i+1} - P_{i+2}$ would not exist. But in that case, the points would be visited out of order. Specifically, P_{i+2} would be between P_i and P_{i+1}.

Now assume that the disk centered at O_{i+2} does not intersect the convex hull. Since the angle $\angle O_iO_{i+1}O_{i+2} \leq 2\beta \leq \pi/6$, this implies that O_{i+2} is in the same halfspace as O_i with respect to the line perpendicular to O_iO_{i+1} passing through O_{i+1}. We extend the convex hull infinitely in that halfspace by allowing the tangent lines to be infinite on that side. By the same argument as before, we know that the disk centered at O_{i+2} must intersect this extended region. But then we would get again that the points are out of order: P_i would be between P_{i+1} and P_{i+2}. □

When $P_i = P_n, P_{i+1} = P_1$ and $P_{i+2} = P_2$, Lemma 9 tells us that it cannot be that P_1 is between P_n and P_2 and both edges P_1P_n and P_1P_2 are bad. Therefore we are done with this case.

Case 3: $P_1 - P_2 - P_3 - P_4$ is a bad triad. This case is similar to Case 1 and cannot happen, since P_1, P_2 and P_3 cannot be collinear.

Case 4: $P_2 - P_3 - P_4 - P_5$ is a bad triad. This case is similar to Case 2 because we have that P_2, P_3 and P_4 are collinear with P_3 between P_2 and P_4 and both P_2P_3 and P_3P_4 being bad edges.

□

C The TSPN that follows a straight line

Here we focus on the second possibility in Theorem 2 in which the optimal TSPN is supported by a straight line that stabs all the disks. We show that in this case, we can return in polynomial time a solution that is within an additive factor of $4R$ from the optimal $TSPN^*$. We note that when the TSPN might not be a line but the disks themselves admit a line transversal, a $\sqrt{2}$-approximation follows from the work of Dumitrescu and Mitchell [14]. We explain the result for completeness.

We start by identifying the centers that are the farthest apart and considering the direction orthogonal to the line going through them. This direction induces parallel segments of length $2R$ in each of the disks (that each go through the centers). It is easy to check that any line transversal through the disks is a line transversal through the segments except for the first and last disk in the associated geometric permutation (for those two disks, the TSPN will stop at the boundary of the disk and never cross the entire circle). Conversely, any line transversal through the segments will automatically also stab the disks. Now compute a shortest line segment that stabs all of these segments in time $O(n \log n)$ using the algorithm of Bhattacharya et al. [5]. We note that this is optimal up to an additive factor of $4R$ that comes from the fact that the optimal $TSPN^*$ might have to travel $4R$ to hit the first and the last two segments in the geometric permutation.

In general, when we know that the disks admit a line transversal, we can output a solution that is a $\sqrt{2}$-approximation [14]. This follows indirectly from an algorithm used for connected regions of the same diameter, when there is a line that stabs all of the diameters. Given the parallel segments of length $2R$ that we constructed earlier, we know that they can also be stabbed by a line. Now consider the smallest perimeter axis-aligned rectangle that intersects all of the segments, of width w and height h. This will be the solution that we return. Arkin and Hassin [2] argued that any tour which touches all four sides of the rectangle must have length at least $2\sqrt{h^2+w^2}$. Since $h+w \leq \sqrt{2}\cdot\sqrt{h^2+w^2}$, we get that the rectangle is a $\sqrt{2}$-approximation.

D Using different thresholds

In this section, we give a more general analysis of the general approximation factor with a parameter $\alpha > 1$ that we will set later in the proof. We include here only the aspects that change. Depending on whether $K \leq \frac{n}{\alpha}$ or not, we will employ different lower bounds on $|TSPN|$, in a similar fashion as before.

Case 1: when $K \geq \frac{n}{\alpha}$. In this situation, we have that:

$$|TSP^*| \leq |TSPN^*| + \frac{1+\cos\beta+2(\alpha-1)}{\alpha}\cdot R \cdot n.$$

Using Lemma 4 gives us that

$$\begin{aligned}|TSP^*| \leq &\Big(1+\frac{4}{\pi}\cdot\frac{1+\cos\beta+2(\alpha-1)}{\alpha}\Big)\cdot|TSPN^*|+\\ &+4\cdot\frac{1+\cos\beta+2(\alpha-1)}{\alpha}R\\ \leq &\Big(1+\frac{4}{\pi}\cdot\frac{1+\cos\beta+2(\alpha-1)}{\alpha}\Big)\cdot|TSPN^*|+8R\\ \leq &\Big(1+\frac{8}{\pi}-\frac{4}{\pi}\cdot\frac{1-\cos\beta}{\alpha}\Big)\cdot|TSPN^*|+8R.\end{aligned}$$

Case 2: when $K < \frac{n}{\alpha}$. Similarly as before, we get that:

$$\begin{aligned}|TSPN| &\geq \sum_{j=1}^{l}|\mathcal{B}_j|\\ &\geq \Big(\frac{1}{\sin(2\beta)}-2\Big)R\cdot(n-K)\\ &\geq \frac{\alpha-1}{\alpha}\cdot\Big(\frac{1}{\sin(2\beta)}-2\Big)R\cdot n.\end{aligned}$$

Since the total detour could be at most $2R$ per edge, we get that:

$$|TSP^*| \leq \Big(1+\frac{\alpha}{\alpha-1}\cdot\frac{2}{\frac{1}{\sin(2\beta)}-2}\Big)\cdot|TSPN|.$$

Setting $\alpha = 1+2/(c/\sin(2\beta)-2c-2)$ for $c = 2.53$ and $\beta = 0.1831$ gives us that both of these cases lead to a 2.53-approximation.

E Algorithm for overlapping disks

We discuss how the analysis from the disjoint case carries over to the case of overlapping disks. As we mentioned before, the best known approximation for this case is by Dumitrescu and Tóth [15]. In general, approaches for this case take advantage of known analyses for the disjoint case and adapt them in a smart way to the overlapping case. We begin by roughly describing the technique of Dumitrescu and Tóth [15] and the show how the analysis changes when we use our framework.

Specifically, Dumitrescu and Tóth [15] start by computing a monotone maximal set of disjoint disks $\mathcal{I}$ by greedily selecting the leftmost disk and deleting all of the other input disks that intersect it. Let k be the size of the set we end up wth. They then compute an approximate TSP tour on the centers of the disks in $\mathcal{I}$, either using the available schemes [3, 25] or Christofides [11]. We call this tour $T_{\mathcal{I}}$. They then augment this tour in such a way that we visit all the input disks, not just the ones in $\mathcal{I}$. Before we discuss the augmentation part, we first define some notation and mention some bounds that follow naturally.

Let the optimal TSP tour on the centers in $\mathcal{I}$ be $TSP^*_{\mathcal{I}}$. The eventual tour $T_{\mathcal{I}}$ that we compute will be an a-approximation to $TSPN^*_{\mathcal{I}}$ so we have that:

$$|T_{\mathcal{I}}| \leq a\cdot|TSP^*_{\mathcal{I}}|. \quad (2)$$

On the other hand, we know that this set of disks also has an associated optimal TSPN tour, which we call $TSPN^*_I$.

Finally, we denote the optimal TSPN tour on all the disks by $TSPN^*$. We know that the tour on $\mathcal{I}$ is a lower bound:

$$|TSPN^*_{\mathcal{I}}| \leq |TSPN^*|. \quad (3)$$

The size of our final solution will be compared to $|TSPN^*|$ and to that end, we use lower bounds on $|TSPN^*_{\mathcal{I}}|$ in conjunction with (3) to get lower bounds on $|TSPN^*|$. This is the part where our new framework will come in, because $|TSPN^*_{\mathcal{I}}|$ is a tour on disjoint disks by definition.

The next step is to augment $T_{\mathcal{I}}$ with detours of length $O(R)$ along the disks in $\mathcal{I}$ such that it touches every other disk not in $\mathcal{I}$. The total length of the solution would then become $|\mathcal{T}_{\mathcal{I}}| + O(1) \cdot |\mathcal{I}| \cdot R$. Specifically, Dumitrescu and Tóth [15] consider short curves around each disk in $\mathcal{I}$ that are guaranteed to cross any of the disks to its right that intersect it. Because the maximal set was chosen from left to right, that covers all the disks that could possibly intersect it. We refer the reader to [15] for the detailed construction. The authors show that the length of the resulting tour T is within $O(1) \cdot |\mathcal{I}| \cdot R$ of $|T_{\mathcal{I}}|$:

$$|T| \leq |T_{\mathcal{I}}| + (A \cdot k + B) \cdot R, \quad (4)$$

where $A = 2 \cdot (\frac{\pi}{6} + \sqrt{3} - 1)$ and $B = 4 - \sqrt{3}$.

Combining 2 and 4, we upper bound the length of the solution $|T|$ in terms of $|TSP^*_{\mathcal{I}}|$ as such:

$$\begin{aligned} |T| &\leq |T_{\mathcal{I}}| + (A \cdot k + B) \cdot R \\ &\leq a \cdot |TSP^*_{\mathcal{I}}| + (A \cdot k + B) \cdot R \end{aligned}$$

In order to complete the analysis, we would need to bound $|TSP^*_{\mathcal{I}}|$ in terms of $|TSPN^*|$ and we do that through $|TSPN^*_{\mathcal{I}}|$. The analysis from Dumitrescu and Tóth [15] uses the bounds from Dumitrescu and Mitchell [14] for the case of disjoint disks. Specifically, they apply Lemma 4 to get that:

$$kR \leq \frac{4}{\pi} \cdot |TSPN^*_{\mathcal{I}}| + 4R.$$

This, together with the bound $|TSP^*_{\mathcal{I}}| \leq |TSPN^*_{\mathcal{I}}| + 2Rk$ and (3) yields:

$$\begin{aligned} |T| &\leq a \cdot |TSP^*_{\mathcal{I}}| + (Ak + B) \cdot R \\ &\leq a \cdot (|TSPN^*_{\mathcal{I}}| + 2Rk) + (Ak + B) \cdot R \\ &\leq a \cdot |TSPN^*_{\mathcal{I}}| + (2a + A) \cdot kR + BR \\ &\leq a \cdot |TSPN^*_{\mathcal{I}}| + (2a + A) \cdot \left(\frac{4}{\pi}|TSPN^*_{\mathcal{I}}| + 4R\right) + BR \\ &\leq \left(a + (2a + A)\frac{4}{\pi}\right) \cdot |TSPN^*_{\mathcal{I}}| + (8a + 4A + B)R \\ &\leq \left((1 + \frac{8}{\pi})a + \frac{4A}{\pi}\right) \cdot |TSPN^*_{\mathcal{I}}| + (8a + 4A + B)R \\ &\leq \left((1 + \frac{8}{\pi})a + \frac{4A}{\pi}\right) \cdot |TSPN^*| + (8a + 4A + B)R \end{aligned}$$

Plugging in the values for A and B gives an overall approximation term of:

$$(1 + \frac{8}{\pi})a + \frac{4A}{\pi} \leq \left(\frac{7}{3} + \frac{8\sqrt{3}}{\pi}\right) \cdot (1 + \epsilon) \leq 6.75 \cdot (1 + \epsilon).$$

Our framework changes the last stage in which we compare $|TSP^*_{\mathcal{I}}|$ with $|TSPN^*_{\mathcal{I}}|$. We do a similar analysis as in the disjoint case, except for the tour on $\mathcal{I}$. We get that **Case 1** would therefore correspond to getting that:

$$|TSP^*_{\mathcal{I}}| \leq |TSPN^*_{\mathcal{I}}| + X \cdot R \cdot k,$$

where $X = 2 - \dfrac{1 - \cos\beta}{\alpha}$ (instead of $2R$). We can then replace it in the analysis and get:

$$\begin{aligned} |T| &\leq a \cdot (|TSPN^*_{\mathcal{I}}| + XRk) + (Ak + B) \cdot R \\ &\leq \left(a + (Xa + A)\frac{4}{\pi}\right) \cdot |TSPN^*_{\mathcal{I}}| + (8a + 4A + B)R \\ &\leq \left((1 + \frac{4X}{\pi})a + \frac{4A}{\pi}\right) \cdot |TSPN^*| + (4Xa + 4A + B)R \end{aligned}$$

In **Case** 2, we have that the overall detour is $2Rk$, but there is a different lower bound on $|TSPN^*_{\mathcal{I}}|$:

$$|TSPN^*_{\rangle}| \geq Y \cdot Rk,$$

where $Y = \frac{\alpha - 1}{\alpha} \cdot \Big(1/(2\sin(2\beta)) - 1\Big)$. Using the fact that $Rk \leq 1/Y \cdot |TSPN^*_{\mathcal{I}}|$, the analysis then becomes:

$$\begin{aligned} |T| &\leq \alpha \cdot |TSPN^*_{\mathcal{I}}| + (2\alpha + A) \cdot kR + BR \\ &\leq \alpha \cdot |TSPN^*_{\mathcal{I}}| + \frac{2\alpha + A}{Y} \cdot |TSPN^*_{\mathcal{I}}| + BR \\ &\leq \left(\alpha + \frac{2\alpha + A}{Y}\right) \cdot |TSPN^*_{\mathcal{I}}| + BR. \end{aligned}$$

If we set α and β like in the previous section, we get that both of the approximation factors are upper bounded by 6.728.

F The Fermat-Weber Point Approach for $n = 3$

In this section, we prove that the Häme, Hyytiä and Hakula conjecture is true for $n = 3$ and discuss a different way of looking at the TSPN tour that we believe might be of independent interest. We start with the observation that the shortest tour on the centers is equivalent to the shortest tour on translates of those centers, as long as all those centers are translated according to the same vector. In other words, if we fix a direction and translate each center along that direction until it reaches its boundary, the shortest tour on the newly obtained points will be exactly the same as the shortest tour on the centers themselves.

Formally, let B_i be the point we obtain by translating the center O_i along a fixed vector of length R. Then the TSP on the points $B_1, B_2, \ldots, B_n$ (in that order) has the same length as the TSP tour on $O_1, O_2, \ldots O_n$ (Figure 6). One advantage of visiting the first set of points (instead of the center points) is that it might be more similar geometrically to what the TSPN actually does. In terms of the following analysis, we would get that:

$$|TSP^*| \leq |TSPN^*| + 2\sum_{i=1}^{n}|P_iB_i|.$$

In this context, a natural question arises about the choice for the points B_i that minimizes the term $\sum_{i=1}^{n}|P_iB_i|$. In order to see what this best choice would be, we transform this input instance into another one by essentially superimposing all the disks on top of each other (Figure 7). Specifically, our

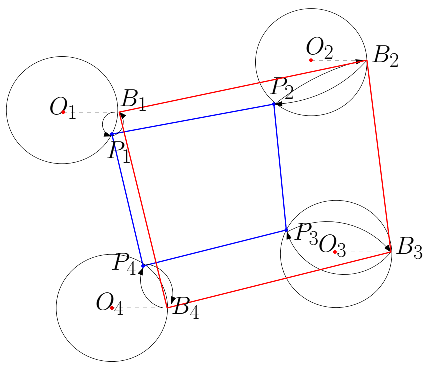

Figure 6: The translated view, when the tour visits the same point on the boundary of each disk.

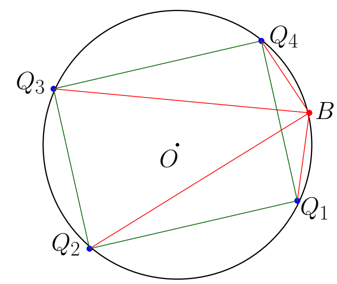

Figure 7: In red, the unified view when we translate each O_i to the same point B on the boundary. In green, the detour when we pick a different B for each pair of points O_i and O_{i+1}.

new instance will consist of one disk of radius R centered at a point O such that the points B_i map to a single point B (corresponding to O translated by the same fixed vector). We then map each point P_i of the TSPN to a corresponding point Q_i on the boundary of this disk such the vector OQ_i is a translate of the vector O_iP_i. We then get that:

$$\sum_{i=1}^{n}|P_iB_i| = \sum_{i=1}^{n}|Q_iB|,$$

and so the best choice for B is the one that minimizes the sum $\sum_{i=1}^{n}|Q_iB|$, otherwise know as the *Fermat-Weber point* or 1-*median* of the points $Q_1, Q_2, \ldots, Q_n$ [32, 31]. We note, however, that while the average distance to the Fermat-Weber point will never be greater than $2R$, there are instances in which this is tight. Consider, for example, the points Q_i to be the vertices of a convex $2n$-gon and notice by triangle inequality that the center of the disk is exactly their Fermat-Weber point (any other point will incur distances greater than the sum of the diagonals).

We can therefore say that when the points B_i are evenly spaced on the boundary of the disk the Fermat-Weber point is exactly the center and so we gain no improvement by moving the centers O_i towards the points B_i. It turns out, however, that the location of the points on the boundary is not as restrictive as the order in which the TSPN visits them. To see that, consider a different transformation in which we only move the centers O_i and O_{i+1} along a fixed vector. In other words, we choose a new vector for each pair of consecutive centers and only compare $|P_iP_{i+1}|$ locally against the newly obtained segment. This does not give us an overall valid tour on the centers, but it allows us to tailor the choice of B for each two points P_i and P_{i+1}. Specifically, we would get that:

$$|O_iO_{i+1}| \leq |P_iP_i + 1| + |Q_iB| + |Q_{i+1}B|.$$

In this case, we know that any point on the segment Q_iQ_{i+1} minimizes the distances in question and so we get that:

$$|O_iO_{i+1}| \leq |P_iP_i + 1| + |Q_iQ_{i+1}| \text{ and}$$
$$|TSP^*| \leq |TSPN^*| + \sum_{i=1}^{n}|Q_iQ_{i+1}|.$$

In other words, the largest detour obtained in this way is when the TSPN visits the points P_i in the order of the Maximum TSP on the associated points Q_i (Figure 7). The case in which all the points are evenly distributed along the boundary no longer becomes that restrictive. We can still construct, however, instances for which the Max TSP is exactly $2Rn$ and that is when the points visited are exactly diametrically opposite each other. Nevertheless, we are able to show that for $n = 3$, the detour is bounded by $3\sqrt{3}R$. Let A, B, C be any three points on the boundary of a circle of disk R centered at O. We then have that that $|AB| + |AC| + |BC| \leq 3\sqrt{3}R$ and the Häme, Hyytiä and Hakula conjecture for $n = 3$ follows:

Theorem 10 *For $n = 3$, we have that any tour which visits the disks in an order σ satisfies the bound*

$$|TSP(\sigma)| \leq |TSPN(\sigma)| + 3\sqrt{3}R.$$