New Approximation Algorithms for Touring Regions

Benjamin Qi
Massachusetts Institute of Technology, Cambridge, MA, USA

Richard Qi
Massachusetts Institute of Technology, Cambridge, MA, USA

Abstract

We analyze the *touring regions problem*: find a $(1+\epsilon)$-approximate Euclidean shortest path in d-dimensional space that starts at a given starting point, ends at a given ending point, and visits given regions $R_1, R_2, R_3, \dots, R_n$ in that order.

Our main result is an $\mathcal{O}\left(\frac{n}{\sqrt{\epsilon}}\log\frac{1}{\epsilon}+\frac{1}{\epsilon}\right)$-time algorithm for touring disjoint disks. We also give an $\mathcal{O}\left(\min\left(\frac{n}{\epsilon},\frac{n^2}{\sqrt{\epsilon}}\right)\right)$-time algorithm for touring disjoint two-dimensional convex fat bodies. Both of these results naturally generalize to larger dimensions; we obtain $\mathcal{O}\left(\frac{n}{\epsilon^{d-1}}\log^2\frac{1}{\epsilon}+\frac{1}{\epsilon^{2d-2}}\right)$ and $\mathcal{O}\left(\frac{n}{\epsilon^{2d-2}}\right)$-time algorithms for touring disjoint d-dimensional balls and convex fat bodies, respectively.

2012 ACM Subject Classification Theory of computation → Computational geometry; Theory of computation → Approximation algorithms analysis

Keywords and phrases shortest paths, convex bodies, fat objects, disks

Digital Object Identifier 10.4230/LIPIcs.SoCG.2023.54

Related Version *Full Version*: https://arxiv.org/abs/2303.06759

Acknowledgements We thank Dhruv Rohatgi, Quanquan Liu, Erik Demaine, Danny Mittal, Spencer Compton, William Kuszmaul, Timothy Qian, Jonathan Kelner, and Chris Zhang for helpful discussions. Special thanks to Xinyang Chen for writing one of the proofs in the full version of this paper and providing inspiration.

1 Introduction

We analyze the *touring regions problem*: find a $(1+\epsilon)$-approximate Euclidean shortest path in d-dimensional space that starts at a given starting point, ends at a given ending point, and visits given regions $R_1, R_2, R_3, \dots, R_n$ in that order. We primarily present algorithms for the cases where the regions R_i are constrained to be convex fat bodies or balls.[1] To the best of our knowledge, we are the first to consider the cases where regions are disjoint convex fat bodies or balls in arbitrary dimensions. Consequently, our algorithms use techniques not previously considered in the touring regions literature (Section 1.4). Our algorithms work under the assumption that a closest point oracle is provided; closest point projection has been extensively used and studied in convex optimization and mathematics [5, 16].

Most prior work focuses on $d=2$ or significantly restricts the convex bodies. The special case where $d=2$ and all regions are constrained to be polygons is known as the *touring polygons problem*. Dror et al. [9] solved the case where every region is a convex polygon exactly, presenting an $\mathcal{O}\left(|V|n\log\frac{|V|}{n}\right)$-time algorithm when the regions are disjoint as well as an $\mathcal{O}\left(|V|n^2\log|V|\right)$-time algorithm when the regions are possibly non-disjoint and the

[1] The full version also contains results for the case where the regions R_i are unions of general convex bodies.

39th International Symposium on Computational Geometry (SoCG 2023).
Editors: Erin W. Chambers and Joachim Gudmundsson; Article No. 54; pp. 54:1–54:16
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

subpath between every two consecutive polygons in the tour is constrained to lie within a simply connected region called a *fence*. Here, $|V|$ is the total number of vertices over all polygons. Tan and Jiang [19] improved these bounds to $\mathcal{O}(|V|n)$ and $\mathcal{O}(|V|n^2)$-time, respectively, without considering subpath constraints.

For touring nonconvex polygons, Ahadi et al. [3] proved that finding an optimal path is NP-hard even when polygons are disjoint and constrained to be two line segments each. Dror et al. [9] showed that approximately touring nonconvex polygons with constraining fences is a special case of 3D shortest path with obstacle polyhedra, which can be solved in $\tilde{\mathcal{O}}\left(\frac{e^4}{\epsilon^2}\right)$ time by applying results of Asano et al. [4], where e is the total number of edges over all polyhedra. Mozafari and Zarei [13] improved the bound for the case of nonconvex polygons with constraining fences to $\tilde{\mathcal{O}}\left(\frac{|V|^2 n^2}{\epsilon^2}\right)$ time. Ahadi et al. [3] also solve the *touring objects problem* exactly in polynomial time, in which the R_i are disjoint, nonconvex polygons and the objective is to visit the border of every region without entering the interior of any region.

For touring disjoint disks, a heuristic algorithm with experimental results was demonstrated by Chou [7]. Touring disjoint *unit* disks was given in a programming contest and was a source of inspiration for this paper; an $\mathcal{O}\left(\frac{n}{\epsilon^2}\right)$-time algorithm was given [1]. The main result that we show for disks is superior to both of these algorithms.

Polishchuk and Mitchell [17] showed the case where regions are constrained to be intersections of balls or halfspaces in d dimensions to be a special instance of a second-order cone program (SOCP), which runs in $\mathcal{O}\left(d^3 c^{1.5} n^2 \log \frac{1}{\epsilon}\right)$ time using SOCP time bounds as a black box. Here, c is the number of halfspace or ball constraints.

1.1 Formal problem description

▶ **Definition 1** (Approximate touring regions problem). *Given n sets of points (regions) $R_1, R_2, \ldots, R_n$ each a subset of $\mathbb{R}^d$, a starting point p_0, and an ending point p_{n+1},*[2] *define the function $D\colon (\mathbb{R}^d)^n \to \mathbb{R}$ as $D(p_1, p_2, \ldots, p_n) \triangleq \sum_{i=0}^{n} \|p_i - p_{i+1}\|_2$.*

Let $\mathcal{A} \triangleq \{(p_1, p_2, \ldots, p_n) \mid \forall i, p_i \in R_i\} \subseteq (\mathbb{R}^d)^n$. Find a tuple of points (tour) $(p'_1, p'_2, \ldots, p'_n) \in \mathcal{A}$ such that $D(p'_1, p'_2, \ldots, p'_n) \leq (1+\epsilon) \min_{x \in \mathcal{A}} D(x)$.

We primarily consider two types of regions: convex fat bodies with constant bounded fatness and balls. Fat objects have been previously considered in a variety of computational geometry settings [12, 10, 15, 14].

▶ **Definition 2** (Bounded fatness). *We say that a convex region $R \subset \mathbb{R}^d$ is* ***fat*** *if there exist balls h, H with radii $0 < r_h \leq r_H$, respectively, that satisfy $h \subseteq R \subseteq H \subset \mathbb{R}^d$ and $\frac{r_H}{r_h} = \mathcal{O}(1)$.*

One element of the problem that has not yet been determined is how we represent the sets of points $R_1, R_2, \ldots, R_n$; this depends on what we restrict the regions to be:

- **Convex fat bodies:** We have access to each of the convex bodies R_i via a *closest point oracle*. This oracle allows us to call the function $closest_i(p)$ on some point p, which returns the point $p' \in R_i$ such that $\|p - p'\|$ is minimized in $\mathcal{O}(1)$ time (note that p' is unique due to convexity). Additionally, for each region, we are given the radius r_h of the inscribed ball (as described in Definition 2), and a constant upper bound on the quantity $\frac{r_H}{r_h}$ over all regions.
- **Balls:** For each ball in the input we are given its center $c \in \mathbb{R}^d$ and its radius $r \in \mathbb{R}_{>0}$.

[2] For convenience, some of our results define the degenerate regions $R_0 \triangleq \{p_0\}$ and $R_{n+1} \triangleq \{p_{n+1}\}$.

We consider the 2-dimensional and general d-dimensional cases separately. In the d-dimensional case, we assume d is a constant (for example, we say $2^d = \mathcal{O}(1)$). We also consider the possibly non-disjoint versus disjoint cases separately, where the latter is defined by the restriction $R_i \cap R_j = \emptyset$ for all $0 \le i < j \le n+1$.

Motivation for our model

When considering general convex bodies, it is natural to augment the model of computation with oracle access to the bodies, including membership, separation, and optimization oracles [11]. In fact, when solving the touring regions problem for general convex bodies, a closest point oracle is necessary even for the case of a single region, where the starting point is the same as the ending point and the optimal solution must visit the closest point in the region to the starting point. Closest point oracles can be constructed trivially when the bodies are constant sized polytopes or balls. Closest point oracles have been used in the field of convex optimization [8, 5].

Our representations for convex fat bodies and balls have the nice structure that the former "contains" the latter: a ball is a specific type of convex fat body, and we can trivially construct a closest point oracle for balls. We justify considering convex fat bodies as they are in some sense "between" balls and general convex bodies: they obey some of the packing constraints of balls.

1.2 Summary of results

Our results and relevant previous results are summarized in Tables 1 and 2. We obtain a $\mathcal{O}\left(\frac{n}{\epsilon^{2d-2}}\right)$ time algorithm for touring disjoint convex fat bodies. Notice that this bound is linear in n; in fact, we show that *any* FPTAS for touring convex fat bodies can be transformed into one that is linear in n (Lemma 14). If the regions are further restricted to be balls, we can apply our new technique of placing points nonuniformly, and the time complexity improves to $\mathcal{O}\left(\frac{n}{\epsilon^{d-1}}\log^2\frac{1}{\epsilon}+\frac{1}{\epsilon^{2d-2}}\right)$, which roughly halves the exponent of $\frac{1}{\epsilon}$ compared to the convex fat bodies algorithm while retaining an additive $\frac{1}{\epsilon^{2d-2}}$ term.

Our 2D-specific optimizations allow us to obtain superior time bounds compared to if we substituted $d = 2$ into our general dimension algorithms. For convex fat bodies, we obtain an algorithm with linear time dependence on both n and $\frac{1}{\epsilon}$. For our main result of touring disjoint disks, we combine our optimizations for convex fat bodies and balls with 2D-specific optimizations.

▶ **Theorem 18.** *There is an* $\mathcal{O}\left(\frac{n}{\sqrt{\epsilon}}\log\frac{1}{\epsilon}+\frac{1}{\epsilon}\right)$*-time algorithm for touring disjoint disks.*

With a new *polygonal approximation* technique, we use the result of [19] for touring polygons as a black box to obtain algorithms with a square root dependence on $\frac{1}{\epsilon}$, most notably an $\mathcal{O}\left(\frac{n^{3.5}}{\sqrt{\epsilon}}\right)$-time algorithm for touring 2D convex bodies and an $\mathcal{O}\left(\frac{n^2}{\sqrt{\epsilon}}\right)$-time algorithm for touring 2D disjoint convex fat bodies.

The $\mathcal{O}\left(c^{1.5}n^2\log\frac{1}{\epsilon}\right)$-time result for touring d dimensional convex bodies given by [17], where each body is an intersection of balls and half spaces (with a total of c constraints) can be applied specifically to balls to yield an $\mathcal{O}\left(n^{3.5}\log\frac{1}{\epsilon}\right)$-time algorithm. Our algorithms for touring disjoint disks and balls all take time linear in n and are thus superior when ϵ is not too small.

Representation	Runtime	Intersecting?	Source
Convex Polygons (Exact)	$\mathcal{O}(\lvert V\rvert n)$, $\mathcal{O}\left(\lvert V\rvert n^2\right)$	No, Yes	Touring Polygons [9], [19]
Convex (Oracle Access)	$\mathcal{O}\left(\frac{n^{2.5}}{\sqrt{\epsilon}}\right)$, $\mathcal{O}\left(\frac{n^{3.5}}{\sqrt{\epsilon}}\right)$	No, Yes	**Theorem 10**
Convex Fat (Oracle Access)	$\mathcal{O}\left(\frac{n}{\epsilon}\right)$, $\mathcal{O}\left(\frac{n^2}{\sqrt{\epsilon}}\right)$	No	**Theorems 16, 17**
Disks	$\mathcal{O}\left(\frac{n}{\sqrt{\epsilon}}\log\frac{1}{\epsilon}+\frac{1}{\epsilon}\right)$	No	**Theorem 18**

Table 1 Previous and new bounds on touring n regions in two dimensions up to multiplicative error $1+\epsilon$, where $\epsilon \leq \mathcal{O}(1)$. For polygons, $|V|$ is the total number of vertices over all polygons.

Representation	Runtime	Intersecting?	Source
Convex Bodies, each an intersection of balls or halfspaces	$\mathcal{O}\left(c^{1.5}n^2\log\frac{1}{\epsilon}\right)$	Yes	SOCP [17]
Convex Fat (Oracle Access)	$\mathcal{O}\left(\frac{n}{\epsilon^{2d-2}}\right)$	No	**Theorem 15**
Balls	$\mathcal{O}\left(\frac{n}{\epsilon^{d-1}}\log^2\frac{1}{\epsilon}+\frac{1}{\epsilon^{2d-2}}\right)$	No	**Theorem 19**

Table 2 Previous and new bounds on touring n regions in $d \geq 2$ dimensions up to multiplicative error $1+\epsilon$, where $\epsilon \leq \mathcal{O}(1)$. Note that d is treated as a constant. For polyhedra, c is the total number of constraints.

1.3 Organization of the paper

We start in Section 2 by introducing the general techniques used by all of our algorithms, including the closest point projection and 2D-specific optimizations. We then use the ideas of packing and grouping to obtain algorithms for convex fat bodies in Section 3. Finally, we optimize specifically for balls in Section 4 by placing points non-uniformly.

1.4 Summary of techniques

Here, we introduce the techniques mentioned in the previous subsection.

Placing points uniformly (Section 2)

A general idea that we use in our approximation algorithms is to approximate a convex body well using a set of points on its boundary. For previous results involving polygons or polyhedra [4, 13], this step of the process was trivial, as points were equally spaced along edges. In order to generalize to convex bodies in arbitrary dimensions, we equally space points on boundaries using the closest point projection oracle with a bounding hypercube (Lemma 4). After discretizing each body into a set of points, we can solve the problem in polynomial time using dynamic programming (DP): for each point, we find and store the optimal path ending at it by considering transitions from all points on the previous region.

2D-specific optimizations (Section 2)

When the input shapes are convex and disjoint, we use properties of Monge matrices to optimize dynamic programming transitions from quadratic to expected linear time (Lemma 5). Previous approximation algorithms for related problems discretize the boundary of each convex region using $\mathcal{O}\left(\frac{1}{\epsilon}\right)$ points. We present a new approach to approximate each boundary using a convex polygon with $\mathcal{O}\left(\frac{1}{\sqrt{\epsilon}}\right)$ vertices (Lemma 9). This allows us to use previous exact algorithms for touring convex polygons as black boxes.

Packing and grouping (Section 3)

The key ideas behind our improvements for disjoint convex fat bodies are *packing* and *grouping*. We use a simple *packing* argument to show that the path length for visiting n disjoint convex fat bodies with radius r must have length at least $\Omega(r \cdot n)$ for sufficiently large n (Lemma 11). This was used by [1] for the case of unit disks. However, it is not immediately clear how to use this observation to obtain improved time bounds when convex fat regions are not all restricted to be the same size. The idea of *grouping* is to split the sequence of regions into smaller contiguous subsequences of regions (groups). In each group, we find the minimum-sized region, called a *representative* region, which allows us to break up the global path into smaller subpaths between consecutive representatives. The earlier packing argument now becomes relevant here, as we can show a lower bound on the total length of the optimal path in terms of the sizes of the representatives.

Placing points non-uniformly (Section 4)

Previous approximation methods rely on discretizing the surfaces of bodies into evenly spaced points. For balls, we use the intuition that the portion of the optimal path from one ball to the next is "long" if the optimal path does not visit the parts of the surfaces that are closest together. This allows us to place points at a lower density on most of the surface area of each ball, leading to improved time bounds. We use this technique in conjunction with packing and grouping. For disks, we additionally apply the aforementioned 2D-specific optimizations.

2 General Techniques

First, we describe the general techniques used by all of our algorithms. We split the discussion into the general d-dimensional case and the 2-dimensional case.

2.1 General dimensions

The first main ingredient is the closest point projection, which allows us to equally space points on each convex body.

▶ **Lemma 3.** *For a convex region C, define $closest_C(p) \triangleq argmin_{c\in C} \|c-p\|$. For any two points p_1 and p_2, $\|closest_C(p_1) - closest_C(p_2)\| \leq \|p_1 - p_2\|$.*

For any closed set X, let ∂X denote the boundary of X.

▶ **Lemma 4** (Equal spacing via closest point projection). *Given a convex body C for which we have a closest point oracle and a hypercube $\mathcal{H}$ with side length r, we can construct a set $S \subset C$ of $\mathcal{O}\left(\frac{1}{\epsilon^{d-1}}\right)$ points such that for all $p \in (\partial C) \cap \mathcal{H}$, there exists $p' \in S$ such that $\|p - p'\| \leq r\epsilon$.*

Proof Sketch. First, we prove the statement for $C = \mathcal{H}$. For this case, it suffices to equally space points on each face of an axis-aligned hypercube defined by $[0, r]^d$. For example, for the face defined by $x_d = 0$, we place points in a lattice at all coordinates $(x_1, x_2, \ldots, x_{d-1}, x_d)$ that satisfy $x_d = 0$ and $x_i = k_i \cdot r\epsilon$ for all integers $k_i \in \left[0, \frac{1}{\epsilon}\right]$. For $C \neq \mathcal{H}$, equally space points on $\mathcal{H}$ as we stated to create a set $S_{\mathcal{H}}$. Then define $S \triangleq \{closest_C(s) \mid s \in S_{\mathcal{H}}\}$. ◀

The proof of Lemma 3 and the remainder of the proof of Lemma 4 are deferred to the full version of this paper.

2.2 Two dimensions

When the convex bodies are constrained to lie in 2D, there are two main avenues for further improvements: first, by speeding up the dynamic programming (DP) transitions when all regions have been discretized into point sets, and second, by approximating convex bodies by convex polygons instead of sets of points.

2.2.1 Dynamic programming speedup

▶ **Lemma 5.** *Given are the vertices of two disjoint convex polygons $B = [b_1, \dots, b_m]$ and $A = [a_1, a_2, \dots, a_n]$ in counterclockwise order and real weights $[w_1, \dots, w_n]$, one for each vertex of A. Define $d(i,j) \triangleq w_j + \|a_j - b_i\|$. Then $\min_{1 \le j \le n} d(i,j)$ may be computed for all $i \in [1,m]$ in $\mathcal{O}(m+n)$ expected time.*

Proof. We first discuss the case where all $w_i = 0$. Aggarwal and Klawe [2] showed how to reduce the computation of $\min_{1 \le j \le n, a_j \text{ visible from } b_i} d(i,j)$ and $\min_{1 \le j \le n, a_j \text{ not visible from } b_i} d(i,j)$ for all $i \in [1,m]$ to computing the row minima of several *Monge partial matrices* with dimensions $m_1 \times n_1, m_2 \times n_2, \dots, m_k \times n_k$ such that $\sum(m_i + n_i) \le \mathcal{O}(m+n)$ in $\mathcal{O}(m+n)$ time. Here, a_j is said to be visible from b_i if the segment $\overline{a_j b_i}$ intersects neither the interiors of polygons A nor B. The definition of Monge partial matrix can be found in [6]. Chan [6] recently introduced an $\mathcal{O}(m+n)$ expected time randomized algorithm for computing the row minima of an $m \times n$ Monge partial matrix.[3] Thus, the case of $w_i = 0$ can be solved in $\mathcal{O}(m+n)$ expected time.

The key claim that Aggarwal and Klawe [2] use to show that all the matrices they construct are Monge partial is as follows:

▷ Claim 6 (Lemma 2.1 of [2], adapted). Assume all $w_j = 0$. Suppose $j \ne j'$ and $i \ne i'$. If $a_j a_{j'} b_{i'} b_i$ form a convex quadrilateral in that order then $d(i,j) + d(i',j') \le d(i,j') + d(i',j)$.

The claim above holds by the triangle inequality, and it is easy to check that it still holds without the assumption $w_j = 0$. Thus the algorithm from [2] generalizes to the case of nonzero w_j with minor modifications. ◀

▶ **Corollary 7.** *The Touring Regions Problem in 2D, where all R_i are sets of finitely many points S_i that each form a convex polygon in counterclockwise order and the convex hulls of all S_i are disjoint, can be solved exactly in $\mathcal{O}\left(\sum_{i=1}^{n} |S_i|\right)$ expected time.*

Using the above techniques, the following result is proven in the full version and is used multiple times in Section 3.

▶ **Theorem 8.** *There is an $\mathcal{O}\left(n^2\left(\log\log n + \frac{1}{\epsilon}\right)\right)$-time algorithm for touring disjoint convex bodies in two dimensions. When the bodies are possibly non-disjoint, the bound is $\mathcal{O}\left(n^3\left(\log\log n + \frac{1}{\epsilon} + \frac{\log 1/\epsilon}{n\epsilon}\right)\right)$ time.*

[3] The Monge partial matrix does not have to be given explicitly; it suffices to provide an oracle that returns the value of any entry of the matrix in $\mathcal{O}(1)$ time.

2.2.2 Polygonal approximation algorithms

Up until now, we have approximated the perimeter of a convex region using points. We can alternatively approximate the perimeter using a convex polygon with fewer vertices. The proof is deferred to the full version.

▶ **Lemma 9** (Polygonal approximation). *Given a closest point oracle for a convex region C and a unit square U, we may select $\mathcal{O}\left(\epsilon^{-1/2}\right)$ points in C such that every point within $C \cap U$ is within distance ϵ of the convex hull of the selected points.*

The polygonal approximation allows us to immediately obtain the following result.

▶ **Theorem 10.** *There is a $\mathcal{O}\left(\frac{n^{2.5}}{\sqrt{\epsilon}}\right)$-time algorithm for touring 2D disjoint convex bodies. When the convex bodies are possibly non-disjoint, the bound is $\mathcal{O}\left(\frac{n^{3.5}}{\sqrt{\epsilon}}\right)$ time.*

Proof Sketch. After using Theorem 8 to find a constant approximation of the optimal path length, we draw a square of this side length around the starting point, and we know the optimal path must lie within the square. Then, we apply Lemma 9 to approximate each region with a convex polygon and use previous exact algorithms for touring polygons [19] to finish. ◀

3 Disjoint convex fat bodies

In this section, we present packing and grouping techniques for touring disjoint convex fat bodies and show how they can be applied to obtain $\mathcal{O}\left(\min\left(\frac{n}{\epsilon}, \frac{n^2}{\sqrt{\epsilon}}\right)\right)$-time algorithms for touring 2D disjoint convex fat bodies.

3.1 Techniques

3.1.1 Packing

A packing argument shows that the length of the optimal path length is at least linear in the number of bodies and the minimum r_h (that is, the minimum radius of any inscribed ball). Intuitively, if we place n disjoint objects of radius at least 1 that are close to being disks on the plane, the length of the optimal tour that visits all of them should be at least linear in n for sufficiently large n. The proof is deferred to the full version.

▶ **Lemma 11** (Packing Lemma). *Assume a fixed upper bound on $\frac{r_H}{r_h}$. Then there exists $n_0 = \mathcal{O}(1)$ such that the optimal path length OPT for touring any $n \geq n_0$ disjoint convex fat objects is $\Omega(n \cdot \min r_h)$. For balls, $n_0 = 3$.*

The packing lemma allows us to obtain a strong lower bound on the length of the optimal tour in terms of the size of the regions, which will be crucial in proving that our algorithms have low relative error.

▶ **Corollary 12.** *Let r_i denote the ith largest r_h. For all $i \geq n_0$, $r_i \leq \mathcal{O}\left(\frac{OPT}{i}\right)$.*

Proof. Consider dropping all regions except those with the i largest inner radii and let OPT_i be the optimal length of a tour that visits the remaining disks in the original order. By Lemma 11, for $i \geq n_0$, $OPT \geq OPT_i \geq \Omega(i \cdot r_i) \implies r_i \leq \mathcal{O}\left(\frac{OPT}{i}\right)$. ◀

▶ **Lemma 13.** *The optimal path length for touring n disjoint convex fat bodies is $\Omega\left(\sum_{i \geq n_0} r_i / \log n\right)$, and there exists a construction for which this bound is tight.*

Proof Sketch. Using Corollary 12,

$$\frac{\sum_{i\geq n_0} r_i}{\log n} \leq \sum_{i\geq n_0} \frac{\mathcal{O}\left(\frac{OPT}{i}\right)}{\log n} \leq \mathcal{O}\left(\frac{OPT}{\log n}\sum_{i=n_0}^{n}\frac{1}{i}\right) \leq \mathcal{O}\left(OPT\right).$$

We display the construction in Figure 1; we defer the full description to the full version. The idea is to place disjoint disks of radii $1/1, 1/2, 1/3, \ldots$ such that they are all tangent to a segment of the x-axis of length $\mathcal{O}(1)$. ◀

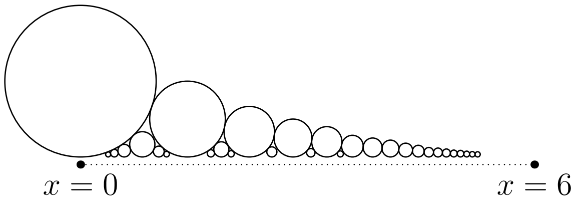

Figure 1 Construction from Lemma 13: placement of the first 30 disks.

3.1.2 Grouping

We now show that we can split up the optimal path into smaller subpaths by splitting the sequence of bodies into groups of consecutive bodies, finding the minimum-sized body in each group, and considering the subpaths between these small bodies. By the packing lemma, the sum of the radii of the representatives is small compared to the total path length.

In particular, using groups of size $\frac{1}{\epsilon}$, we can compress the smallest sized region into a single point, meaning that we can consider touring regions between these points independently from each other. This allows us to turn *any* polynomial time approximation scheme for touring disjoint convex fat bodies into one that is linear in n.

▶ **Lemma 14** (Grouping Lemma)**.** *Given an algorithm for touring disjoint convex fat bodies in d dimensions that runs in $f(n, \epsilon)$ time, where f is a polynomial, we can construct an algorithm that runs in $\mathcal{O}(n\epsilon + 1) \cdot f\left(\frac{1}{\epsilon}, \epsilon\right)$ time (for $\epsilon \leq \mathcal{O}(1)$).*

Proof. We describe an algorithm achieving a $(1 + \mathcal{O}(\epsilon))$-approximation. To achieve a $(1 + \epsilon)$-approximation, scale down ϵ by the appropriate factor.

Define $s \triangleq \left\lceil \frac{1}{\epsilon} \right\rceil$ and let n_0 be the constant defined in the statement of Lemma 11. We will prove the statement for all ϵ satisfying $\frac{1}{\epsilon} \geq n_0$. First, we divide the $n+2$ regions (including R_0 and R_{n+1}) into $k = \max\left(\left\lceil \frac{n+2}{s} \right\rceil, 2\right) \leq \mathcal{O}(n\epsilon + 1)$ consecutive subsequences, each with exactly s regions (except the starting and ending subsequences, which are allowed to have fewer). Let M_i be the region with minimum inscribed radius r_h in the ith subsequence; note that $M_1 = R_0$ and $M_k = R_{n+1}$. For each $i \in [1, k]$, pick an arbitrary point $p_i \in M_i$. Let OPT' be the length of the shortest tour of $R_0, \ldots, R_{n+1}$ that passes through all of the p_i. The $p_1, \ldots, p_k$ form $k-1$ subproblems, each with at most $2s$ regions. Therefore, we can $(1+\epsilon)$-approximate OPT' by $(1+\epsilon)$-approximating each subproblem in $(k-1) \cdot f(2s, \epsilon) \leq \mathcal{O}(n\epsilon + 1) \cdot f\left(\frac{1}{\epsilon}, \epsilon\right)$ time.

It remains to show that OPT' is a $(1+O(\epsilon))$-approximation for OPT. Let r_i be shorthand for the radius r_h of M_i ($r_1 = r_k = 0$). By the definition of fatness, the distance between any two points in M_i is at most $\mathcal{O}(r_i)$. By following through OPT and detouring to each point p_i, we get a path through points p_i with length at most $OPT + \mathcal{O}\left(\sum r_i\right)$, and OPT' is at most this amount.

The last remaining step is to show $\sum r_i \le \mathcal{O}(\epsilon \cdot OPT)$. We apply Lemma 11 to each subsequence, and obtain that $r_i s \le \mathcal{O}(OPT_i)$, where OPT_i is the optimal distance to tour regions in subsequence i. Note that although the starting and ending subsequences can have sizes less than s, they satisfy $r_i = 0$, so this bound holds for all subsequences. Therefore, $\sum r_i \le \mathcal{O}(\epsilon \cdot \sum OPT_i) \le \mathcal{O}(\epsilon \cdot OPT)$. ◀

3.2 Algorithms for convex fat bodies

Using a similar grouping argument, but using constant sized instead of $\frac{1}{\epsilon}$ sized groups, along with earlier methods of using estimates of the path length to place points on the boundaries of the convex fat bodies yields the following results.

▶ **Theorem 15.** *There is an $\mathcal{O}\left(\frac{n}{\epsilon^{2d-2}}\right)$-time algorithm for touring disjoint convex fat bodies in d dimensions.*

Proof. We proceed in a similar fashion as Lemma 14, except we define $s \triangleq n_0$, i.e., using constant sized groups instead of $\lceil \frac{1}{\epsilon} \rceil$ sized groups. Let the M_i be defined as in the proof of Lemma 14, and define m_i to be the outer radius of M_i.

For each pair of regions M_i, M_{i+1}, pick arbitrary points $a \in M_i, b \in M_{i+1}$, and use the d-dimensional analog of Theorem 8[4] to obtain a 4-approximation D_{approx} of the length of the shortest path from a to b in $\mathcal{O}(1)$ time. Suppose that the optimal path uses $p \in M_i, q \in M_{i+1}$ and the shortest path from a to b has distance $OPT_{a,b}$; by the triangle inequality, we must have

$$\frac{1}{4} D_{approx} \le OPT_{a,b} \le OPT_i + 2m_i + 2m_{i+1}.$$

Now, consider the path where we start at p and then travel along the line segment from p to a, the approximate path of length D_{approx} from a to b (visiting the regions in between M_i and M_{i+1}), and the line segment from b to q. This path has length at most $D_{approx} + 2m_i + 2m_{i+1}$, and upper bounds the length of the optimal path between p and q. So, the entire path between p and q lies within a ball of radius $D_{approx} + 4m_i + 2m_{i+1}$ centered at a; call this ball L. Note that L has radius $l = D_{approx} + 4m_i + 2m_{i+1} \le \mathcal{O}(OPT_i + m_i + m_{i+1})$.

For each region R_j between M_i and M_{i+1} inclusive, we apply Lemma 4 with the region and a hypercube containing L, which has side length $2l$. Note that points are placed twice on each M_i; this is fine. Lemma 4 guarantees the existence of a point in R_j that is $2l\epsilon$ close to the point OPT uses by placing $\mathcal{O}\left(\frac{1}{\epsilon^{d-1}}\right)$ points on each region.

We now bound the difference between the optimal and the shortest paths using only the points we placed. The difference is at most

$$\sum_{i=1}^{k} (2l_i \epsilon \cdot n_0) = \epsilon \cdot \mathcal{O}\left(\sum_{i=1}^{k} l_i\right) = \epsilon \cdot \mathcal{O}\left(OPT + \sum_{i=1}^{k} m_i\right) = \mathcal{O}(\epsilon \cdot OPT),$$

where the last step is due to Corollary 12 applied on each subsequence: in particular, the optimal path length visiting all the regions in subsequence i has length at least $\Omega(m_i)$, so summing this inequality over all subsequences, we have $\sum_{i=1}^{k} m_i \le \mathcal{O}(OPT)$.

We have now reduced the problem to the case where each region has only finitely many points. We finish with dynamic programming. Since we have $\mathcal{O}\left(\frac{1}{\epsilon^{d-1}}\right)$ points on each of the n regions, the runtime is $\mathcal{O}\left(\frac{n}{\epsilon^{2d-2}}\right)$, as desired. ◀

[4] This theorem may be found in Table 2 of the full version.

▶ **Theorem 16.** *There is an $\mathcal{O}\left(\frac{n}{\epsilon}\right)$-time algorithm for touring 2D disjoint convex fat bodies.*

Proof. This is almost the same as Theorem 15, where $\mathcal{O}\left(\frac{1}{\epsilon^{d-1}}\right) = \mathcal{O}\left(\frac{1}{\epsilon}\right)$ points are placed on each body, except that we use Corollary 7 to more efficiently solve the case where each region is a finite point set. ◀

▶ **Theorem 17.** *There is an $\mathcal{O}\left(\frac{n^2}{\sqrt{\epsilon}}\right)$-time algorithm for touring 2D disjoint convex fat bodies.*

Proof. Theorem 16 through the construction of Theorem 15 places $\mathcal{O}\left(\frac{1}{\epsilon}\right)$ points on an arc of length R on each convex fat body to guarantee additive error $\leq \epsilon R$. We can achieve the same additive error using a convex polygon with $\mathcal{O}\left(\epsilon^{-1/2}\right)$ vertices using Lemma 9. Then, recall that [19] gives an $\mathcal{O}(|V|n)$-time exact algorithm for touring convex polygons, so we can recover a solution in $\mathcal{O}(|V|n) = \mathcal{O}\left((n \cdot \epsilon^{-1/2}) \cdot n\right)$ time. ◀

4 Balls

We can improve the results in previous sections by discretizing the surfaces non-uniformly, placing fewer points on areas of each hypersphere that are farther away from the previous and next ball in the sequence. This reduces the dependence on ϵ by a square root compared to Theorem 15 and Theorem 16. We first state the results:

▶ **Theorem 18.** *There is an $\mathcal{O}\left(\frac{n}{\sqrt{\epsilon}} \log \frac{1}{\epsilon} + \frac{1}{\epsilon}\right)$-time algorithm for touring disjoint disks.*

▶ **Theorem 19.** *There is an $\mathcal{O}\left(\frac{n}{\epsilon^{d-1}} \log^2 \frac{1}{\epsilon} + \frac{1}{\epsilon^{2d-2}}\right)$-time algorithm for touring disjoint balls in d dimensions.*

The crucial lemma we use for these results follows. We defer its proof to the full version.

▶ **Lemma 20.** *A tour of disjoint balls is globally optimal if and only if for each intermediate ball, the tour either passes straight through the ball or perfectly reflects off its border (see Figure 2 for an example).*

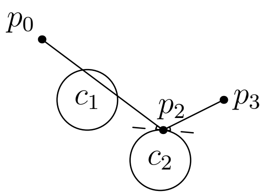

Figure 2 Lemma 20: An optimal tour of two unit disks. The tour starts at p_0, passes through c_1, reflects off c_2 at p_2, and ends at p_3.

We start with the special case of unit disks and then generalize to non-unit disks (Theorem 18). First, we provide intuition through a simple example where $n = 1$ and R_1 is a line.

▶ **Example 21.** Given start and endpoints $p_0 = (-1, 1)$ and $p_2 = (1, 1)$, select p_1 from the x-axis such that $OPT = \|p_0 - p_1\| + \|p_1 - p_2\|$ is minimized.

Solution. To solve this exactly, choose $p_1 = (0, 0)$ such that the path perfectly reflects off the x-axis. This gives $OPT = 2\sqrt{2}$.

Now suppose that we are only interested in an approximate solution. Tile the x-axis with points at regular intervals such that every two consecutive points are separated by distance d, and round p_1 to the closest such point p_1'. Since $\|p_1 - p_1'\| \le d$,

$$\begin{aligned}
OPT' &\triangleq \|p_0 - p_1'\| + \|p_1' - p_2\| \\
&\le \sqrt{1+(1-d)^2} + \sqrt{1+(1+d)^2} \le \sqrt{2-2d+d^2} + \sqrt{2+2d+d^2} \\
&\le \sqrt{2}(1 - d/2 + 1 + d/2 + \mathcal{O}\left(d^2\right)) \le 2\sqrt{2}(1+\mathcal{O}\left(d^2\right)).
\end{aligned}$$

So, to attain $OPT' \le (1+\epsilon)OPT$, it suffices to take $d = \Theta(\sqrt{\epsilon})$ rather than $d = \Theta(\epsilon)$ because $p_1' - p_1$ is parallel to the x-axis. We can apply a similar idea to replace the middle region with a point set when R_1 is a circle rather than a line since circles are locally linear. However, this doesn't quite work when either $\|p_0 - p_1\|$ or $\|p_1 - p_2\|$ is small. For example, if p_0 was very close to the x-axis (say, $p_0 = (-d, d)$) then rounding p_1 to the nearest p_1' could cause OPT' to increase by $\Theta(d) \gg d^2$. So when we replace each circle with a point set, we need to be careful about how we handle two circles that are close to touching; the solution is to space points more densely near where they touch. ◀

▶ **Theorem 22.** *There is an $\mathcal{O}\left(\frac{n}{\sqrt{\epsilon}}\log\frac{1}{\epsilon}\right)$-time algorithm for touring disjoint unit disks.*

Proof. We describe how to place a set of $\mathcal{O}\left(\frac{1}{\sqrt{\epsilon}}\log\frac{1}{\epsilon}\right)$ points S_i on each unit circle c_i so that the length of an optimal path increases by at most $\mathcal{O}\left(n\epsilon\right)$ after rounding each p_i to the nearest $p_i' \in S_i$. Define $\mathrm{unit}(x) = \frac{x}{\|x\|}$. Let $o_i \triangleq p_i' - p_i$ for all $i \in [0, n+1]$ (note that $o_0 = o_{n+1} = 0$), where o stands for offset. Also, define vectors

$$d_i \triangleq p_{i+1}' - p_i' = p_{i+1} + o_{i+1} - p_i - o_i$$

and scalars

$$a_i \triangleq d_i \cdot \mathrm{unit}(p_{i+1} - p_i) = \|p_{i+1} - p_i\| + (o_{i+1} - o_i) \cdot \mathrm{unit}(p_{i+1} - p_i),$$

where a_i is the component of d_i along the direction of $p_{i+1} - p_i$. Then the total path length after rounding each p_i to p_i' is:

$$\begin{aligned}
\sum_{i=0}^{n} \|d_i\| &= \sum_{i=0}^{n} \sqrt{[d_i \cdot \mathrm{unit}(p_{i+1} - p_i)]^2 + [d_i \cdot \mathrm{unit}(p_{i+1} - p_i)^{\perp}]^2} \\
&= \sum_{i=0}^{n} \sqrt{a_i^2 + [(o_{i+1} - o_i) \cdot \mathrm{unit}(p_{i+1} - p_i)^{\perp}]^2} \\
&= \sum_{i=0}^{n} \left[a_i + \left(\sqrt{a_i^2 + [(o_{i+1} - o_i) \cdot \mathrm{unit}(p_{i+1} - p_i)^{\perp}]^2} - a_i\right)\right] \\
&= OPT + \sum_{i=1}^{n} \overbrace{o_i \cdot (\mathrm{unit}(p_i - p_{i-1}) - \mathrm{unit}(p_{i+1} - p_i))}^{\mathrm{extra}_1(i)} \\
&\quad + \sum_{i=0}^{n} \overbrace{\left(\sqrt{a_i^2 + [(o_{i+1} - o_i) \cdot \mathrm{unit}(p_{i+1} - p_i)^{\perp}]^2} - a_i\right)}^{\mathrm{extra}_2(i)} \\
&= OPT + \sum_{i=1}^{n} \mathrm{extra}_1(i) + \sum_{i=0}^{n} \mathrm{extra}_2(i).
\end{aligned}$$

We defer the construction of the sets S_i so that both extra terms are small to Lemma 23. Then we can finish with dynamic programming (Corollary 7). ◀

▶ **Lemma 23.** *It is possible to choose S_i in the proof of Theorem 22 such that $|S_i| \leq \mathcal{O}\left(\frac{1}{\sqrt{\epsilon}} \log \frac{1}{\epsilon}\right)$, extra$_1(i) \leq \mathcal{O}(\epsilon)$, and extra$_2(i) \leq \mathcal{O}(\epsilon)$ for all i.*

Proof. First, we present the construction. For every pair of adjacent disks i and $i+1$ we describe a procedure to generate points on their borders. Then we set S_i to be the union of the generated points on the border of disk i when running the procedure on disks $(i, i+1)$ and on disks $(i-1, i)$. Finally, we show that extra$_1(i)$ and extra$_2(i)$ are sufficiently small for all i for our choice of S_i.

Procedure. Reorient the plane that $c_i = (0, y)$ and $c_{i+1} = (0, -y)$ for some $y > 1$. Let *spacing*: $\mathbb{R}_{\geq 0} \to \mathbb{R}_{>0}$ be a function that is nonincreasing in $|\phi|$ that we will define later. Given *spacing*, we use the following process to add points to S_i (and symmetrically for S_{i+1}):

1. Set $\phi = 0$.
2. While $\phi \leq \pi$:
   - Add $(\sin\phi, y - \cos\phi)$ to S_i.
   - ϕ += $\mathit{spacing}(\phi)$.
3. Repeat steps 1-2 but for ϕ from 0 to $-\pi$.

This procedure has the property that for any $\phi \in [-\pi, \pi]$, the point $(\sin\phi, y - \cos\phi)$ is within distance $\mathit{spacing}(|\phi|)$ of some point in S_i. In particular, if the optimal path has $p_i = (\sin\phi_i, y - \cos\phi_i)$ then it is guaranteed that $\|o_i\| \leq \mathit{spacing}(\phi_i)$. To compute $|S_i|$, note that as long as $\mathit{spacing}(\phi)$ is sufficiently smooth that $\frac{\mathit{spacing}(\phi)}{\mathit{spacing}(\phi + \mathit{spacing}(\phi))} = \Theta(1)$ for all ϕ, the number of points added to S_i will be at most a constant factor larger than the value of the definite integral $\int_{-\pi}^{\pi} \frac{1}{\mathit{spacing}(\phi)}\, d\phi$.

Next, we construct *spacing* so that $|S_i| = \mathcal{O}\left(\frac{1}{\sqrt{\epsilon}} \log \frac{1}{\epsilon}\right)$. Intuitively, by Example 21, we should have $\mathit{spacing}(\phi) = \Theta(\epsilon)$ closer to circle $i+1$ (when $\phi \approx 0$) and $\mathit{spacing}(\phi) = \Theta(\sqrt{\epsilon})$ farther from circle $i+1$ (when $\phi = \Theta(1)$). Thus, we set $\mathit{spacing}(\phi) = \max(\epsilon, \sqrt{\epsilon}\phi)$. The total number of added points is on the order of:

$$\int_0^{\pi} \frac{1}{\mathit{spacing}(\phi)}\, d\phi = \frac{1}{\sqrt{\epsilon}} \left(\int_0^{\sqrt{\epsilon}} \frac{1}{\sqrt{\epsilon}}\, d\phi + \int_{\sqrt{\epsilon}}^{\pi} \frac{1}{\phi}\, d\phi \right)$$
$$= \frac{1}{\sqrt{\epsilon}} \left(1 + \log\left(\frac{\pi}{\sqrt{\epsilon}}\right)\right) \leq \mathcal{O}\left(\frac{1}{\sqrt{\epsilon}} \log \frac{1}{\epsilon}\right).$$

Finally, we show that both extra terms are small for our choice of S_i.

Part 1: extra$_1(i)$. We note that $\mathrm{unit}(p_i - p_{i-1}) - \mathrm{unit}(p_{i+1} - p_i)$ must be parallel to $p_i - c_i$ for an optimal solution p. To verify this, it suffices to check the two cases from Lemma 20:

1. The points p_{i-1}, p_i, p_{i+1} are collinear, in which case $\mathrm{unit}(p_i - p_{i-1}) - \mathrm{unit}(p_{i+1} - p_i) = 0$.
2. The path reflects perfectly off circle i, in which case $\mathrm{unit}(p_i - p_{i-1}) - \mathrm{unit}(p_{i+1} - p_i)$ is parallel to $p_i - c_i$.

If we ensure that $\mathrm{spacing}(\phi) \leq \sqrt{\epsilon}$ for all ϕ, then $|o_i \cdot \mathrm{unit}(p_i - c_i)| \leq \epsilon$ because o_i is always nearly tangent to the circle centered at c_i at point p_i. The conclusion follows because $\mathrm{extra}_1(i) \leq 2|o_i \cdot \mathrm{unit}(p_i - c_i)| \leq 2\epsilon$.

Part 2: extra$_2(i)$. We upper bound $\text{extra}_2(i)$ by the sum of two summands, the first associated only with o_i and the second associated only with o_{i+1}.

▷ Claim 24. Letting $\text{ycoord}(\cdot)$ denote the y-coordinate of a point,

$$\text{extra}_2(i) \leq 2 \cdot \left(\min\left(\|o_i\|, \frac{4\|o_i\|^2}{\text{ycoord}(p_i)}\right) + \min\left(\|o_{i+1}\|, \frac{4\|o_{i+1}\|^2}{-\text{ycoord}(p_{i+1})}\right)\right).$$

Proof. We do casework based on which term is smaller on each of the mins.

1. $\|o_i\| \geq \frac{\text{ycoord}(p_i)}{4}$, $\|o_{i+1}\| \geq \frac{-\text{ycoord}(p_{i+1})}{4}$
   The result, $\text{extra}_2(i) \leq 2(\|o_i\|+\|o_{i+1}\|)$, follows by summing the following two inequalities:

$$\begin{aligned}&\sqrt{a_i^2 + [(o_{i+1}-o_i)\cdot \text{unit}(p_{i+1}-p_i)^\perp]^2} - \|p_{i+1}-p_i\| \\ &= \|p_{i+1}-p_i+o_{i+1}-o_i\| - \|p_{i+1}-p_i\| \leq \|o_i\| + \|o_{i+1}\|\end{aligned}$$

   and $\|p_{i+1}-p_i\| - a_i \leq \|o_i\| + \|o_{i+1}\|$.
2. $\|o_i\| \leq \frac{\text{ycoord}(p_i)}{4}$, $\|o_{i+1}\| \leq \frac{-\text{ycoord}(p_{i+1})}{4}$
   Then $\|o_i\|, \|o_{i+1}\| \leq \frac{\|p_{i+1}-p_i\|}{4}$ so $a_i \geq \frac{\|p_{i+1}-p_i\|}{2}$, and

$$\begin{aligned}\text{extra}_2(i) &\leq \frac{\|o_{i+1}-o_i\|^2}{2a_i} \leq \frac{2(\|o_{i+1}\|^2 + \|o_i\|^2)}{2a_i} \\ &\leq 2\cdot\frac{\|o_{i+1}\|^2 + \|o_i\|^2}{\|p_i - p_{i+1}\|} \leq 2\cdot\left(\frac{\|o_i\|^2}{\text{ycoord}(p_i)} + \frac{\|o_{i+1}\|^2}{-\text{ycoord}(p_{i+1})}\right).\end{aligned}$$

3. $\|o_i\| \leq \frac{\text{ycoord}(p_i)}{4}$, $\|o_{i+1}\| \geq \frac{-\text{ycoord}(p_{i+1})}{4}$
   Define $\text{extra}'(i)$ to be the same as $\text{extra}_2(i)$ with o_{i+1} set to 0. Then

$$\begin{aligned}\text{extra}'(i) &\triangleq \|p_{i+1}-p_i-o_i\| - (\|p_{i+1}-p_i\| - o_i\cdot\text{unit}(p_{i+1}-p_i)) \\ &= \sqrt{(\|p_{i+1}-p_i\| - o_i\cdot\text{unit}(p_{i+1}-p_i))^2 + [o_i\cdot\text{unit}(p_{i+1}-p_i)^\perp]^2} \\ &\quad - (\|p_{i+1}-p_i\| - o_i\cdot\text{unit}(p_{i+1}-p_i)) \\ &\leq \frac{\|o_i\|^2}{2\cdot\frac{3}{4}\|p_i-p_{i+1}\|} \leq \frac{\|o_i\|^2}{2\cdot\frac{3}{4}\cdot\text{ycoord}(p_i)}\end{aligned}$$

   and by similar reasoning as case 1, $\text{extra}_2(i) - \text{extra}'(i) \leq 2\|o_{i+1}\|$.
4. $\|o_i\| \geq \frac{\text{ycoord}(p_i)}{4}$, $\|o_{i+1}\| \leq \frac{-\text{ycoord}(p_{i+1})}{4}$
   Similar to case 3. ◀

Now that we have a claim showing an upper bound on $\text{extra}_2(i)$, it remains to show that $\min\left(\|o_i\|, \frac{\|o_i\|^2}{\text{ycoord}(p_i)}\right) \leq \mathcal{O}(\epsilon)$ for our choice of *spacing*. Indeed, when $\phi \leq \sqrt{\epsilon}$ we have $\|o_i\| \leq spacing(\phi) \leq \epsilon$, while for $\phi > \sqrt{\epsilon}$ we have $\frac{\|o_i\|^2}{\text{ycoord}(p_i)} \leq \mathcal{O}\left(\frac{spacing(\phi)}{\phi^2}\right) \leq \mathcal{O}(\epsilon)$. ◀

With small modifications to the proof of Lemma 23, we have the following corollary:

▶ **Corollary 25.** *Consider the case of non-unit disks. If the ith disk has radius r_i, then we can place $\mathcal{O}\left(\frac{1}{\sqrt{\epsilon_i}}\log\frac{1}{\epsilon_i}\right)$ points on its border such that the additive error associated with c_i —specifically, $extra_1(i)$ plus the components of $extra_2(i-1)$ and $extra_2(i)$ associated with $\|o_i\|$ —is $\mathcal{O}(r_i\epsilon_i)$. Consequently, $OPT + \sum_{i=1}^n extra_1(i) + \sum_{i=0}^n extra_2(i) \leq OPT + \sum_{i=1}^n r_i\epsilon_i$.*

Now, we finally prove Theorems 18 and 19.

Proof of Theorem 18 (Non-Unit Disks). We first present a slightly weaker result, and then show how to improve it. Recall that by Corollary 12, the ith largest disk has radius $\mathcal{O}\left(\frac{OPT}{i}\right)$ for $i \geq 3$. So if we set $\epsilon_i = \epsilon' = \frac{\epsilon}{\log n}$ for each of the ith largest disks for $i \geq 3$, the total additive error contributed by these disks becomes

$$\mathcal{O}\left(\sum_{i=3}^{n} \frac{OPT}{i} \cdot \epsilon_i\right) \leq \mathcal{O}\left(OPT \cdot \epsilon' \cdot \sum_{i=3}^{n} \frac{1}{i}\right) \leq \mathcal{O}\left(\epsilon OPT\right)$$

by Corollary 25. For the two largest disks, we use the previous naive discretization (placing $\mathcal{O}\left(\frac{1}{\epsilon}\right)$ points uniformly on the intersection of the circles with a square of side length $\mathcal{O}\left(OPT\right)$ centered about the starting point). We may assume we have already computed a constant approximation to OPT in $\mathcal{O}\left(n\right)$ time by applying Theorem 16 with $\epsilon = 1$. After selecting the point sets, we can finish with Corollary 7. The overall time complexity is $\mathcal{O}\left(\frac{n}{\sqrt{\epsilon'}} \log \frac{1}{\epsilon'} + \frac{1}{\epsilon}\right) \leq \mathcal{O}\left(\frac{n\sqrt{\log n}}{\sqrt{\epsilon}} \log\left(\frac{\log n}{\epsilon}\right) + \frac{1}{\epsilon}\right)$.

We can remove the factors of $\log n$ by selecting the ϵ_i to be an increasing sequence. Set $\epsilon_i = \Theta\left(\frac{\epsilon i^{2/3}}{n^{2/3}}\right)$ for each $i \in [3, n]$ such that more points are placed on larger disks. Then the total added error remains

$$\mathcal{O}\left(OPT \cdot \left(\epsilon + \sum_{i=3}^{n} \frac{\epsilon_i}{i}\right)\right) = \mathcal{O}\left(OPT \cdot \left(\epsilon + \sum_{i=3}^{n} \frac{1}{i} \cdot \frac{\epsilon i^{2/3}}{n^{2/3}}\right)\right)$$
$$= \mathcal{O}\left(OPT\epsilon \cdot \left(1 + n^{-2/3} \cdot \sum_{i=3}^{n} i^{-1/3}\right)\right) \leq \mathcal{O}\left(OPT\epsilon\right),$$

and the factors involving $\log n$ drop out from the time complexity:

$$\mathcal{O}\left(\sum_{i=3}^{n} \frac{1}{\sqrt{\epsilon_i}} \log\left(\frac{1}{\epsilon_i}\right) + \frac{1}{\epsilon}\right) \leq \mathcal{O}\left(\int_{i=3}^{n} \frac{1}{\sqrt{\epsilon}} n^{1/3} i^{-1/3} \log\left(\frac{n^{2/3}}{i^{2/3}\epsilon}\right) di + \frac{1}{\epsilon}\right)$$
$$\leq \mathcal{O}\left(\frac{3n^{1/3}}{2\sqrt{\epsilon}} i^{2/3} \left(\log \frac{n^{2/3}}{i^{2/3}\epsilon} + 1\right)\Bigg|_3^n + \frac{1}{\epsilon}\right)$$
$$\leq \mathcal{O}\left(\frac{n}{\sqrt{\epsilon}} \log\left(\frac{1}{\epsilon}\right) + \frac{1}{\epsilon}\right). \qquad \blacktriangleleft$$

To extend to multiple dimensions, we generalize the construction from Lemma 23.

Proof of Theorem 19 (Balls). As in Lemma 23, set $spacing(\phi) = \max(\epsilon, \sqrt{\epsilon}\phi)$ for a point p_i satisfying $m\angle p_i c_i c_{i+1} = \phi$, meaning that there must exist $p_i' \in S_i$ satisfying $\|p_i - p_i'\| \leq r_i \cdot spacing(\phi)$. The total number of points $|S_i|$ placed on the surface of a d-dimensional sphere is proportional to

$$\int_0^\pi \frac{\sin^{d-2}(\phi)}{spacing(\phi)^{d-1}}\, d\phi \leq \frac{1}{(\sqrt{\epsilon})^{d-1}} \int_0^\pi \frac{\phi^{d-2}}{\max(\sqrt{\epsilon}, \phi)^{d-1}}\, d\phi$$
$$= \frac{1}{\epsilon^{(d-1)/2}} \left(\int_0^{\sqrt{\epsilon}} \frac{\phi^{d-2}}{(\sqrt{\epsilon})^{d-1}}\, d\phi + \int_0^{\sqrt{\epsilon}} \frac{1}{\phi}\, d\phi\right) \leq \mathcal{O}\left(\frac{1}{\epsilon^{(d-1)/2}} \log \frac{1}{\epsilon}\right),$$

where the derivation of the integration factor $\sin^{d-2}(\phi)$ can be found in [18].

It remains to describe how to space points so that they satisfy the given spacing function. For each spacing $s = \epsilon, 2\epsilon, 4\epsilon, \ldots, \sqrt{\epsilon}$, we can find a d-dimensional hypercube of side length

$O(s/\sqrt{\epsilon})$ that encloses all points on the hypersphere with required spacing at most $2s$. Evenly space points with spacing s across the surface of this hypercube according to Lemma 4, and project each of these points onto the hypersphere. There are a total of $\mathcal{O}\left(\log \frac{1}{\epsilon}\right)$ values of s, and each s results in $\mathcal{O}\left(\frac{1}{\epsilon^{(d-1)/2}}\right)$ points being projected onto the hypersphere, for a total of $\mathcal{O}\left(\frac{1}{\epsilon^{(d-1)/2}} \log \frac{1}{\epsilon}\right)$ points. ◀

References

1. Xxii open cup, grand prix of siberia, 2022. Last accessed 9 April 2022. URL: `https://codeforces.com/blog/entry/96710`.
2. Alok Aggarwal and Maria Klawe. Applications of generalized matrix searching to geometric algorithms. *Discrete Applied Mathematics*, 27(1-2):3–23, 1990.
3. Arash Ahadi, Amirhossein Mozafari, and Alireza Zarei. Touring a sequence of disjoint polygons: Complexity and extension. *Theoretical Computer Science*, 556:45–54, 2014.
4. Tetsuo Asano, David Kirkpatrick, and Chee Yap. Pseudo approximation algorithms with applications to optimal motion planning. *Discrete & Computational Geometry*, 31(1):139–171, 2004.
5. Heinz H. Bauschke and Jonathan M. Borwein. On projection algorithms for solving convex feasibility problems. *SIAM Review*, 38(3):367–426, 1996. `doi:10.1137/S0036144593251710`.
6. Timothy M Chan. (near-) linear-time randomized algorithms for row minima in monge partial matrices and related problems. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1465–1482. SIAM, 2021.
7. Chang-Chien Chou. An exact algorithm for computing the shortest path touring n circles in 2d. In *International Conference on Computational Problem-Solving*, pages 98–103, 2010.
8. P. L. Combettes and H. J. Trussell. Method of successive projections for finding a common point of sets in metric spaces, December 1990. `doi:10.1007/bf00939646`.
9. Moshe Dror, Alon Efrat, Anna Lubiw, and Joseph S. B. Mitchell. Touring a sequence of polygons. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '03, pages 473–482, New York, NY, USA, 2003. Association for Computing Machinery. `doi:10.1145/780542.780612`.
10. Alon Efrat, Günter Rote, and Micha Sharir. On the union of fat wedges and separating a collection of segments by a line. *Comput. Geom. Theory Appl.*, 3(5):277–288, November 1993. `doi:10.1016/0925-7721(93)90018-2`.
11. Peter Gritzmann and Victor Klee. *Computational convexity*. CRC Press, Taylor; Francis Group, a Chapman; Hall Book, 2017.
12. Matthew J. Katz. 3-d vertical ray shooting and 2-d point enclosure, range searching, and arc shooting amidst convex fat objects. *Computational Geometry*, 8(6):299–316, 1997. `doi:10.1016/S0925-7721(96)00027-2`.
13. Amirhossein Mozafari and Alireza Zarei. Touring polygons: An approximation algorithm. In *IWOCA*, 2012.
14. Mark H. Overmars. Point location in fat subdivisions. *Information Processing Letters*, 44(5):261–265, 1992. `doi:10.1016/0020-0190(92)90211-D`.
15. Mark H. Overmars and A.F. van der Stappen. Range searching and point location among fat objects. *J. Algorithms*, 21:629–656, 1996.
16. R. R. Phelps. Convex sets and nearest points. *Proceedings of the American Mathematical Society*, 8(4):790–797, 1957. URL: `http://www.jstor.org/stable/2033300`.
17. Valentin Polishchuk and Joseph SB Mitchell. Touring convex bodies-a conic programming solution. In *CCCG*, pages 290–293, 2005.
18. David J. Smith and Mavina K. Vamanamurthy. How small is a unit ball? *Mathematics Magazine*, 62(2):101–107, 1989. URL: `http://www.jstor.org/stable/2690391`.

19 Xuehou Tan and Bo Jiang. Efficient algorithms for touring a sequence of convex polygons and related problems. In T.V. Gopal, Gerhard Jäger, and Silvia Steila, editors, *Theory and Applications of Models of Computation*, pages 614–627, Cham, 2017. Springer International Publishing.