

Going Around in Circles

Adrian Dumitrescu*

December 20, 2011

Abstract

Let $\varepsilon > 0$ and let Ω be a disk of sufficiently large radius R in the plane, *i. e.*, $R \geq R(\varepsilon)$. We first show that the set of lattice points inside Ω can be connected by a (possibly self-intersecting) spanning tour (Hamiltonian cycle) consisting of straight line edges such that the turning angle at each point on the tour is at most ε . This statement remains true for any large and evenly distributed point set (suitably defined) in a disk. This is the first result of this kind that suggests far-reaching generalizations to arbitrary regions with a smooth boundary. Our methods are constructive and lead to an efficient algorithm for computing such a tour. On the other hand, it is shown that such a result does not hold for convex regions without a smooth boundary.

Keywords: Hamiltonian cycle, turning angle, robot path planning, geometric graph, integer lattice.

1 Introduction

A *spanning tour* (spanning path) on n points is a directed Hamiltonian cycle (Hamiltonian path, respectively), drawn with straight line edges. In the Euclidean traveling salesman problem (TSP), given a set of points in the plane, one seeks a shortest spanning tour. Particularly in the last decade, there has been an increased interest in studying tours that optimize objective functions related to angles between consecutive edges in the tour, rather than the length. The problem has applications in motion planning, where restrictions on turning angles have to be enforced. For example, an aircraft or a boat moving at high speed, required to pass through a set of given locations, cannot make sharp turns in its motion. This and other applications to planning curvature-constrained paths for auto-vehicles and aircraft are discussed in [2, 3, 8, 11, 12].

Consider a spanning tour (or path) on a set of $n \geq 2$ points. When three consecutive points, p_1 , p_2 , and p_3 , are traversed in this order, the *turning angle* at p_2 , denoted by $\text{turn}(p_1, p_2, p_3)$, is the supplement of the angle in $[0, \pi]$ determined by the segments p_2p_1 and p_2p_3 ; observe that $\text{turn}(p_1, p_2, p_3) \in [0, \pi]$. If p_3 is on the left (resp. right) side of the oriented line $\overrightarrow{p_1p_2}$, we say that the tour (or path) makes a *left* (resp. *right*) *turn* at p_2 . If all of its turning angles are at least $\pi/2$, we call it an *acute* tour (or path). If all turning angles are at most $\pi/2$, the tour (or path) is *obtuse*; see Figure 1.

Fekete and Woeginger [10] proved that every n -element point set S admits an acute spanning path. It is easy to see that in some cases, such a path cannot be completed to an acute *tour*: indeed, if all points are on a line and n is *odd*, then along any (spanning) tour, one of the turning angles must be equal to 0.

*Department of Computer Science, University of Wisconsin–Milwaukee, Email: dumitres@uwm.edu. Supported in part by NSF CAREER grant CCF-0444188 and NSF grant DMS-1001667.




Figure 1: Left: A left turn at p_2 and its turning angle θ . Middle: An acute tour on 6 points. Right: An obtuse tour on 9 points.

Most desirable in applications, particularly in motion planning, are spanning paths or tours from the other extreme, that we study here, namely where *each turning angle is small*. A rough approximation is provided by paths or tours that are obtuse. However not all point sets admit obtuse tours or even obtuse paths. For instance, some point sets require turning angles at least as large as $5\pi/6$ in any spanning path [10]: the 3 vertices of an equilateral triangle and its center make a small 4-point example. Moreover, other point sets (*e. g.*, collinear) require the maximum turning angle possible, namely π , in any spanning cycle.

Related problems and results. Aggarwal *et al.* [2] studied the following variant of TSP. Given n points in the plane, compute a Hamiltonian tour of the points that minimizes the total turning angle. The total turning angle of a tour is the sum of the turning angles at each of the n points. They proved that this problem is NP-hard. On the positive side, they gave a polynomial-time algorithm with approximation ratio $O(\log n)$.

Various angle conditions imposed on *geometric graphs* (graphs with straight-line edges) drawn on a given set of points have been studied in [4, 5, 6, 7]. For instance, Bárány, Pór, and Valtr [7] proved that any point set admits a (possibly self-intersecting) Hamiltonian *path* in which each turning angle is at most $8\pi/9$. Fekete and Woeginger [10] had conjectured earlier that this holds for turning angles at most $5\pi/6$, and showed that no smaller value would do. Dumitrescu *et al.* [9] have recently shown that every set S of n points in the plane ($n \geq 4$, even) admits a (possibly self-intersecting) Hamiltonian *cycle* consisting of n straight line edges such that the turning angle at each point on the cycle is at least $\pi/3$.

Aichholzer *et al.* [4] studied similar questions for *planar* geometric graphs. For instance, they showed that any point set in general position in the plane admits a non-intersecting Hamiltonian (spanning) path with the property that each turning angle is at least $\pi/4$. They also conjectured that this value can be replaced by $\pi/2$. Arkin *et al.* [6] introduced the notion of *reflexivity* of a point set, as the minimum number of reflex vertices in a polygonization (*i. e.*, simple polygon) of the set. They gave estimates for the maximum reflexivity of an n -element point set. Recently, Ackerman *et al.* [1] made further progress on this problem. Other variants of Euclidean TSP can be found in the survey article by Mitchell [14].

Our results.

Theorem 1. *For any $\varepsilon > 0$, there exists $R(\varepsilon) = O(\varepsilon^{-3})$ with the following property. Let Ω be a disk of radius $R \geq R(\varepsilon)$ in the plane. Then the set of lattice points $\mathbb{Z}^2 \cap \Omega$ admits a Hamiltonian tour with each turning angle at most ε . Such a tour can be computed in $O(|\mathbb{Z}^2 \cap \Omega|)$ time.*

This statement remains true for any large and evenly distributed point set (suitably defined) in a disk; see Theorem 3 below. This shows that the result is not just an artifact of some grid

structure. The grid point example illustrates the ideas of the proof in a cleaner and simpler way, while for applications, the more general Theorem 3 concerning “arbitrary” even-distributed point sets is much more relevant. Further extensions are discussed below. Theorems 1 and 2 suggest far-reaching extensions for regions with a smooth boundary: Conjectures 1 and 2 below. On the other hand, it is obvious that the result does not hold for small (lattice) point sets. Theorem 2 shows that one cannot expect such a result either for convex regions without a smooth boundary.

We think that the statement in Theorem 1 holds for two reasons: (1) the lattice points are evenly distributed in the disk, and (2) the disk has a smooth boundary. We next make these notions precise.

We say that a region, X , bounded by a closed Jordan curve, has a *smooth boundary*, if there is a unique tangent to the curve at each boundary point. Obviously the disk is the simplest region with a smooth boundary, and that is why we prove Theorem 1 for disks. Theorem 2 below shows that the smooth boundary condition is necessary for such a result to hold for convex regions, irrespective of the size of the point set.

Theorem 2. *Let X be a convex region without a smooth boundary in the plane. Then there exist $\varepsilon_0 > 0$, and $\lambda_0 > 0$, with the following property: For any $\lambda \geq \lambda_0$, and any similar region λX of X , each Hamiltonian cycle on the set $\mathbb{Z}^2 \cap \lambda X$ has a turning angle larger than ε_0 .*

Consider a fixed (possibly disconnected) region X in the plane, bounded by finitely many closed smooth pairwise-disjoint (boundary) curves. For brevity, let us refer to this type of region as *bounded region with a smooth boundary*. Let $\rho > 0$, and $c \geq 1$. Let S be a finite set of points in X . We say that S is (ρ, c) -*evenly distributed* in X if (i) the disks of radius ρ centered at the points in S are pairwise interior-disjoint, and (ii) the disks of radius $c\rho$ centered at the points in S collectively cover X . For instance, the set of lattice points $\mathbb{Z}^2 \cap \Omega$ in a disk Ω is $(1/2, \sqrt{2})$ -evenly distributed in Ω . Our result in Theorem 1 remains true for any large and evenly distributed point set in a disk.

Theorem 3. *Let $\rho > 0$, and $c \geq 1$ be two constants. For any $\varepsilon > 0$, there exists $R(\rho, c, \varepsilon) = O(\varepsilon^{-3})$ with the following property. Let Ω be a disk of radius $R \geq R(\rho, c, \varepsilon)$ in the plane. Let S be a (ρ, c) -evenly distributed point set in Ω . Then S admits a Hamiltonian tour with each turning angle at most ε . Such a tour can be computed in $O(|S|)$ time.*

Let X be a bounded region with a smooth boundary in the plane. We believe that the statement in Theorem 3 holds for any evenly distributed point set contained in any sufficiently large similar copy λX of X , and likewise for large *random* point sets, uniformly selected from X :

Conjecture 1. *Let $\rho > 0$, and $c \geq 1$ be two constants, and X be a bounded region with a smooth boundary in the plane. Let S be a (ρ, c) -evenly distributed point set in a similar copy λX of X . Then, for any $\varepsilon > 0$, the point set S admits a Hamiltonian tour with each turning angle at most ε , provided that λ is sufficiently large.*

Conjecture 2. *Let X be a bounded region with a smooth boundary in the plane, and let S be a set of n points, randomly and uniformly selected from X . Then, for any $\varepsilon > 0$, the point set S almost surely admits a Hamiltonian tour with each turning angle at most ε , as n tends to infinity.*

From the proof of Theorem 2 we obtain the following.

Corollary 1. *Let $P = p_1 \dots p_n$ be a convex polygon with angles $\beta_i = \pi - \alpha_i$, where $\alpha_i \in (0, \pi)$ for $i = 1, \dots, n$. Let $\alpha = \max\{\alpha_i \mid i = 1, \dots, n\}$. Then there exists $\lambda_0 > 0$, with the following property: For any $\lambda \geq \lambda_0$, and any similar region λP of P , each Hamiltonian cycle on the set $\mathbb{Z}^2 \cap \lambda P$ has some turning angle that is $\Omega(\alpha)$.*

We think that this lower bound is asymptotically tight:

Conjecture 3. *Let $P = p_1 \dots p_n$ be a convex polygon with angles $\beta_i = \pi - \alpha_i$, where $\alpha_i \in (0, \pi)$ for $i = 1, \dots, n$. Let $\alpha = \max\{\alpha_i \mid i = 1, \dots, n\}$. Then there exist $\lambda_0 > 0$ and $c \geq 1$, with the following property: For any $\lambda \geq \lambda_0$, and any similar region λP of P , the set $\mathbb{Z}^2 \cap \lambda P$ can be traversed by a Hamiltonian cycle with maximum turning angle at most $c\alpha$.*

Remarks. By placing a vertex of λP which attains the minimum angle, say β_i , at a lattice point forces a turning angle at least $\alpha_i = \pi - \beta_i$. This shows that if Conjecture 3 holds, the inequality $c \geq 1$ is needed. Obviously any turning angle is at most π , so the interest in Conjecture 3 is for α close to zero (unbounded from below). Observe also that $\sum_{i=1}^n \alpha_i = 2\pi$ which implies $\alpha \geq 2\pi/n$. In particular, for the regular n -gon, $\alpha = 2\pi/n$, and the question is whether the set $\mathbb{Z}^2 \cap \lambda P$ can be traversed by a Hamiltonian cycle with maximum turning angle $O(1/n)$, provided that λ is sufficiently large. By Corollary 1, the maximum turning angle is $\Omega(1/n)$, regardless of how λP is placed.

Definitions and notations. For a positive integer m , let $[m] = \{1, 2, \dots, m\}$. A lattice point (i, j) is *special* if both i and j are congruent to 0 modulo 4. For a bounded region R in the plane, let $\text{Area}(X)$ and $\text{per}(X)$ denote the area and respectively, the perimeter, of X . Let ∂X denote the boundary of X . We will refer to a *circular annulus* also as a *circular ring*, or simply *ring*, and to an *annular sector* also as a *ring sector*. A ring sector will be also called *block*. The *length* of a block with inner and outer radii r_1, r_2 and center angle β is $\frac{r_1+r_2}{2}\beta$. When there is no danger of confusion, a block may also refer to the actual set of lattice points within.

Given a fixed (possibly disconnected) region X in the plane, bounded by finitely many closed curves, and $\lambda > 0$, denote by λX any *similar* copy of X , that is, a possibly rotated copy of X that is scaled by a factor of λ .

2 Tour construction: proofs of Theorems 1 and 3

We are to make a Hamiltonian cycle for the set of lattice points enclosed in a large disk. The main idea is to follow a spiral path and go around in very thin rings of large radius which forces turning angles to be small. However a special plan is needed to visit the points near the disk center, where the rings have a small radius, and thus would be unsuitable. We also need to ensure that the constructed path reconnects smoothly to the start point.

2.1 Proof of Theorem 1

We first describe the cycle construction. The rest of the proof (correctness and analysis) is divided into a sequence of five lemmas. We will assume w.l.o.g. that $\varepsilon \leq \frac{1}{4}$.

Parameters. We will use a set of interdependent parameters, that we collect here for easy reference; recall that R is the radius of Ω , and $R \geq R(\varepsilon)$ is assumed:

- $R(\varepsilon) = 10^9 \varepsilon^{-3}$,
- $k = \lceil \frac{100\pi}{\varepsilon} \rceil$, $\beta = \frac{2\pi}{2k+1}$; in particular, $\beta \leq \frac{\varepsilon}{100}$.
- $r_0 = \frac{R\beta}{120}$,

- $v = \lfloor \frac{R-r_0}{20} \rfloor$, $w = \frac{R-r_0}{v}$; this setting implies that $20 \leq w \leq 21$, as shown below.
- $u = \lceil \frac{2r_0}{w} \rceil$.

Since $R \geq 500$, we have

$$R - r_0 = R \left(1 - \frac{\beta}{120}\right) \geq 500 \left(1 - \frac{1}{1000}\right) \geq 450,$$

$$\text{hence } v = \left\lfloor \frac{R - r_0}{20} \right\rfloor \geq \left\lfloor \frac{450}{20} \right\rfloor = 22.$$

Consequently,

$$w = \frac{R - r_0}{v} = \frac{R - r_0}{20} \cdot \frac{20}{v} \leq (v + 1) \cdot \frac{20}{v} = 20 + \frac{20}{v} \leq 21.$$

Elements of the construction. Let Ω and ω be two concentric disks centered at o of radius R and r_0 , respectively. Denote by ℓ the horizontal line through o , by A and D the two intersection points of $\partial\Omega$ with ℓ , and by B and C the two intersection points of $\partial\omega$ with ℓ . (A, B, C, D are ordered from left to right on ℓ .) Partition $\Omega \setminus \omega$ into v concentric rings (annuli) $\Gamma_1, \dots, \Gamma_v$, of equal width w , with common boundary points being assigned arbitrarily to one of the two rings involved. Refer to Fig. 2. The rings $\Gamma_1, \dots, \Gamma_v$ are ordered by increasing (inner) radii. The inner and outer radii of Γ_i are $r_0 + (i - 1)w$, and respectively, $r_0 + iw$, for $i = 1, 2, \dots, v$. Partition each of the concentric rings $\Gamma_1, \dots, \Gamma_v$ into $2k + 1$ ring sectors (blocks) of equal center angle β , where this partition is conforming with ℓ : that is, the blocks of each ring are ordered counterclockwise from ℓ , with the first and the last blocks separated by ℓ , and with the first block above ℓ . The blocks of Γ_i are denoted $\Gamma_{i,1}, \dots, \Gamma_{i,2k+1}$.

The idea behind using blocks is as follows: the block length must be large enough compared to the width to allow small turning angles when skipping a block, or connecting to an adjacent ring. And for the same reason, the block length must be small enough compared to the average radius of the ring; that is, its center angle must be small.

Consider the disk Ω_1 of radius $(R - 3r_0)/2$, and the disk Ω_2 of radius $(R + r_0)/2$, both centered at the midpoint of BD . Let Λ denote the ring of width w and outer radius $(R - r_0)/2$ centered at the midpoint of AB . Cover the ring $\Omega_2 \setminus \Omega_1$ by u concentric rings Φ_1, \dots, Φ_u , of equal width w ; recall that $u = \lceil 2r_0/w \rceil$. The rings Φ_1, \dots, Φ_u are ordered by increasing (inner) radii. Observe that these u rings completely cover ω , and also that Λ is tangent to Φ_u . Partition each of the concentric rings Φ_1, \dots, Φ_u into $2k + 1$ blocks of equal center angle β , where this partition is also conforming with ℓ : its blocks are ordered counterclockwise from ℓ , with the first and the last blocks separated by ℓ , and with the first block above ℓ . The blocks of Φ_i are denoted $\Phi_{i,1}, \dots, \Phi_{i,2k+1}$. Similarly partition Λ into $2k + 1$ blocks of equal center angle β , labeled in the same convention, by $\Lambda_1, \dots, \Lambda_{2k+1}$. Note that, since $\beta R/3 \geq 2r_0$, it follows that $\Phi_i \cap \omega \subset \Phi_{i,k+1}$, for each $i \in [u]$.

The paths P_1 and P_2 . We construct a Hamiltonian cycle by concatenating two paths: P_1 and P_2 . P_1 covers all lattice points in ω , while leaving most of the other points in $\Omega \setminus \omega$ untouched, via a clockwise outward spiral, and then connects to a point in Γ_v . With the exception of the points in ω , the other points in P_1 , chosen from $\Phi_1 \cup \dots \cup \Phi_u$ are sparse. More precisely, outside ω , P_1 traverses only *special* lattice points. P_2 traverses all points left in $\Gamma_1 \cup \dots \cup \Gamma_v$ via a counterclockwise inward spiral path, and then reconnects to P_1 . The basic elements used in constructing the cycle are *circular rings*, see Fig. 2.




Figure 2: Left: Concentric rings $\Gamma_1, \dots, \Gamma_v$, and smaller disk ω (darkly shaded). Right: Concentric smaller rings Φ_1, \dots, Φ_u superimposed, containing the paths covering ω . The ring Λ with the short path (in the upper half) for switching the direction is tangent from the left to Φ_u and to ω .

P_1 starts at a point in $\Phi_{1,k+1} \cap \omega$ above ℓ and goes upward and clockwise inside this inner ring for a full circle, returning to $\Phi_{1,k+1}$, by skipping every other block in the ring, or two blocks in the last step: if k is even, via blocks $k-1, k-3, \dots, 1, 2k, 2k-2, \dots, k+2$; if k is odd, via blocks $k, k-2, \dots, 1, 2k, 2k-2, \dots, k+3$. (If $\Phi_{1,k+1} \cap \omega$ is empty of lattice points, an arbitrary lattice point in $\Phi_{1,k+1}$ is selected as starting point). The path continues clockwise around Φ_1 , by visiting one new point from $\Phi_1 \cap \omega$ at each rotation, until all points in $\Phi_1 \cap \omega$ are traversed. Within the same ring, the path visits a point in every other block, and in certain cases skips two or three blocks. The path then connects to a point in $\Phi_2 \cap \omega$, and goes clockwise around Φ_2 , by visiting one new point from $\Phi_2 \cap \omega$ at each rotation, until all points in $\Phi_2 \cap \omega$ are traversed. Each of the subsequent rings Φ_2, \dots, Φ_u is repeatedly traversed until all points in $(\Omega_2 \setminus \Omega_1) \cap \omega$ are traversed. The turning angle at each connection point between a ring and the next concentric ring is at most ε by Lemma 1, below. For convenience (to simplify some later calculation), we will assume that the last point traversed from $(\Omega_2 \setminus \Omega_1) \cap \omega$, specifically from $\Phi_{u,k+1} \cap \omega$ lies below ℓ . This condition can be easily ensured.

Observe that for any $i \in [u]$, the number of *special* lattice points in any block $\Phi_{i,j}$ is at least $Rw\beta/2$, which is larger than the number of lattice points in $\Phi_i \cap \omega$, namely about $2wr_0$; so there is always an available point to extend the path P_1 with, in any desired block, as long as necessary.

From the last point traversed in $(\Omega_2 \setminus \Omega_1) \cap \omega$, the path P_1 moving upwards switches to the ring Λ tangent to the left and continues counterclockwise for a half-circle, until it reaches the largest ring Γ_v from the other family of rings in a block near point A . Here P_1 ends and P_2 starts. P_2 circles around counterclockwise in Γ_v until all points left in Γ_v are traversed. It then switches to the second largest ring Γ_{v-1} , continuing counterclockwise until all points left in Γ_{v-1} are traversed, and so on until the last ring Γ_1 . Generating the paths visiting the points in $\Gamma_v, \Gamma_{v-1}, \dots, \Gamma_1$ is done according to Lemma 5. The conditions in the lemma ensure that these paths can be linked together. Once all points left in $\Gamma_1 \cup \dots \cup \Gamma_v$ are traversed, P_2 closes the Hamiltonian cycle by reconnecting to the start point of P_1 .

Ensuring small turning angles. Next we bound from above the turning angles at the points on the tour. Then we describe the construction of the path P_2 in detail. The first lemma handles connections between two adjacent concentric rings, such as Γ_i, Γ_{i-1} , or Φ_j, Φ_{j+1} . Two such adjacent rings make a ring of width $2w$.

Lemma 1. *Let Γ be a circular ring with inner radius $r-2w$, outer radius r , and center o , partitioned by rays into congruent blocks of center angle β . Assume that $r_0 + 2w \leq r \leq R$. Let $p_1 \in B_1$, and $p_2 \in B_2$, be two points in two distinct blocks, where: (i) the block B_1 precedes the block B_2 in clockwise order, and (ii) B_1 and B_2 are separated by j other blocks, where $j \in \{1, 2, 3\}$. Let ℓ_i , $i = 1, 2$ be the two lines perpendicular to op_1 and op_2 , respectively. Then the angle between ℓ_i and p_1p_2 is at most $\varepsilon/2$, for $i = 1, 2$.*

Proof. By symmetry, it suffices to prove the bound for $i = 1$. Refer to Fig. 3. Let q_2 denote the




Figure 3: B_1 and B_2 are separated by two other blocks ($j = 2$). The angle between p_1p_2 and ℓ_1 is at most $\varepsilon/2$.

intersection point between op_2 and the inner circle of radius $r - 2w$. Denote by a the projection of p_2 onto ℓ_1 , and by b the intersection point between the extension of op_2 and the line ℓ_1 . If α denotes the angle between ℓ_1 and p_1p_2 , we have

$$\sin \alpha = \frac{|p_2a|}{|p_1p_2|} \leq \frac{|p_2b|}{|p_1p_2|} \leq \frac{|q_2b|}{|p_1p_2|}.$$

Recall that the points p_1 and p_2 do *not* lie in adjacent blocks, hence

$$|p_1p_2| \geq 2(r - 2w) \sin \frac{\beta}{2}.$$

Write $\gamma = \angle p_1op_2$. Since B_1 and B_2 are separated by at most three other blocks, it follows that $\gamma \leq 5\beta$. We have $|op_1| = |ob| \cos \gamma$, hence $|ob| = |op_1| / \cos \gamma$. The length $|q_2b|$ is bounded from above as follows:

$$|q_2b| = |ob| - |oq_2| = \frac{|op_1|}{\cos \gamma} - (r - 2w) \leq \frac{r}{\cos \gamma} - r + 2w = \frac{r}{\cos \gamma} (1 - \cos \gamma) + 2w.$$

By our choice of parameters, $\cos \gamma \geq 2/3$, hence

$$|q_2b| \leq \frac{3r}{2} (1 - \cos \gamma) + 2w = 3r \sin^2 \frac{\gamma}{2} + 2w.$$

By Jensen's inequality [13] (or see [15, p. 24]), $\gamma \leq 5\beta$ yields

$$\sin \frac{\gamma}{2} \leq \sin \frac{5\beta}{2} \leq 5 \sin \frac{\beta}{2}.$$

Putting these inequalities together yields

$$\sin \alpha \leq \frac{|q_2 b|}{|p_1 p_2|} \leq \frac{75r \sin^2 \frac{\beta}{2} + 2w}{2(r-2w) \sin \frac{\beta}{2}}.$$

By our choice of parameters,

$$\frac{r}{r-2w} \leq \frac{r_0}{r_0-2w} \leq 1.01,$$

hence

$$\frac{75r \sin^2 \frac{\beta}{2}}{2(r-2w) \sin \frac{\beta}{2}} \leq 38 \sin \frac{\beta}{2} \leq 38 \frac{\beta}{2} = 19\beta \leq \frac{19\varepsilon}{100},$$

and

$$\frac{w}{(r-2w) \sin \frac{\beta}{2}} \leq \frac{21}{r_0 \sin \frac{\beta}{2}} \leq \frac{21}{r_0 \frac{6\beta}{7}} = \frac{49}{r_0 \beta} = \frac{49 \cdot 120}{R\beta^2} \leq \frac{6000 \cdot 10^4 \cdot \varepsilon^3}{10^9 \varepsilon^2} = \frac{6 \cdot 10^7 \cdot \varepsilon^3}{10^9 \varepsilon^2} \leq \frac{6\varepsilon}{100}.$$

Consequently,

$$\sin \alpha \leq \frac{|q_2 b|}{|p_1 p_2|} \leq \frac{19\varepsilon}{100} + \frac{6\varepsilon}{100} = \frac{\varepsilon}{4}, \text{ hence } \alpha \leq \frac{\varepsilon}{2},$$

as required. \square

The second lemma handles connections between two tangent rings, as they occur in the cycle: from Φ_u to Λ , and from Λ to Γ_v .

Lemma 2. *Consider the two tangent rings Φ_u and Λ , centered at o_1 and o_2 , respectively. Refer to Fig. 4. Let $p_1 \in \Phi_{u,k+1}$, and $p_2 \in \Lambda_2$, where p_1 lies below ℓ . Let ℓ_i , $i = 1, 2$ be the two lines perpendicular to $o_1 p_1$ and $o_2 p_2$, respectively. Let $p_3 \in \Lambda_{k+1}$, and $p_4 \in \Gamma_{v,k+3}$, where p_3 lies above ℓ . Let ℓ_i , $i = 3, 4$ be the two lines perpendicular to $o_2 p_3$ and $o p_4$, respectively. Then: (i) the angle between ℓ_i and $p_1 p_2$ is at most $\varepsilon/2$, for $i = 1, 2$. (ii) the angle between ℓ_i and $p_3 p_4$ is at most $\varepsilon/2$, for $i = 3, 4$.*

Proof. (i). Denote by r_1 and r_2 the outer radii of Φ_u and Λ , respectively. By construction, we have $R/3 \leq r_1, r_2 \leq R$. Also by construction, the angle made by ℓ_1 with the vertical direction belongs to the interval $[0, \beta/2]$, and the angle made by ℓ_2 with the vertical direction belongs to the interval $[\beta, 2\beta]$. Let α_{12} be the angle made by $p_1 p_2$ with the vertical direction. This angle attains its maximum if p_1 lies on ℓ on the inner circle defining Φ_u , and p_2 lies at its lowest position in the second block of Λ on the ray from o_2 separating the first two blocks of Λ . For these placements we have: $r_2 \leq R/2$, and $r_2 - w \geq R/3$, hence

$$\tan \alpha_{12} = \frac{2w + (r_2 - w)(1 - \cos \beta)}{(r_2 - w) \sin \beta} \leq \frac{42 \cdot 3}{R \sin \beta} + \frac{2 \sin^2 \frac{\beta}{2}}{\sin \beta} \leq \frac{\varepsilon}{100} + \tan \frac{\beta}{2} \leq \frac{\varepsilon}{100} + \frac{\varepsilon}{100} \leq \frac{\varepsilon}{4},$$

hence $\alpha_{12} \leq \varepsilon/3$. It follows that the angle between ℓ_1 and $p_1 p_2$ is at most $\alpha_{12} + \beta/2 \leq \varepsilon/3 + \beta/2 \leq \varepsilon/2$, as claimed. We also get that the angle between ℓ_2 and $p_1 p_2$ is at most $\alpha_{12} + 2\beta \leq \varepsilon/3 + 2\beta \leq \varepsilon/2$, and this concludes the proof of part (i).




Figure 4: The angle between p_1p_2 and ℓ_i is at most $\varepsilon/2$, for $i = 1, 2$. The angle between ℓ_i and p_3p_4 is at most $\varepsilon/2$, for $i = 3, 4$.

(ii). By construction, the angle made by ℓ_3 with the vertical direction belongs to the interval $[0, \beta/2]$, and the angle made by ℓ_4 with the vertical direction belongs to the interval $[3\beta/2, 5\beta/2]$. Let α_{34} be the angle made by p_3p_4 with the vertical direction. This angle attains its maximum if $p_3 = A$ on the line ℓ , and p_4 lies at its lowest position in $\Gamma_{v,k+3}$ on the inner circle of Γ_v . For these placements a similar calculation gives:

$$\tan \alpha_{34} = \frac{w + (R - w)(1 - \cos \frac{5\beta}{2})}{(R - w) \sin \frac{5\beta}{2}} \leq \frac{\varepsilon}{100} + \tan \frac{5\beta}{4} \leq \frac{\varepsilon}{100} + \frac{3\beta}{2} \leq \frac{\varepsilon}{100} + \frac{3\varepsilon}{200} \leq \frac{\varepsilon}{4},$$

hence $\alpha_{34} \leq \varepsilon/3$. It follows that the angle between ℓ_3 and p_3p_4 is at most $\alpha_{34} + \beta/2 \leq \varepsilon/3 + \beta/2 \leq \varepsilon/2$, as claimed. We also get that the angle between ℓ_4 and p_3p_4 is at most $\alpha_{34} + 5\beta/2 \leq \varepsilon/3 + 5\beta/2 \leq \varepsilon/2$, and this concludes the proof of part (ii). \square

The third lemma handles the reconnection from the end of P_2 to the first point of P_1 , namely from Γ_1 to Φ_1 .

Lemma 3. *Let $p_1 \in \Phi_{1,k+1}$, and $p_2 \in \Gamma_{1,2k}$, where p_1 lies above ℓ . Let o_1 be the center of Φ_1 and o be the center of Γ_1 (and Ω). Let ℓ_i , $i = 1, 2$ be the two lines perpendicular to o_1p_1 and op_2 , respectively. Then the angle between ℓ_i and p_1p_2 is at most $\varepsilon/2$, for $i = 1, 2$.*

Proof. Observe first that the two rings intersect, *i. e.*, $\Phi_1 \cap \Gamma_1 \neq \emptyset$. By construction, the angle made by ℓ_1 with the vertical direction belongs to the interval $[0, \beta/2]$, and the angle made by ℓ_2 with the vertical direction belongs to the interval $[\beta, 2\beta]$. Let α_{12} be the angle made by p_1p_2 with the vertical direction. This angle attains its maximum if p_1 lies on the line ℓ at C , and p_2 lies at its lowest position in $\Gamma_{1,2k}$ on the inner circle of Γ_1 . For these placements the calculation gives:

$$\tan \alpha_{12} \leq \frac{r_0(1 - \cos 2\beta)}{r_0 \sin 2\beta} = \tan \beta \leq \frac{3\beta}{2} \leq \frac{3\varepsilon}{200} \leq \frac{\varepsilon}{4},$$

hence $\alpha_{34} \leq \varepsilon/3$. As in the previous proofs, it follows that the angle between ℓ_i and p_1p_2 is at most $\varepsilon/2$, for $i = 1, 2$. \square

Block sizes and smoothing them out. The next lemma estimates the number of lattice points in a block, and as corollary, the number of special lattice points in a block. These key elements are relevant for the construction of the path P_2 (details in Lemma 5) which deals with blocks of different but similar sizes.

Lemma 4. *Consider a ring Γ_i , $i \in [v]$ and its partition into congruent ring sectors (annular sectors) of angle β : $\Gamma_{i,j}$, with $j = 1, \dots, 2k + 1$. Let $B_j = \mathbb{Z}^2 \cap \Gamma_{i,j}$, for $j = 1, \dots, 2k + 1$. Then*

$$||B_j| - \text{Area}(\Gamma_{i,j})| \leq 0.15 \cdot \text{Area}(\Gamma_{i,j}). \quad (1)$$

Proof. Recall that the inner and outer radii of Γ_i are $r_0 + (i - 1)w$, and respectively, $r_0 + iw$, for $i = 1, 2, \dots, v$. We have

$$\text{Area}(\Gamma_{i,j}) = w \cdot r_{\text{avg}} \cdot \beta, \text{ and } \text{per}(\Gamma_{i,j}) = 2r_{\text{avg}} \cdot \beta + 2w,$$

where $r_{\text{avg}} = r_0 + (i - 0.5)w$ is the average of the inner and outer radii of Γ_i . Consequently

$$\frac{\text{per}(\Gamma_{i,j})}{\text{Area}(\Gamma_{i,j})} = \frac{2r_{\text{avg}} \cdot \beta + 2w}{w \cdot r_{\text{avg}} \cdot \beta} = \frac{2}{w} + \frac{2}{r_{\text{avg}} \cdot \beta} \leq \frac{2}{w} + \frac{2}{r_0 \cdot \beta} \leq \frac{2}{w} + \frac{1}{1000}. \quad (2)$$

Let $z_{i,j} = |\mathbb{Z}^2 \cap \Gamma_{i,j}|$. We need the following inequality which relates $z_{i,j}$ to the area and perimeter of $\Gamma_{i,j}$.

$$\text{Area}(\Gamma_{i,j}) - \sqrt{2}\text{per}(\Gamma_{i,j}) - 2\sqrt{2}\pi \leq z_{i,j} \leq \text{Area}(\Gamma_{i,j}) + \sqrt{2}\text{per}(\Gamma_{i,j}) + 2\sqrt{2}\pi. \quad (3)$$

Indeed, an enlarged copy of $\Gamma_{i,j}$ at distance $\sqrt{2}$ around its boundary contains all grid cells intersected by $\partial(\Gamma_{i,j})$; this implies the upper bound. On the other hand, all grid cells that intersect a shrunk copy of $\Gamma_{i,j}$ at distance $\sqrt{2}$ are contained in $\Gamma_{i,j}$; this implies the lower bound. Since $\sqrt{2}(2/w + 1/1000) < 0.15$, from the inequalities (2) and (3), we obtain

$$0.85 \cdot \text{Area}(\Gamma_{i,j}) \leq z_{i,j} \leq 1.15 \cdot \text{Area}(\Gamma_{i,j}),$$

as required. \square

Denote by $z'_{i,j}$ the number of special lattice points in $\mathbb{Z}^2 \cap \Gamma_{i,j}$. By the same argument used for estimating $z_{i,j} = |\mathbb{Z}^2 \cap \Gamma_{i,j}|$ in the proof of Lemma 4, we have

$$z'_{i,j} \leq 1.15 \cdot \frac{\text{Area}(\Gamma_{i,j})}{16} \leq 0.071 \cdot \text{Area}(\Gamma_{i,j}). \quad (4)$$

Assume now that the path P_1 covering the points in ω has been generated; recall that outside ω , P_1 visits only *special* lattice points. Consider a ring Γ_i , $i \in [v]$. Let now the blocks B_j contain the current (yet not traversed) points in $\Gamma_{i,j}$: $B_j = \mathbb{Z}^2 \cap \Gamma_{i,j} \setminus P_1$, for $j = 1, \dots, 2k + 1$. By the previous estimate in (4), the current block sizes satisfy

$$0.85 \cdot \text{Area}(\Gamma_{i,j}) - 0.071 \cdot \text{Area}(\Gamma_{i,j}) \leq |B_j| \leq 1.15 \cdot \text{Area}(\Gamma_{i,j}),$$

or equivalently

$$0.779 \cdot \text{Area}(\Gamma_{i,j}) \leq |B_j| \leq 1.15 \cdot \text{Area}(\Gamma_{i,j}).$$

Obviously $\text{Area}(\Gamma_{i,j}) = \text{Area}(\Gamma_i)/(2k + 1)$, for $j = 1, \dots, 2k + 1$. Since $1.15/0.779 \leq 3/2$, there exist positive integers $b_i \geq 10$, so that

$$b_i \leq |B_j| \leq \frac{3b_i}{2}, \quad i \in [v], \quad j = 1, \dots, 2k + 1. \quad (5)$$

This means that after P_1 has been generated, for each of the rings Γ_i , $i \in [v]$, the blocks have about the same size, as described by (5). We show next that the inward spiral path P_2 traverses the remaining points and closes the cycle by meeting the requirement of small turning angles.

Lemma 5. *Let $k \geq 2$, and $b \geq 6$ be positive integers, and $s, t \in [2k + 1]$. Let Γ_i be a circular ring with inner radius $r - w$, outer radius r , and center o , partitioned by rays into $2k + 1$ congruent ring sectors $\Gamma_{i,j}$ of center angle β , labeled clockwise. Let B_j be a set of lattice points in $\Gamma_{i,j}$, and assume that $b \leq |B_j| \leq 3b/2$, for each $j = 1, \dots, 2k + 1$. Let $n = |\bigcup_{j=1}^{2k+1} B_j|$. Then the set of lattice points $\bigcup_{j=1}^{2k+1} B_j$ can be traversed by a spanning path p_1, \dots, p_n starting at a point $p_1 \in B_s$ and ending at a point $p_n \in B_t$, satisfying the following three conditions: (i) The angle between the first edge of the path and the tangent to the circle of radius op_1 centered at o is at most $\varepsilon/2$. (ii) The angle between the last edge of the path and the tangent to the circle of radius op_n centered at o is at most $\varepsilon/2$. (iii) Each of the turning angles, at p_i , $2 \leq i \leq n - 1$, is at most ε .*

Proof. For convenience relabel the blocks in clockwise order so that $s = 1$. We construct a spanning path P of the form

$$P = (P_\sigma)^*(1, 3, \dots, 2k + 1, 2, 4, \dots, 2k)^m P', \text{ where } m \geq 1. \quad (6)$$

That is, P is obtained by concatenation of several paths denoted P_σ , constructed iteratively, followed by m paths of the form $(1, 3, \dots, 2k + 1, 2, 4, \dots, 2k)$, and finally by a path P' . The numbers enclosed by the second pair of parentheses appearing in the description of P are block labels.

Initially P' is chosen as follows: If $t = 1$ then $P' = 3, 5, \dots, 2k - 1, 1$. If $t \neq 1$ is even, then $P' = 2, 4, \dots, 2k$; and if $t \neq 1$ is odd, then $P' = 3, \dots, t$. One arbitrary point from each block labeled as above is selected and included in the path P' . These points are subsequently removed from the corresponding blocks. Formally, $B_i \leftarrow B_i \setminus \{p \mid p \in P'\}$. Note that the sizes of the blocks are either unchanged or reduced by one, after removing the points in P' .

While the ranges of the current sizes of the blocks are not the same, new paths P_σ are constructed. Assume that in the current iteration, we have $m \leq |B_i| \leq M$, for $i = 1, \dots, 2k + 1$, where $m < M$ are the minimum and the maximum block sizes, respectively. A path P_σ is constructed after which the difference $M - m$ between the (new) values of M and m is reduced by at least one unit. This is achieved by going four times around the ring and visiting two points from each block of size M , and exactly one point from each other block. These points are chosen arbitrarily from the remaining ones. P_σ starts at a point in B_1 .

1. In round 1, traverse only the odd indexes clockwise: $i = 1, 3, \dots, 2k + 1$. If the current block i has size M , or $i \equiv 1 \pmod{4}$, output i .
2. In round 2, traverse only the even indexes clockwise: $i = 2, 4, \dots, 2k$. If the current block i has size M , or $i \equiv 2 \pmod{4}$, output i .
3. In round 3, traverse only the odd indexes clockwise: $i = 1, 3, \dots, 2k + 1$. If the current block i has size M , or $i \equiv 3 \pmod{4}$, output i .
4. In round 4, traverse only the even indexes clockwise: $i = 2, 4, \dots, 2k$. If the current block i has size M , or $i \equiv 0 \pmod{4}$, output i .

Once this path is constructed, the visited points are removed from the corresponding blocks, and the updated blocks are ready for the next iteration. Formally, $B_i \leftarrow B_i \setminus \{p \mid p \in P_\sigma\}$. Since the difference $M - m$ strictly decreases after each iteration (path P_σ), equality $M = m$ is reached after at most $b/2$ iterations.

Assume now that after at most that many iterations, each block has the same number, say m , of points (still to be visited). Using m paths of the form $(1, 3, \dots, 2k + 1, 2, 4, \dots, 2k)$, exhausts all these points.

Observe now that the path P in (6) is Hamiltonian, in clockwise order, starts at a point in B_s ($= B_1$) and ends at a point in B_t . Moreover, it can be checked that each edge in the path connects points in two distinct blocks, separated clockwise by one, two or three other blocks. By Lemma 1, this implies that the three conditions (i), (ii), and (iii) are met. \square

Algorithm. Standard list representations are used for storing the current set of lattice points in each block, corresponding to the ring sectors $\Gamma_{i,j}$, $\Phi_{i,j}$, and Λ_i . To append a new point to the path, given a block label, an arbitrary point is selected from the block (with some exceptions, as specified in Lemma 2 and Lemma 3). Once a point is traversed, it is removed from the corresponding block (list). Since $|P_1| \leq |P_2|$, the overlap in all the lists does not exceed $|\mathbb{Z}^2 \cap \Omega|$. Consequently, the time complexity of the algorithm is (linearly) proportional to the number of points traversed, namely $|\mathbb{Z}^2 \cap \Omega|$.

This concludes the proof of Theorem 1. \square

2.2 Proof of Theorem 3

(Sketch.) The overall procedure for cycle construction is the same, however, some of the parameters need to be adjusted. The path P_1 is constructed in the same way. To construct the path P_2 and complete the cycle, it is enough to derive a stronger analogue of Lemma 5. The condition

$$b \leq |B_j| \leq 3b/2, \quad j = 1, \dots, 2k + 1 \quad (7)$$

is replaced by the condition

$$b \leq |B_j| \leq c_1 b, \quad j = 1, \dots, 2k + 1 \quad (8)$$

where $c_1 = c_1(\rho, c)$ is another constant. This new inequality can be inferred from the even distribution condition, by using packing and covering arguments similar to those we used to derive (7) in the first place. To construct P_2 , the modulus 4 in the calculation has to be replaced with a larger, but still constant modulus, depending on c_1 . Instead of skipping one, two, or three blocks when connecting consecutive points on the path, a larger but still constant skip range needs to be allowed. We omit the details. \square

3 Limitations: proof of Theorem 2

The idea of the proof is simple: if X has a vertex, there is not enough space to turn in the vicinity of that vertex; it remains to give a precise technical argument to implement it. For convenience, rather than considering a fixed grid and placing a large similar copy of X over it, we fix X and place an arbitrarily oriented square grid of small side-length δ over it. Let o be a vertex of X , *i. e.*, a point of ∂X where the clockwise tangent line ℓ_1 differs from the counterclockwise tangent line ℓ_2 . Choose a coordinate system such that o is a point of minimum y -coordinate in X , $\angle o = \pi - 2\alpha$, where $\alpha \in (0, \pi/2)$, and the bisector of $\angle o$ is a vertical line. Refer to Fig. 5. Denote by h_2 the horizontal line through o and by V the upward wedge bounded by ℓ_1 and ℓ_2 . Observe that ℓ_1 and ℓ_2 make an angle of α (in absolute value) with h_2 .

We first find a horizontal line h_1 sufficiently close to o such that the triangle Δ bounded by ℓ_1 , ℓ_2 and h_1 approximates closely the part of X below h_1 . Specifically, let b_1 and d_1 be the two intersection points of h_1 with ∂X ; h_1 is chosen so that the angles made by ob_1 and od_1 with h_2 are each at most 1.01α in absolute value. Now select δ small enough (this is equivalent to selecting λ large enough) so that both, the base (along h_1) and the height of Δ are at least 4δ , namely:

$$|h_1 \cap V| \geq \max(4\delta, 8\delta \cot \alpha). \quad (9)$$




Figure 5: The wedge containing region X .

Translate the line h_1 downwards until (9) becomes an equality: either the base or the height measures 4δ (while the other is at least that long). Note that by the convexity of X , the angles made by ob_1 and od_1 with the x -axis are each still at most 1.01α .

Now re-scale the whole figure so that $\delta = 1$, that is, we have a standard (not necessarily axis-parallel) unit grid superimposed over X . The above choice ensures that the set of grid points in X below h_1 is nonempty. Indeed, the largest empty disk in a unit rectangular lattice has diameter $\sqrt{2}$, and Δ contains a disk of diameter 2.




Figure 6: The turning angle at p is at most β .

Let S be the set of grid points contained in X . Refer to Fig. 6. Let $p \in S$ be a lowest point, *i. e.*, one with a minimum y -coordinate. If there are multiple points, pick the leftmost one. Let h denote the horizontal line through p . Note that h lies below h_1 ; moreover, by (9), the vertical distance between h_1 and h is at least the vertical distance between h and h_2 . Let b and d be the two intersection points of h with ∂X . As before, by the convexity of X , the angles made by ob and od with h_2 , α_1 and α_2 , are each at least α and at most 1.01α . Denote by m the midpoint of bd . By symmetry, we can assume that p is contained in the segment bm . Write $x = |pm|$, and $z = |bm| = |md|$.

Let $\varepsilon_0 = \beta = \alpha/100$. Then by Jensen's inequality, we have $\tan \beta / \tan \alpha \leq 1/100$. Consider a Hamiltonian cycle H of S . We will show that H has some turning angle larger than β . For purpose of contradiction assume that this does not hold. It is easy to dismiss the case $\alpha \in [\pi/4, \pi/2)$: by following fewer than 10 links from p on H , the (almost horizontal) polygonal path must exit V and X , a contradiction. We therefore assume that $\alpha \in (0, \pi/4]$ for the rest of the proof. Since the turning angle at p is at most β , each of the two edges of H incident to p makes an angle at most β (in absolute value) with the line h . In particular, b and d are not in S (since otherwise, their turning angles in H would be larger than β).

The following two properties characteristic to lattices will be used repeatedly:

(*) if p_1 and p_2 are lattice points, the reflection of p_1 with respect to p_2 is also a lattice point.

(**) if p_1, p_2, p_3 are lattice points, then $p_1 + \overrightarrow{p_2 p_3}$ is also a lattice point.

Let o', b' and d' denote the reflections of o, b and d with respect to p . Since p is contained in the segment bm , d' lies left of b and b' lies left of d . By construction we have $|b'd| = |d'b|$. Also by construction, and by property (*), we deduce that the triangle $\Delta o'd'b'$ is empty of points in S in its interior.

Denote the left and right neighbors of p in H by q and r . Note that q lies strictly above h (since p is the leftmost lattice point on h). In particular, q, p and r are non-collinear. Denote by a the intersection between $d'o'$ and the extension of ob , and by y the (vertical) distance between a and h . The horizontal segment $d'b$ is subdivided into two segments, of lengths s_1 and s_2 by the projection of a onto h ; write $s = |d'b|$, so $s = s_1 + s_2$. By construction we have $|pd| = |pd'|$, or $x + z = (z - x) + s$, hence

$$|b'd| = |d'b| = s = 2x. \quad (10)$$

Let γ be the angle made by pa with the horizontal line h . Since pq lies above pa , we have $\gamma \leq \beta$. We can express $\tan \gamma$ as follows.

$$\tan \gamma = \frac{y}{|pd'| - s_1} = \frac{y}{|pd| - s_1} = \frac{y}{(x+z) - s_1} \geq \frac{y}{x+z} = \frac{s_2}{x+z} \cdot \frac{y}{s_2} = \frac{s_2}{x+z} \tan \alpha_1.$$

This implies $\tan \beta \geq \tan \gamma \geq \frac{s_2}{x+z} \tan \alpha_1$, and further that

$$s_2 \leq \frac{\tan \beta}{\tan \alpha_1} (x+z) \leq \frac{\tan \beta}{\tan \alpha} (x+z) \leq \frac{x+z}{100}.$$

Since $\tan \alpha_2 = y/s_1$, and $\tan \alpha_1 = y/s_2$, we also have (recall that $\alpha \leq \pi/4$):

$$s_1 = \frac{\tan \alpha_1}{\tan \alpha_2} s_2 \leq \frac{\tan 1.01\alpha}{\tan \alpha} s_2 \leq \frac{\tan 1.01\pi/4}{\tan \pi/4} s_2 \leq 1.02 s_2 \leq \frac{x+z}{98}.$$

Summing these two inequalities yields $s = s_1 + s_2 \leq (x+z)/49$. By (10), we also have $x = s/2 \leq (x+z)/98$. Finally, we obtain upper bounds on x and s in terms of z :

$$x \leq \frac{z}{97}, \text{ and } s \leq \frac{z}{48}. \quad (11)$$

We first argue that p is the only point of S contained in the interior of the segment bb' . Assume that there is another such point, say p_1 ; we select the rightmost one. If p_1 lies in the interior of the segment pb' , then by the property (*), we get that the reflection of p_1 with respect to p is another point in $S \cap h$, left of p , a contradiction. Hence $p_1 \in b'd$, if it exists. In particular, this means that $|S \cap \Delta o'd'b'| = 1$, that is, with the exception of p , there are no other points of S in the triangle $\Delta o'd'b'$.

We next show that $|bq|$ and $|b'r|$ are both small with respect to z . Denote by q_1 the intersection of the extensions of ob and pq , and observe that $|bq| \leq |bq_1|$. By the Law of Sines in the triangle Δpbq_1 , we deduce that

$$|bq_1| \leq \frac{\sin \beta}{\sin(\alpha_1 - \beta)} |bp| = \frac{\sin \beta}{\sin(\alpha_1 - \beta)} (z - x) \leq \frac{\beta}{(\alpha_1 - \beta)/2} z \leq \frac{2\beta}{(\alpha - \beta)} z = \frac{2z}{99},$$

thus $|bq| \leq 2z/99 \leq z/48$.

Similarly, denote by r_1 the intersection of $o'b'$ with pr , and by r_2 the intersection of the extensions of od and pr . By the Law of Sines in the triangle Δpdr_2 , we find that

$$|dr_2| \leq \frac{\sin \beta}{\sin(\alpha_2 - \beta)} |pd| = \frac{\sin \beta}{\sin(\alpha_2 - \beta)} (z + x) \leq \frac{2\beta}{(\alpha - \beta)} \cdot \frac{98z}{97} = \frac{2}{99} \cdot \frac{98z}{97} \leq \frac{z}{48}.$$

Therefore, by the triangle inequality,

$$|b'r| \leq |b'r_2| \leq |b'd| + |dr_2| \leq z/48 + z/48 = z/24.$$

Finally, consider the reflection \hat{p} of p with respect to the midpoint of qr , say g . By property (**), \hat{p} is a lattice point. Recall that q , p and r are non-collinear, so \hat{p} lies strictly above h , hence it is distinct from p . Moreover,

$$|p\hat{p}| = 2|pg| = |\vec{bq} + \vec{b'r}| \leq |bq| + |b'r| \leq z/48 + z/24 = z/16.$$

It follows that $\hat{p} \in \Delta o'bb'$, thus $\hat{p} \in S \cap \Delta o'd'b'$, a contradiction. This concludes the proof of Theorem 2. \square

4 Concluding remarks

While our proofs of Theorems 1 and 3 are tailored for the disk case, we believe that the general ideas used there will be applicable for the general case (Conjectures 1 and 2) in constructing Hamiltonian tours with small turning angles. The two basic ideas are:

- (i) Partitioning the input region into smaller sub-regions that can be conveniently linked.
- (ii) Visiting only some points in a sub-region and going back to that sub-region over and over again, until all points are exhausted while maintaining the turning angle constraint.

We outline two further directions.

1. The proofs of Theorems 1 and 3 are constructive and lead to linear-time algorithms for computing a Hamiltonian tour with each turning angle at most ε : for an input point set S , such a tour can be computed in $O(|S|)$ time. However, some constant factors in our tour construction are too large and need to be reduced. These constants are of little concern in the existence part of Theorem 1, but are relevant for the algorithmic part, because they impose large disks (regions, in general) as inputs.

2. Let S be a set of n points in the plane, and $\alpha \in [0, \pi]$. A necessary condition for S to admit a Hamiltonian cycle with each turning angle at most $\pi - \alpha$ is the following:

- [T] Each point $q \in S$ determines a triangle Δpqr with $\angle pqr \geq \alpha$; or equivalently, that the turning angle at q in the sequence p, q, r is at most $\pi - \alpha$.

It is easy to construct examples which show that the above condition does not suffice to guarantee a Hamiltonian cycle with each turning angle at most $\pi - \alpha$: for instance: (i) an equilateral triangle with its center, for $\alpha = \pi/3$, and (ii) a square with its center, for $\alpha = \pi/2$. Suppose that S satisfies the above condition [T]. In particular, [T] implies that the interior angle at each point on $\text{conv}(S)$ is at least α . Does any such S admit a Hamiltonian cycle with each turning angle at most $\pi - c \cdot \alpha$, for some absolute constant $c > 0$? By the above examples, $c \leq 1/2$, if it exists.

Acknowledgment. The author is grateful to an anonymous reviewer for suggesting the argument yielding a sharper bound in the proof of Theorem 2, as materialized in Corollary 1.

References

- [1] E. Ackerman, O. Aichholzer, and B. Keszegh, Improved upper bounds on the reflexivity of point sets, *Computational Geometry: Theory and Applications*, **42(3)** (2009), 241–249.
- [2] A. Aggarwal, D. Coppersmith, S. Khanna, R. Motwani, and B. Schieber, The angular-metric traveling salesman problem, *SIAM Journal on Computing*, **29(3)** (1999), 697–711.
- [3] P. Aggarwal, P. Raghavan, and H. Tamaki, Motion planning for a steering-constrained robot through moderate obstacles, *Proceedings of the 27th ACM Symposium on the Theory of Computing*, 1995, pp. 343–352.
- [4] O. Aichholzer, T. Hackl, M. Hoffmann, C. Huemer, A. Pór, F. Santos, B. Speckman, and B. Vogtenhuber, Maximizing maximal angles for plane straight-line graphs, *Proceedings of the 10th International Workshop on Algorithms and Data Structures*, 2007, vol. 4619/2007 of LNCS, pp. 458–469.
- [5] E. M. Arkin, M. A. Bender, E. D. Demaine, S. P. Fekete, J. S. B. Mitchell, and S. Sethia, Optimal covering tours with turn costs, *SIAM Journal on Computing*, **35(3)** (2005), 531–566.
- [6] E. M. Arkin, S. Fekete, F. Hurtado, J. Mitchell, M. Noy, V. Sacristán, and S. Sethia, On the reflexivity of point sets, in *Discrete and Computational Geometry: The Goodman-Pollack Festschrift* (B. Aronov, S. Basu, J. Pach, M. Sharir, editors), Springer, 2003, pp. 139–156.
- [7] I. Bárány, A. Pór, and P. Valtr, Paths with no small angles, *SIAM Journal on Discrete Mathematics*, **23(4)** (2009), 1655–1666.
- [8] J. Boissonat, A. Cérézo, and J. Leblond, Shortest paths of bounded curvature in the plane, *Journal of Intelligent and Robotic Systems: Theory and Applications*, **11** (1994), 5–20.
- [9] A. Dumitrescu, J. Pach, and G. Tóth, Drawing Hamiltonian cycles with no large angles, *Proceedings of the 17th International Symposium on Graph Drawing*, Chicago, Sept. 2009, Volume 5849 of LNCS, Springer Verlag, March 2010, pp. 3–14.
- [10] S. P. Fekete and G. J. Woeginger, Angle-restricted tours in the plane, *Computational Geometry: Theory and Applications*, **8(4)** (1997), 195–218.
- [11] T. Frachard, Smooth trajectory planning for a car in a structured world, *Proceedings of the IEEE International Conference on Robotics*, 1989, pp. 318–323.
- [12] P. Jacobs and J. Canny, Planning smooth paths for mobile robots, *Proceedings of the IEEE International Conference on Robotics*, 1989, pp. 2–7.
- [13] J. L. W. V. Jensen, Sur les fonctions convexes et les inégalités entre les valeurs moyennes, *Acta Mathematica*, **30(1)** (1906), 175–193.
- [14] J. S. B. Mitchell, Geometric shortest paths and network optimization, in *Handbook of Computational Geometry (J.-R. Sack and J. Urrutia, editors)*, Chap. 15, Elsevier, 2000, pp. 633–701.
- [15] M. Mitzenmacher, E. Upfal, *Probability and Computing*, Cambridge University Press, 2005.