The Great Wall of David Shin

Tiankai Liu

12–15 June 2015

On 9 May 2010, David Shin posed the following puzzle in a Facebook note:

Problem 1. You're blindfolded, disoriented, and standing one mile from the Great Wall of China. How far do you have to walk to guarantee that you will run into the wall? Assume that the wall is infinitely long and straight.

Though Ian Le posted the correct answer on 7 August, no one posted a proof of optimality. After various unhelpful remarks by the author, such as "yan and i convinced each other that we have an convincing argument," the problem seems to have languished for nearly half a decade.

Update (19 January 2016): Never mind, the proof was given in H. Joris, "Le chasseur perdu dans la forêt," *Elemente der Mathematatik* **35** (1), 10 January 1980. It is available at `http://infoscience.epfl.ch/record/130395/files/PPN378850199_0035__0_0.pdf`.

1 The shortest path

Before giving the answer to Problem 1, let us slightly reformulate the problem. Let O be the origin in the real plane $\mathbb{R}^2$, and let Δ be the closed unit disk,

$$\Delta = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leqslant 1\}.$$

Problem 1 asks for the shortest curve C starting at O that

(a) meets every line tangent to $\partial\Delta$.

We claim this is equivalent to saying

(b) the convex hull $\mathcal{H}$ of C contains Δ.

Proof. To show (a) implies (b): If $\mathcal{H}$ misses some point $P \in \Delta$, then, since $\mathcal{H}$ is the intersection of all open half-planes containing C, there must exist some open half-plane $\mathcal{R}^+$ containing C but not P. Say $\mathcal{R}^+$ is bounded by the line ℓ'; then let ℓ be the unique line parallel to ℓ', tangent to $\partial\Delta$, and lying outside $\mathcal{R}^+$. Then evidently C fails to meet ℓ.

To show (b) implies (a): If C fails to meet a line ℓ tangent to $\partial\Delta$ at a point P, then C lies entirely in the open half-plane $\mathcal{R}^+$ containing O and bounded by ℓ. Since $\mathcal{R}^+$ is convex, the convex hull $\mathcal{H}$ also lies in $\mathcal{R}^+$, and in particular $P \notin \mathcal{H}$. $\square$

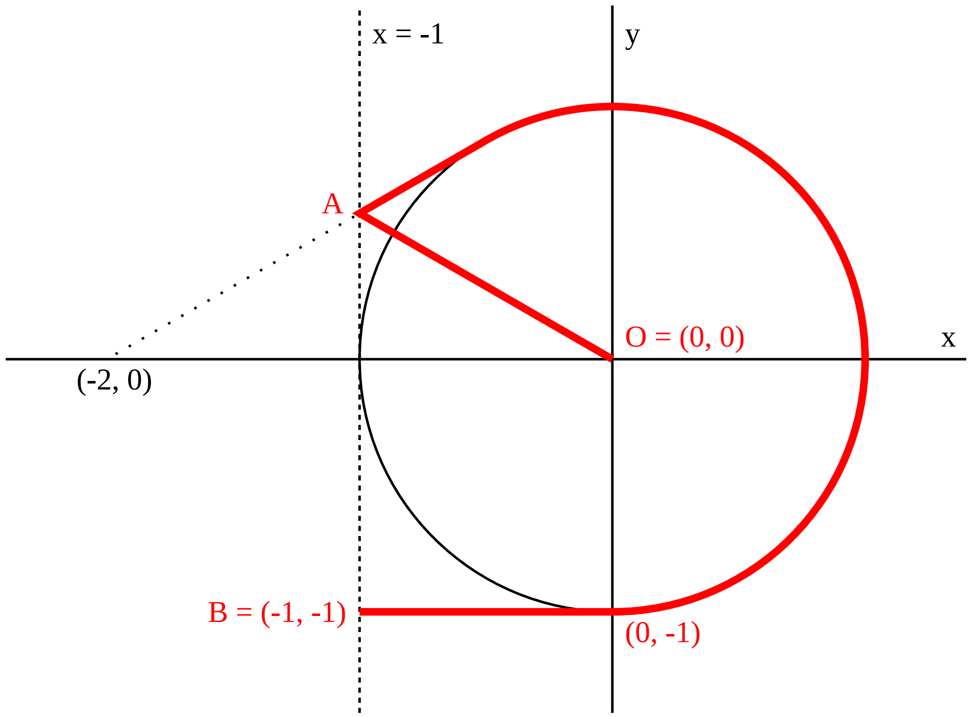

Figure 1: The optimal path

Let us say a curve C starting at O *touches the Great Wall of Shin* if it starts at O and satisfies the equivalent conditions (a) and (b). Problem 1 asks for the minimum possible length of C.

The answer is $\sqrt{3} + \frac{7\pi}{6} + 1$ and can be realized by the following curve (Figure 1). Start at $(0, 0)$, and walk straight for $\frac{2}{\sqrt{3}}$ miles, say, to $A := (-1, \frac{1}{\sqrt{3}})$. Make a $120°$ turn and walk straight for $\frac{1}{\sqrt{3}}$ mile, say, to $(\frac{-1}{2}, \frac{\sqrt{3}}{2})$, which lies on $\partial\Delta$. Without making an abrupt turn, walk $\frac{7\pi}{6}$ miles around $\partial\Delta$, to $(0, -1)$. Continue straight for 1 mile, to $B := (-1, 1)$, and stop. This curve clearly satisfies condition (b), and therefore touches the Great Wall of Shin.

2 The proof of optimality

We would like to show that the curve of length $\sqrt{3} + \frac{7\pi}{6} + 1$ in answer to Problem 1 cannot be beaten. For technical reasons, we would like to restrict our attention to contenders that are polygonal paths.[1] Here's why this reduction is justified:

Lemma 2. *If C is a finite-length curve that touches the Great Wall of Shin, then for every $\epsilon > 0$ there exists a polygonal path C', also touching the Great Wall of Shin, that has length at most $|C| + \epsilon$.*

Proof. Fix some $0 < \delta \ll 1$. Let

$$O = P_0,\ P_1,\ \ldots,\ P_n$$

[1] I'll use "path" interchangeably with "rectifiable curve," so "paths" may intersect themselves. I only feel like doing this because it sounds weird to talk about straight paths as "curves."

be points on C, spaced less than δ apart in arc length, such that P_n is the endpoint of C. Then every point of C is within $\delta/2$ of one of the points P_i. Therefore, every point of the convex hull of C, being a weighted average of points of C, is also with $\delta/2$ of the convex hull H of $\{P_0, \dots, P_n\}$. It follows that the disk of radius $1 - \frac{\delta}{2}$ centered at the origin lies within H. Therefore, if we let C' be the polygonal path obtained from $P_0 \cdots P_n$ via a dilation of ratio $1/(1 - \frac{\delta}{2})$ from the origin, then

$$|C'| < \frac{|C|}{1 - \frac{\delta}{2}}$$

and C' touches the Great Wall of Shin. For sufficiently small δ depending on $|C|$, we can ensure that $|C'| \leqslant |C| + \epsilon$, as desired. □

Therefore, the solution to Problem 1 is completed by the following Proposition:

Proposition 3. *For any $n \in \mathbb{N}$ and any points $P_1, \dots, P_n \in \mathbb{R}^2$, let $\mathcal{P}$ denote the polygonal path $OP_1 \cdots P_n$, and $\mathcal{H}$ the convex hull of $\{O, P_1, \dots, P_n\}$. Suppose that $\mathcal{P}$ touches the Great Wall of Shin. Then $\mathcal{P}$ must have length greater than*

$$\sqrt{3} + \frac{7\pi}{6} + 1.$$

Proof. Our proof proceeds by a sequence of reductions. Let us say

- $\mathcal{P}$ is *non-redundant* if, for any $i \in \{1, \dots, n\}$, the convex hull of $\{O, P_1, \dots, \widehat{P_i}, \dots, P_n\}$ does not contain Δ.
- $\mathcal{P}$ is *non-self-intersecting* if the *interiors* of the line segments $\overline{OP_1}$, $\overline{P_1P_2}$, …, $\overline{P_{n-1}P_n}$ are pairwise disjoint. (Caution: *a priori* we allow the endpoint of one segment to lie in the interior of another segment, so our definition may be a bit different from standard terminology.)
- The *violence* of $\mathcal{P}$ is the number of segments $\overline{P_{k-1}P_k}$ $(2 \leqslant k \leqslant n)$ that cut Δ in two.[2] We say $\mathcal{P}$ is *non-violent* if its violence equals zero.

Given any polygonal path $\mathcal{P}$ that touches the Great Wall of Shin, we will explain how to construct another one, $\mathcal{P}'$, of lesser or equal length, that is non-redundant, non-self-intersecting, and non-violent. Then we will prove Proposition 3 for $\mathcal{P}'$ under these additional assumptions.

Step 1: Reduction to non-redundant $\mathcal{P}$. This should be pretty obvious; we just repeatedly throw away one redundant vertex at a time.

Henceforth, assume $\mathcal{P}$ is non-redundant. Let us pause and introduce some new notation. The convex hull $\mathcal{H}$ of $\{O, P_1, \dots, P_n\}$ is a polygon with exactly n vertices, namely $P_1, \dots, P_n$. However, a priori, these vertices may not be in order. So let us write $A_0, \dots, A_{n-1}$, with indices taken modulo n, to be the vertices of $\mathcal{H}$ in some order (either clockwise or counter-clockwise).

The following lemma is pretty obvious, and I will dispense with the proof.

[2] We say that a closed line segment σ *cuts* a convex region $R \subset \mathbb{R}^2$ *in two* if $R \backslash \sigma$ is disconnected.

Lemma 4. *With $A_0, \ldots, A_{n-1}$ arising as above from a non-redundant path that touches the Great Wall of Shin, let i and j be two different members of $\mathbb{Z}/n\mathbb{Z}$. Then there are two mutually exclusive possibilities:*

(a) $j = i \pm 1$, and the segment $\overline{A_i A_j}$ is either disjoint from Δ or tangent to the unit circle $\partial\Delta$.

(b) $j \neq i \pm 1$, and the segment $\overline{A_i A_j}$ cuts Δ in two. It also cuts $\mathcal{H}$ in two, and any line segment joining points of these two components together must meet the segment $\overline{A_i A_j}$.

Step 2: Reduction to non-self-intersecting $\mathcal{P}$. Recall that we are assuming $\mathcal{P}$ is non-redundant. For this step, imagine that the set $\{P_1, \ldots, P_n\}$ is fixed. There are $n!$ polygonal paths that start at O and visit all the points of $\{P_1, \ldots, P_n\}$ in some order. Each such path is non-redundant and touches the Great Wall of Shin. Therefore, we may assume that, among all $n!$ such paths, $\mathcal{P} = OP_1 \cdots P_n$ happens to be one of minimum length. We claim that this implies that $\mathcal{P}$ is non-self-intersecting.

If not, there exist $1 \leqslant i < j \leqslant n$ such that the interiors of segments $\overline{P_{i-1}P_i}$ and $\overline{P_{j-1}P_j}$ meet, where we temporarily write P_0 for O for the duration of Step 2. There are two cases to consider.

Case 1: The segments $\overline{P_{i-1}P_i}$ and $\overline{P_{j-1}P_j}$ meet transversely at a point X. Then by the triangle inequality

$$\begin{aligned}|P_{i-1}P_i| + |P_{j-1}P_j| &= |P_{i-1}X| + |XP_i| + |P_{j-1}X| + |XP_j| \\ &= |P_{i-1}X| + |XP_{j-1}| + |P_iX| + |XP_j| \\ &> |P_{i-1}P_{j-1}| + |P_iP_j|,\end{aligned}$$

so

$$\mathcal{P}' := \underbrace{P_0 \cdots P_{i-1}}_{\text{increasing indices}} \quad \overbrace{P_{j-1} \cdots P_i}^{\text{decreasing indices}} \quad \underbrace{P_j \cdots P_n}_{\text{increasing indices}}$$

is strictly shorter than $\mathcal{P}$, a contradiction.

Case 2: The segments $\overline{P_{i-1}P_i}$ and $\overline{P_{j-1}P_j}$ are collinear. This is an unusual situation: since $P_1, \ldots, P_n$ are all vertices of the convex polygon $\mathcal{H}$, it must be that $(i, j) = (0, 1)$, and $O \in \overline{P_1P_2}$. In particular, $\overline{P_1P_2}$ cuts Δ in two. Therefore, by Lemma 4, $\overline{P_1P_2}$ also cuts $\mathcal{H}$ in two. There must then exist some $k \in \{3, \ldots, n\}$ such that P_{k-1} and P_k lie in different components of $\mathcal{H}\backslash\overline{P_1P_2}$. But this means that $\overline{P_1P_2}$ and $\overline{P_{k-1}P_k}$ intersect transversely, as in Case 1 above, which was shown to be impossible. Therefore, Case 2 is also impossible.

Henceforth, we may assume that $\mathcal{P}$ is non-self-intersecting.

Step 3: Reduction to non-violent $\mathcal{P}$. Recall that we are assuming $\mathcal{P}$ is non-redundant and non-self-intersecting. Assuming that the violence v of $\mathcal{P}$ is positive, we will construct a non-redundant, non-self-intersecting path

$$\mathcal{Q} = OQ_1 \cdots Q_m$$

that touches the Great Wall of Shin, has violence equal to $v-1$, and has length at most the length of $\mathcal{P}$. (Possibly $m > n$, but this does not matter.)

Since $v > 0$, we may let $k \in \{2, \ldots, n\}$ be the largest index for which $\overline{P_{k-1}P_k}$ cuts Δ (and therefore also $\mathcal{H}$) in two. By maximality of k, Lemma 4 implies that $P_k, P_{k+1}, \ldots, P_n$ must traverse some of the vertices of the convex polygon $\mathcal{H}$ in order. So without loss of generality, let us suppose $A_i = P_i$ for $k \leqslant i \leqslant n$. Since $\mathcal{P}$ is non-self-intersecting, Lemma 4(b) tells us that $P_1, \ldots, P_{k-2}$ must all lie in one component of $\mathcal{H} \backslash \overline{P_{k-1}P_k}$, and $P_{k+1}, \ldots, P_n$ must all lie in the other component. It follows that $A_1 = P_{k-1}$. And since $\mathcal{P}$ touches the Great Wall of Shin, we must have

$$3 \leqslant k \leqslant n-1.$$

Therefore, it makes sense to let $\mathcal{R}^-$ and $\mathcal{R}^+$ be the two open half-planes, separated by the line $P_{k-1}P_k$, that contain P_{k-2} and P_n, respectively.

Set $Q_i = P_i$ for $i \leqslant k-1$. If it happens that

$$|P_{k-1}P_n| = |A_1A_n| \leqslant |A_1A_k| = |P_{k-1}P_k|, \tag{1}$$

then we may simply let

$$\mathcal{Q} := O \underbrace{P_1 \cdots P_{k-1}}_{\text{increasing indices}} \overbrace{P_n \cdots P_k}^{\text{decreasing indices}},$$

and be done; $\mathcal{Q}$ has violence $v-1$ because $\overline{P_{k-1}P_n} = \overline{A_1A_n}$ no longer cuts $\mathcal{H}$ in two, and inequality (1) ensures that $\mathcal{Q}$ is not longer than $\mathcal{P}$.

Unfortunately, we cannot simply assume that inequality (1) holds, but there is an easy fix. Let ℓ be the unique line through P_{k-1} that is tangent to $\partial\Delta$ at a point $T \in \partial\Delta \cap \mathcal{R}^+$. Note that the path $P_k \cdots P_n$ must meet ℓ at some point, because $\overline{P_kP_{k-1}}$ cuts Δ in two but $\overline{P_nP_{k-1}}$ doesn't. Intuitively, we want to consider the shortest curve γ from P_k to anywhere on the line ℓ that lies entirely within the closure of $\mathcal{R}^+ \backslash \Delta$, and replace the path $P_k \cdots P_n$ by a sufficiently good polygonal approximation $P'_k \cdots P'_m$ of C, so that by setting

$$\mathcal{Q} := O \underbrace{P_1 \cdots P_{k-1}}_{\text{increasing indices}} \overbrace{P'_m \cdots P'_k}^{\text{decreasing indices}}, \tag{2}$$

we can carry on just as above.

To say this rigorously, we claim that there is a unique shortest path γ from P_k to ℓ within $\operatorname{clos}(\mathcal{R}^+ \backslash \Delta)$, namely:

- Case 1: γ is the perpendicular from P_k to ℓ, if this perpendicular lies in $\operatorname{clos}(\mathcal{R}^+ \backslash \Delta)$.
- Case 2: Otherwise, let $X, Y \in \partial\Delta \cap \mathcal{R}^+$ be the unique points such that $\overline{P_kX}$ is tangent to $\partial\Delta$, and $OY \parallel \ell$. Let Z be the foot of the perpendicular from Y to ℓ, so that $\overline{YZ}$ is tangent to $\partial\Delta$. Then γ is the union of the segment $\overline{P_kX}$, the arc XY of Δ in $\mathcal{R}^+$, and the segment $\overline{YZ}$.

The first case is trivial. In the second case, it is easy to deduce from Proposition 5 in the Appendix that the shortest path γ exists and is (in the notation of that Proposition) of the form $\gamma(P_k, W; \Delta, \mathcal{R}^+)$ for some W on the ray $\overrightarrow{P_kT}$ beyond T. Such a γ comes in three parts: a line segment starting from P_k, an arc of Δ, and finally a line segment ending on W. The final line segment must be perpendicular to ℓ, or else γ could be shortened by replacing a tiny portion of the end with a segment perpendicular to ℓ. Thus, W must be the point Z defined above, as claimed.

So indeed we can define $P'_k \cdots P'_m$ and $\mathcal{Q}$ by equation (2) above, such that P'_m is the endpoint of γ. (In case 2, we can perform the polygonal approximation in such a way that $\overline{P'_{i-1}P'_i}$ is tangent to $\partial\Delta$ for every $i \in \{k+1, \ldots, m\}$.) We must verify that

$$|P_{k-1}P_k| \geqslant |P_{k-1}P'_m|.$$

To this end, note that, since $\mathcal{P}$ is non-self-intersecting, the interiors of segments $\overline{OP_1}$ and $\overline{P_{k-1}P_k}$ are disjoint, so O doesn't lie in $\mathcal{R}^+$. Then it is easy to see that the distance to P_{k-1} decreases monotonically as we travel along γ from P_k to P'_m. So indeed we have replaced $\mathcal{P}$ by a shorter, less violent $\mathcal{Q}$.

Henceforth, we may assume $\mathcal{P}$ is non-violent.

Step 4: Solution for non-violent $\mathcal{P}$. Since $\mathcal{P}$ is non-redundant and non-violent, the path $P_1 \cdots P_n$ must avoid $\text{int}(\Delta)$ entirely, and therefore Lemma 4 implies that it must wind around Δ in order, so that we may assume $P_i = A_i$ for all $i \in \{1, \ldots, n\}$. If $\overline{P_nP_1}$ isn't tangent to Δ, then we can always move P_n closer to P_{n-1} so that it *is* tangent. Let $\ell = \overleftrightarrow{P_nP_1}$, and let $\mathcal{R}^+$ be the open half-plane bounded by ℓ and containing most of Δ. Then Proposition 5 applies directly, and it suffices to show that, if (using the notation of that Proposition)

$$C := \overline{OA} \cup \gamma(A, B; \Delta, \mathcal{R}^+)$$

for two points $A, B \in \ell$ on opposite sides of the point $\ell \cap \Delta$, then then $|C| \geqslant \sqrt{3} + \frac{7\pi}{6} + 1$.

Let's impose some coordinates to help with the calculation (see Figure 1). Say ℓ is the line $x = -1$ and $A = (-1, a)$ for some $a > 0$. As in Step 3, our ability to wiggle B along ℓ implies that the final straight segment of γ is perpendicular to ℓ, so that $B = (-1, -1)$. By reflecting O across ℓ, we see that A is optimal when the ray from $(-2, 0)$ to A is tangent to Δ, or in other words $a = \frac{1}{\sqrt{3}}$. Then C is exactly the curve given as the answer in Section 1. This completes the proof of Proposition 3 and the solution to Problem 1. $\square$

Appendix: Paths avoiding convex bodies

Here we prove a basic fact (Proposition 5) about paths with given endpoints that avoid a given convex region of the plane, which was used in steps 3 and 4 of the proof of Proposition 3 above.

Let $A, B \in \mathbb{R}^2$ be two distinct points, and $\mathcal{R}^+ \subset \mathbb{R}^2$ one of the open half-planes bounded by the line $\overleftrightarrow{AB}$. Let $\mathcal{D} \subset \mathbb{R}^2$ be a compact, strictly[3] convex set such that $A, B \notin \mathcal{D}$, but $\mathcal{D}$ meets both $\mathcal{R}^+$ and the line segment $\overline{AB}$.

[3] By "strictly" we mean every line either splits $\mathcal{D}$ in two, or else meets $\mathcal{D}$ in at most one point. However, the strictness assumption is not essential, and we include it only to avoid a few technical annoyances.

Let A' denote the unique point on $\partial\mathcal{D} \cap \mathcal{R}^+$ such that $\overleftrightarrow{AA'} \cap \mathcal{D} = \{A'\}$. Define $B' \in \partial\mathcal{D} \cap \mathcal{R}^+$ likewise with respect to B.

Let $\gamma(A, B; \mathcal{D}, \mathcal{R}^+)$ denote the path from A to B that consists of the line segment $\overline{AA'}$, the arc along $\partial\mathcal{D}$ from A' to B' inside $\mathcal{R}^+$, and the line segment $\overline{B'B}$.

Proposition 5. *With the above notation and assumptions, $\gamma := \gamma(A, B; \mathcal{D}, \mathcal{R}^+)$ is the shortest path from A to B that lies in $\mathcal{R}^+$ (except at its endpoints) and avoids the interior of $\mathcal{D}$.*

Proof. Suppose for the sake of contradiction that there were a shorter alternative β to γ.

Step 1: Reduction to the case where β is polygonal. Fix $\eta := 1 - \frac{|\beta|}{|\gamma|} > 0$. Choose any point $O \in \mathcal{D} \cap \overline{AB}$, and let

$$\epsilon := \min\left\{\operatorname{dist}(\{A, B\}, \mathcal{D}),\ \frac{\eta}{3}\operatorname{dist}(O, \beta),\ \frac{\eta^2}{6}|\gamma|\right\} > 0.$$

Choose a sequence of points

$$A = P_0,\ P_1,\ \ldots,\ P_n,\ P_{n+1} = B$$

along β, spaced less than ϵ apart in arc length along β. For each $i \in \{1, \ldots, n\}$, let Q_i be the point on the ray $\overrightarrow{OP_i}$ lying a distance 3ϵ beyond P_i. Let β' be the polygonal path

$$AQ_1 \cdots Q_nB.$$

We claim that β' avoids $\mathcal{D}$.

Since $\epsilon \leqslant \operatorname{dist}(\{A, B\}, \mathcal{D})$ and, by assumption,

$$|AP_0|,\ |P_nB| < \epsilon,$$

it is easy to see that the segments $\overline{AQ_1}$ and $\overline{Q_nB}$ avoid $\mathcal{D}$.

We claim that all segments $\overline{Q_iQ_{i+1}}$ avoid $\mathcal{D}$ as well. Indeed, suppose for the sake of contradiction that $X \in \mathcal{D} \cap \overline{Q_iQ_{i+1}}$. By the law of cosines,

$$|OP_i|^2 + |OP_{i+1}|^2 - 2|OP_i| \cdot |OP_{i+1}| \cdot \cos\angle P_iOP_{i+1} = |P_iP_{i+1}|^2,$$

so

$$0 \leqslant 1 - \cos\angle P_iOP_{i+1} = \frac{|P_iP_{i+1}|^2 - (|OP_i| - |OP_{i+1}|)^2}{2|OP_i| \cdot |OP_{i+1}|} \leqslant \frac{|P_iP_{i+1}|^2}{2|OP_i| \cdot |OP_{i+1}|}.$$

Moreover,

$$\begin{aligned}|Q_iQ_{i+1}|^2 &= |OQ_i|^2 + |OQ_{i+1}|^2 - 2|OQ_i| \cdot |OQ_{i+1}| \cdot \cos\angle Q_iOQ_{i+1} \\ &= (3\epsilon + |OP_i|)^2 + (3\epsilon + |OP_{i+1}|)^2 - 2(3\epsilon + |OP_i|)(3\epsilon + |OP_{i+1}|)\cos\angle P_iOP_{i+1} \\ &= |P_iP_{i+1}|^2 + 6\epsilon(3\epsilon + |OP_i| + |OP_{i+1}|)(1 - \cos\angle P_iOP_{i+1}),\end{aligned}$$

Since

$$|OP_i|,\ |OP_{i+1}| \geqslant \operatorname{dist}(O, \beta) \geqslant \frac{3\epsilon}{\eta},$$

we therefore have

$$1 \leqslant \frac{|Q_iQ_{i+1}|^2}{|P_iP_{i+1}|^2} \leqslant 1 + \frac{3\epsilon(3\epsilon + |OP_i| + |OP_{i+1}|)}{|OP_i| \cdot |OP_{i+1}|} \leqslant 1 + \eta^2 + \eta + \eta = (1+\eta)^2,$$

and

$$1 \leqslant \frac{|Q_iQ_{i+1}|}{|P_iP_{i+1}|} \leqslant 1 + \eta. \tag{3}$$

In particular,

$$|Q_iX| \leqslant |Q_iQ_{i+1}| \leqslant (1+\eta)|P_iP_{i+1}| < 2\epsilon,$$

and

$$|OX| \geqslant |OQ_i| - |Q_iX| > |OP_i| + 3\epsilon - 2\epsilon = |OP_i| + \epsilon.$$

Now between P_i and P_{i+1}, by assumption, the arc of β has length less than ϵ, and yet this arc must meet the ray $\overrightarrow{OX}$ at some point Y outside $\mathrm{int}(\mathcal{D})$. So

$$|P_iY| \geqslant |OY| - |OP_i| \geqslant |OX| - |OP_i| > \epsilon,$$

by the previous inequality. This is a contradiction, so our assumption of the existence of X must have been false.

So indeed the polygonal path β' is disjoint from $\mathcal{D}$. Moreover, inequality (3) implies that

$$\begin{aligned} |\beta'| &< |P_1Q_1| + (1+\eta)|\beta| + |P_nQ_n| \\ &= (1+\eta)|\beta| + 6\epsilon \\ &\leqslant (1+\eta)(1-\eta)|\gamma| + \eta^2|\gamma| \\ &= |\gamma|. \end{aligned}$$

We have thus reduced to the case where β is a polygonal path, which henceforth we will call

$$\mathcal{P} := P_0 \cdots P_{n+1},$$

where $P_0 = A$ and $P_{n+1} = B$.

Step 2: Reduction to the case where each segment $\overline{P_iP_{i+1}}$ meets $\mathcal{D}$ at one point. Let us induct primarily on n (the number of intermediate vertices in the polygonal path $\mathcal{P}$), and, in cases of equal n, on the number of segments $\overline{P_iP_{i+1}}$ $(0 \leqslant i \leqslant n)$ disjoint from $\mathcal{D}$.

If, for some pair of indices $(i,j) \neq (0,n+1)$ with $j \geqslant i+2$, the segment $\overline{P_iP_j}$ fails to cut Δ in two, then throw away $P_{i+1},\ldots,P_{j-1}$. The resulting polygonal path $\mathcal{P}' := P_0 \cdots P_iP_j \cdots P_n$ is no longer than the original $\mathcal{P}$, so we are done by induction on n. Therefore, we may assume that no such pair (i,j) exists.

If there exists a segment of the form $\overline{P_iP_{i+1}}$ disjoint from $\mathcal{D}$, then either $i \neq 0$ or $i \neq n$. Without loss of generality, assume the former. Then, as either $(i-1,i+1) = (0,n)$ or $\overline{P_{i-1}P_{i+1}}$ cuts $\mathcal{D}$ in two, there must exist a unique point P_i' strictly between P_i and P_{i+1} such that $\overline{P_{i-1}P_i'}$ meets $\mathcal{D}$ in exactly one point. Then replacing P_i by P_i', we are done by induction, since n stays the same, and the polygonal path $\mathcal{P}' := P_0 \cdots P_{i-1}P_i'P_{i+1} \cdots P_{n+1}$ is shorter than $\mathcal{P}$ and has fewer segments disjoint from $\mathcal{D}$.

So we are reduced to the case where $\overline{P_iP_{i+1}}$ meets $\mathcal{D}$ at a unique point W_i, for every $i \in \{0,\ldots,n\}$. Then $W_0 = A'$ and $W_n = B'$. Moreover, since all segments of the form $\overline{P_iP_{i+2}}$ cut $\mathcal{D}$ in two, we must have $W_i \neq W_{i+1}$. It follows that $W_0,\ldots,W_n$ are in order along the arc of $\partial\mathcal{D}$ in $\mathcal{R}^+$.

Step 3: a local calculation. It remains to show that, for $i \in \{1, \ldots, n\}$,

$$|W_{i-1}W_i|_{\partial\mathcal{D}} < |W_{i-1}P_i| + |P_iW_i|,$$

where $|W_{i-1}W_i|_{\partial\mathcal{D}}$ denotes the length of the arc between W_{i-1} and W_i along $\partial\mathcal{D}$ in $\mathcal{R}^+$. To this end, let us introduce some coordinates. Suppose that $P_i = (0,0) \in \mathbb{R}^2$, and the y-axis bisects $\angle W_{i-1}P_iW_i$, so that $W_{i-1} = (-a, ta)$ and $W_i = (b, tb)$ for some $a, b, t > 0$. The arc of $\partial\mathcal{D}$ in question is the graph of a convex function

$$f : [-a, b] \to \mathbb{R}_{\geqslant 0}$$

with $f(-a) = ta$ and $f(b) = tb$; this graph remains above the graph of

$$x \mapsto t|x|$$

for all $x \in [-a, b]$.

Then if f is differentiable, convexity says

$$-t \leqslant f'(a) \leqslant f'(x) \leqslant f(b) \leqslant t,$$

so that

$$|W_{i-1}W_i|_{\partial\mathcal{D}} = \int_{-a}^{b} \sqrt{1 + |f'(x)|^2}\, dx \leqslant (a+b)\sqrt{1+t^2} = |W_{i-1}P_i| + |P_iW_i|,$$

as needed.

If f is not differentiable, the same idea works by taking piecewise linear approximations to f. $\square$