Shortcuts for the circle

Computational Geometry 79 (2019) 37–54

Contents lists available at ScienceDirect

Computational Geometry: Theory and Applications

www.elsevier.com/locate/comgeo

Shortcuts for the circle ☆,☆☆

Sang Won Bae [a], Mark de Berg [b], Otfried Cheong [c,*], Joachim Gudmundsson [d], Christos Levcopoulos [e]

[a] Division of Computer Science and Engineering, Kyonggi University, Republic of Korea
[b] Department of Computer Science, TU Eindhoven, the Netherlands
[c] School of Computing, KAIST, Republic of Korea
[d] University of Sydney, Australia
[e] Lund University, Sweden

ARTICLE INFO

Article history:
Received 12 March 2018
Accepted 29 October 2018
Available online 24 January 2019

Keywords:
Geometric network
Graph augmentation
Graph diameter

ABSTRACT

Let C be the unit circle in $\mathbb{R}^2$. We can view C as a plane graph whose vertices are all the points on C, and the distance between any two points on C is the length of the smaller arc between them. We consider a graph augmentation problem on C, where we want to place $k \geqslant 1$ *shortcuts* on C such that the diameter of the resulting graph is minimized.
We analyze for each k with $1 \leqslant k \leqslant 7$ what the optimal set of shortcuts is. Interestingly, the minimum diameter one can obtain is not a strictly decreasing function of k. For example, with seven shortcuts one cannot obtain a smaller diameter than with six shortcuts. Finally, we prove that the optimal diameter is $2 + \Theta(1/k^{2/3})$ for any k.

1. Introduction

Graph augmentation problems have received considerable attention over the years. The goal in such problems is typically to add extra edges to a given graph G in order to improve some quality measure. One natural quality measure is the (vertex- or edge-)connectivity of G. This has led to work where one tries to find the minimum number of edges that can be added to the graph to ensure it is k-connected, for a desired value of k. Another natural measure is the diameter of G, that is, the maximum distance between any pair of vertices. The goal then becomes to reduce the diameter as much as possible by adding a given number of edges, or to achieve a given diameter with a small number of extra edges; see for example the papers by Erdös, Rényi, and Sós [1,2].

Chung and Garey [3] studied this problem for the special case where the original graph is the n-vertex cycle. They showed that if k edges are added, then the diameter of the resulting graph is at least $\frac{n}{k+2} - 3$ for even k and $\frac{n}{k+1} - 3$ for odd k, and that there is a way to add k edges so that the resulting graph has diameter at most $\frac{n}{k+2} - 1$ for even k and $\frac{n}{k+1} - 1$ for odd k. (For paths, slightly better bounds are known [4].)

☆ An extended abstract of this work was published in the 28th International Symposium on Algorithms and Computation (ISAAC 2017).

☆☆ SWB was supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (2015R1D1A1A01057220 and 2018R1D1A1B07042755). MdB was supported by the Netherlands' Organisation for Scientific Research (NWO) under project no. 024.002.003. OC was supported by NRF grant 2011-0030044 (SRC-GAIA) funded by the government of Korea. JG was supported under Australian Research Council Discovery Projects funding scheme (project numbers DP150101134 and DP180102870).

* Corresponding author.
E-mail addresses: swbae@kgu.ac.kr (S.W. Bae), m.t.d.berg@tue.nl (M. de Berg), otfried@kaist.airpost.net (O. Cheong), joachim.gudmundsson@gmail.com (J. Gudmundsson), christos.levcopoulos@cs.lth.se (C. Levcopoulos).

https://doi.org/10.1016/j.comgeo.2019.01.006
0925-7721/

The algorithmic problem of finding a set of $k \geqslant 1$ edges that minimizes the diameter of the augmented graph was first asked by Chung [5] in 1987. Since then many papers have considered the problem for general graphs, see [6–9,4]. Große et al. [10] were the first to consider the diameter minimization problem in the geometric setting where the graph is embedded in the Euclidean plane. They presented an $O(n \log^3 n)$ time algorithm that determined the optimal shortcut that minimizes the diameter of a polygonal path with n vertices. The running time was later improved to $O(n \log n)$ by Wang [11].

In the above papers only the discrete setting is considered, that is, shortcuts connect two vertices and the diameter is measured between vertices. In the continuous setting all points along the edges of the network are taken into account when placing a shortcut and when measuring distances in the augmented network. In the continuous setting, Yang [12] studied the special case of adding a single shortcut to a polygonal path and gave several approximation algorithms for the problem. De Carufel et al. [13] considered the problem for paths and cycles. For paths they showed that an optimal shortcut can be determined in linear time. For cycles they showed that a single shortcut can never decrease the diameter, while two shortcuts always suffice. They also proved that for convex cycles the optimal pair of shortcuts can be computed in linear time. Recently, Cáceres et al. [14] gave a polynomial time algorithm that can determine whether a plane geometric network admits a reduction of the continuous diameter by adding a single shortcut.

We are interested in a geometric continuous variant of this problem. Let C be a unit circle in the plane. We define the *distance* $d(p,q)$ between two points $p,q \in C$ to be the length of the smaller arc along C that connects p to q. Thus the diameter of C in this metric is π. We now want to add a number of *shortcuts*—a shortcut is a chord of C—to improve the diameter. Here the distance $d_S(p,q)$ between p and q for a given collection S of shortcuts is defined as the length of the shortest path between p and q that can travel along C and along the shortcuts where, if two shortcuts intersect in their interior, we do not allow the path to switch from one shortcut to the other at the intersection point. In other words, if the path uses a shortcut, it has to traverse it completely. Note that if we view the original circle C as a graph with infinitely many vertices (namely all points on C) where the graph distance is the distance along C, then adding shortcuts corresponds to adding edges to the graph. For a set S of shortcuts, define $\operatorname{diam}(S) := \max_{p,q \in C} d_S(p,q)$ to be the diameter of the resulting "graph." We are interested in the following question: given k, the number of shortcuts we are allowed to add, what is the best diameter we can achieve? In other words, we are interested in the quantity $\operatorname{diam}(k) := \inf_{|S|=k} \operatorname{diam}(S)$.

It is obvious that $\pi = \operatorname{diam}(0) \geqslant \operatorname{diam}(1) \geqslant \cdots \geqslant \operatorname{diam}(k) \geqslant \cdots \geqslant \lim_{k\to\infty} \operatorname{diam}(k) = 2$.

Our main results are as follows.

- For $1 \leqslant k \leqslant 7$, we determine $\operatorname{diam}(k)$ exactly. Our results show that $\operatorname{diam}(k)$ is not strictly decreasing as a function of k. This not only holds at the very beginning—it is easy to see that $\operatorname{diam}(1) = \operatorname{diam}(0)$—but, interestingly also for certain larger values of k. In particular, we show that $\operatorname{diam}(7) = \operatorname{diam}(6)$.
- We have $\operatorname{diam}(8) < \operatorname{diam}(7)$.
- We show that $\operatorname{diam}(k) = 2 + \Theta(1/k^{\frac{2}{3}})$.

We rely on a number of numerical calculations. A Python script that performs these calculations can be found at http://github.com/otfried/circle-shortcuts, and the output of the script is included as Appendix A.

2. The umbra and the region of a shortcut

A shortcut s is a chord of C. A shortcut of length $a = |s| \in [0,2]$ spans an angle of $\alpha(a) \in [0,\pi]$, where $\alpha(a) := 2\arcsin\left(\frac{a}{2}\right)$. The following function $\delta : [0,2] \mapsto [0, \pi/2 - 1]$ will play a key role in our arguments:

$$\delta(a) := \frac{\alpha(a) - a}{2} = \arcsin\left(\frac{a}{2}\right) - \frac{a}{2}.$$

Note that both $\alpha(a)$ and $\delta(a)$ are increasing and convex functions, and $\alpha(a) = a + 2\delta(a)$. See Fig. 1. To simplify the notation, we will allow shortcuts themselves as the function argument, with the understanding that $\alpha(s) = \alpha(|s|)$ and $\delta(s) = \delta(|s|)$.

We parameterize the points on the circle C using their polar angle in $[0, 2\pi)$. For a shortcut s with endpoints u and v we will write $s = uv$ if the *counter-clockwise* arc $\widehat{uv}$ is the *shorter* arc of C connecting u and v. Only for $|s| = 2$, we have $s = uv = vu$; in this case u and v are *antipodal points*, that is $v = u + \pi$.

The *inner umbra* of a shortcut $s = uv$ is the arc $\widehat{u_1 v_1}$ where $u_1 = u + \delta(s)$ and $v_1 = v - \delta(s)$. The *outer umbra* is the set of antipodal points of the inner umbra, that is the arc $\widehat{u_1' v_1'}$ where $x' = x + \pi$. Together they form the *umbra* $U(s)$ of s. Since $\alpha(s) = |s| + 2\delta(s)$, the inner and outer umbra have length $|s|$. The *radiance* of s consists of the two arcs $\widehat{vu'}$ and $\widehat{v'u}$. For $|s| = 2$, we cannot distinguish inner and outer umbra, and the radiance consists of two isolated points, see Fig. 1.

Let $p \in U(s)$. Then a path going from p to one endpoint of s and traversing the shortcut is at least as long as going directly from p to the other endpoint—so the shortcut is not useful. This gives us the following observation:

Observation 1. *Given a set S of shortcuts, if the shortest path γ from p to q uses shortcuts $s_1, s_2, \ldots, s_m \in S$ in this order, then $p \notin U(s_1)$ and $q \notin U(s_m)$. Let u_i and v_i be the endpoints of s_i such that γ traverses s_i from u_i to v_i. Then $v_i \notin U(s_{i+1})$ and $u_{i+1} \notin U(s_i)$ for $i = 1, \ldots, m-1$.*

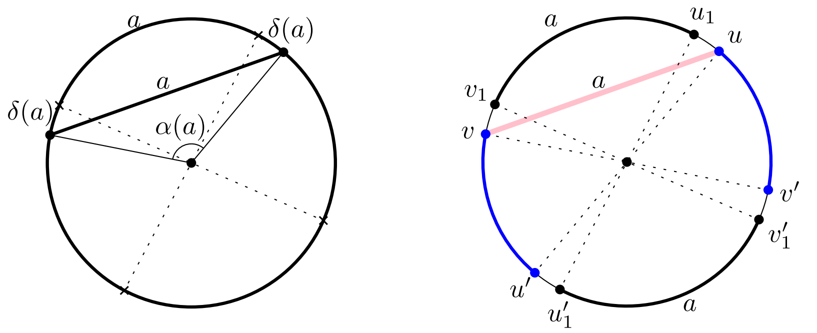

Fig. 1. For a shortcut s of length $a = |s|$, (left) $\alpha(a) = a + 2\delta(a)$ and (right) the umbra $U(s)$ (consisting of two arcs of length a in thick black) and radiance (in thick blue). (For interpretation of the colors in this and following figures, the reader is referred to the web version of this article.)

(For the boundary cases, we will assume that a shortest path uses the minimum number of shortcuts possible.) An immediate implication is that one shortcut alone cannot help to improve the diameter, that is, $\mathrm{diam}(1) = \mathrm{diam}(0) = \pi$.

Another useful observation is the following (remember that $d(p,q) = \min(|\widehat{pq}|, |\widehat{qp}|)$ is the distance along C without shortcuts):

Observation 2. *Given a set S of shortcuts, if the shortest path from p to q uses the set of shortcuts $\{s_1, s_2, \ldots, s_m\} \subseteq S$, then $d_S(p,q) \geqslant d(p,q) - 2\sum_{i=1}^{m} \delta(s_i)$.*

Indeed, if γ is the shortest path, we can replace each shortcut s_i by walking along the circle instead, increasing the path length by exactly $2\delta(s_i)$.

Observation 1 does not exclude the possibility that the shortest path for $p, q \in U(s)$ uses s as an intermediate shortcut (not the first or last one). We therefore define the *deep umbra* $\hat{U}(s)$ of s as the set of points p in the inner umbra of s such that $|\overline{pu}| + |s| \geqslant d(p,v)$, where u is the endpoint of s closer to p and v is the other endpoint.

Observation 3. *If $p \in \hat{U}(s)$ or $q \in \hat{U}(s)$ then the shortest path from p to q does not use shortcut s.*

The following lemma shows that the deep umbra covers nearly the entire inner umbra.

Lemma 4. *The deep umbra $\hat{U}(s)$ of a shortcut s is non-empty and has length at least $|s| - 4\delta(\delta(s)) > |s| - 0.02$.*

Proof. Let $s = uv$ and let $u_2 = u + \alpha(\delta(s))$ and $v_2 = v - \alpha(\delta(s))$. Since $\delta(s) < \alpha(\delta(s)) < \alpha(s)/2$, the arc $\widehat{u_2 v_2}$ is non-empty and lies inside the inner umbra of s.

We claim that $\widehat{u_2 v_2} \subset \hat{U}(s)$. Indeed, for $p \in \widehat{u_2 v_2}$ such that $d(p,u) < d(p,v)$, we have $|\overline{pu}| \geqslant |\overline{u_2 u}| = \delta(s)$, and so $d(p,v) = |\widehat{pv}| \leqslant |\widehat{u_2 v}| = \alpha(s) - \alpha(\delta(s)) < |s| + \delta(s) \leqslant |s| + |\overline{pu}|$.

The arc length of $\widehat{u_2 v_2} \subset \hat{U}(s)$ is $\alpha(s) - 2\alpha(\delta(s))$. Since $\alpha(x) = x + 2\delta(x)$, this is $|s| + 2\delta(s) - 2(\delta(s) + 2\delta(\delta(s))) = |s| - 4\delta(\delta(s))$.

Finally, we observe that the function $x \mapsto \delta(\delta(x))$ is increasing and $\delta(\delta(2)) < 0.005$. □

Let us now fix a target diameter of the form $\pi - \delta^*$, for some $\delta^* \in [0, \pi - 2]$. To achieve the target diameter, pairs of points $p, q \in C$ that span an angle of at most $\pi - \delta^*$ do not need a shortcut, so it suffices to consider pairs of points $p, q \in C$ where $q = p + \pi + \xi$, for $-\delta^* \leqslant \xi \leqslant \delta^*$. We represent these point pairs by the rectangle $\mathfrak{S}(\delta^*) = [0, 2\pi] \times [-\delta^*, +\delta^*]$, where (θ, ξ) corresponds to the pair of points $p = \theta - \xi/2$ and $q = \theta + \pi + \xi/2$, as illustrated in Fig. 2. So the counter-clockwise angle from p to q is $\pi + \xi$.

$\mathfrak{S}(\delta^*)$ is topologically a cylinder: the right edge $\theta = 2\pi$ is identified with the left edge $\theta = 0$. Furthermore, if the point pair $(p,q) \in C \times C$ corresponds to (θ, ξ), then the point pair (q,p) corresponds to $(\theta + \pi, -\xi)$. Since $d_S(p,q) = d_S(q,p)$, we could therefore identify the middle segment $\theta = \pi$ with the left edge $\theta = 0$, but with opposite orientation, resulting in a Möbius strip topology. As will become clear shortly, for our purposes it is easier to work with the cylinder topology, but keep in mind that, for instance, the upper boundary $\xi = \delta^*$ and the lower boundary $\xi = -\delta^*$ of $\mathfrak{S}(\delta^*)$ really represent the same point pairs.

For a shortcut s, we define the region $\mathfrak{R}(s, \delta^*) \subset \mathfrak{S}(\delta^*)$ consisting of those pairs $(\theta, \xi) \in \mathfrak{S}$ where $d_s(\theta - \xi/2, \theta + \pi + \xi/2) \leqslant \pi - \delta^*$. (To simplify the notation, we will use $d_s(p,q)$ for $d_{\{s\}}(p,q)$.)

Let us fix a shortcut s of length $a > 0$, and let $\alpha = \alpha(a)$ and $\delta = \delta(a)$. Rotating a shortcut around the origin means translating $\mathfrak{R}(s, \delta^*)$ horizontally in (the cylinder) $\mathfrak{S}(\delta^*)$. We can thus choose s to be vertical and connect the points $-\alpha/2$

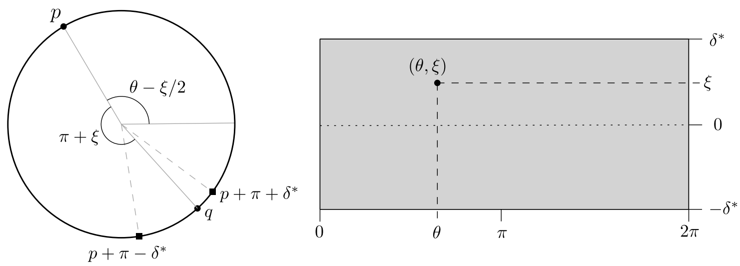

Fig. 2. (left) (θ, ξ) corresponds to the pair of points $p = \theta - \xi/2$ and $q = \theta + \pi + \xi/2$. (right) $\mathfrak{S}(\delta^*)$ represents all pair of points $p = \theta - \xi/2$ and $q = \theta + \pi + \xi/2$ with $-\delta^* \leqslant \xi \leqslant \delta^*$.

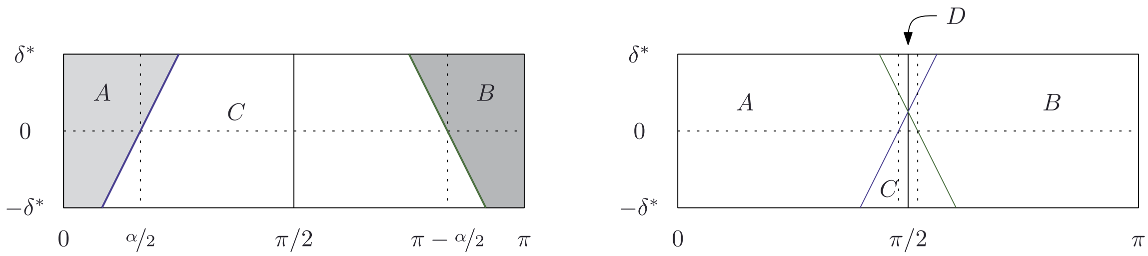

Fig. 3. The regions A, B, C, D.

and $\alpha/2$. This implies that the umbra of s consists of the two intervals $[-\alpha/2 + \delta, \alpha/2 - \delta]$ and $[\pi - \alpha/2 + \delta, \pi + \alpha/2 - \delta]$. The radiance of s consists of the two intervals $[\alpha/2, \pi - \alpha/2]$ and $[\pi + \alpha/2, 2\pi - \alpha/2]$.

The following function gives the length of the path from p to q that uses the shortcut s from top to bottom, that is, from the point $\alpha/2$ to $-\alpha/2$:

$$f(\theta, \xi) := |\alpha/2 - p| + a + |q - (2\pi - \alpha/2)|, \qquad \text{where} \quad (p, q) = (\theta - \xi/2, \theta + \pi + \xi/2).$$

By the observation about the Möbius strip topology above, it suffices to understand $\mathfrak{R}(s, \delta^*)$ for $0 \leqslant \theta \leqslant \pi$. We claim that for $0 \leqslant \theta \leqslant \pi$ we have $d_s(p, q) < \pi - \delta^*$ if and only if $f(\theta, \xi) < \pi - \delta^*$.

This is clearly true if the shortest path from p to q uses s from top to bottom, or not at all, because the length of the shorter circle arc between p and q is $\pi - |\xi| \geqslant \pi - \delta^*$. It remains to consider the case when the shortest path uses s from bottom to top. This can only happen when p is closer to the bottom end of s than to its top end—in other words, when $\pi < p < 2\pi$. Since $0 \leqslant \theta \leqslant \pi$ and $p = \theta - \xi/2$, this implies either $\theta < \delta^*/2$ and $\xi > 2\theta$, or $\theta > \pi - \delta^*/2$ and $\xi < -2(\pi - \theta)$. Since $q = \theta + \pi + \xi/2$, the first case implies $\pi \leqslant q \leqslant \pi + \delta^* < 2\pi$, while the second case implies $\pi < 2\pi - \delta^* \leqslant q \leqslant 2\pi$. In both cases, q lies closer to the bottom end of the shortcut than to its top end, a contradiction to the shortcut being used from bottom to top to go from p to q.

It follows that for $0 \leqslant \theta \leqslant \pi$, we have $(\theta, \xi) \in \mathfrak{R}(s, \delta^*)$ if and only if $f(\theta, \xi) \leqslant \pi - \delta^*$. To analyze f, we partition the rectangle $[0, \pi] \times [-\delta^*, \delta^*]$ into regions, depending on the signs of $\alpha/2 - p$ and $q - (2\pi - \alpha/2)$. First, we have $p < \alpha/2$ if and only if $\xi > 2\theta - \alpha$. This is the lightly shaded region A above the blue line in Fig. 3(left). Second, we have $q > 2\pi - \alpha/2$ if and only if $\xi > 2\pi - \alpha - 2\theta$. This is the darker region B above the green line in Fig. 3(left). If the two regions do not intersect then we get three regions as shown in Fig. 3(left). Otherwise, if $\alpha > \pi - \delta^*$, or equivalently, $\delta^* > \pi - a - 2\delta$, then the regions intersect and we get four regions as illustrated in Fig. 3(right). We now study $\mathfrak{R}(s, \delta^*)$ independently for each of the three or four regions.

In region A, we have $p < \alpha/2$ and $q < 2\pi - \alpha/2$. It follows that

$$\begin{aligned} f(\theta, \xi) &= \alpha/2 - p + \alpha - 2\delta + 2\pi - \alpha/2 - q \\ &= -\theta + \xi/2 + \alpha - 2\delta + 2\pi - \theta - \pi - \xi/2 \\ &= \pi - 2\delta + 2(\alpha/2 - \theta). \end{aligned}$$

This implies that $f(\theta, \xi) \leqslant \pi - \delta^*$ if and only if $\theta \geqslant \alpha/2 + \delta^*/2 - \delta = a/2 + \delta^*/2$. This is the blue area as shown in Fig. 4.

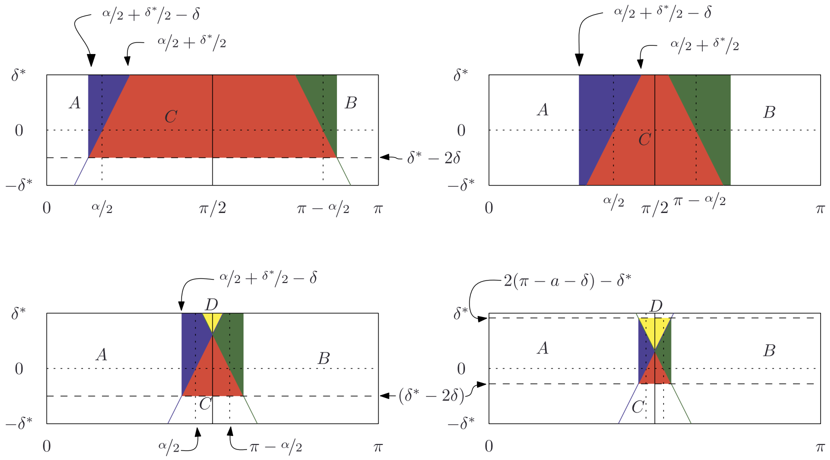

Fig. 4. The region $\mathfrak{R}(s, \delta^*)$ in four different situations.

In region B, we have $p \geqslant \alpha/2$ and $q > 2\pi - \alpha/2$. This implies

$$\begin{aligned} f(\theta,\xi) &= p - \alpha/2 + \alpha - 2\delta + q - 2\pi + \alpha/2 \\ &= \theta - \xi/2 + \alpha - 2\delta + \theta + \pi + \xi/2 - 2\pi \\ &= 2(\theta - (\pi - \alpha/2)) + \pi - 2\delta, \end{aligned}$$

and so we have $f(\theta,\xi) \leqslant \pi - \delta^*$ if and only if $\theta \leqslant \pi - \alpha/2 - \delta^*/2 + \delta$. This is the green area in Fig. 4.

Next, in region C, we have $p \geqslant \alpha/2$ and $q \leqslant 2\pi - \alpha/2$. Therefore,

$$\begin{aligned} f(\theta,\xi) &= p - \alpha/2 + \alpha - 2\delta + 2\pi - \alpha/2 - q \\ &= \theta - \xi/2 - 2\delta + 2\pi - \theta - \pi - \xi/2 \\ &= \pi - 2\delta - \xi. \end{aligned}$$

We have $f(\theta,\xi) \leqslant \pi - \delta^*$ if and only if $\xi \geqslant \delta^* - 2\delta$. This is the red area in Fig. 4.

When $\alpha > \pi - \delta^*$ regions A and B intersect in region D, as shown in Fig. 3(right). In region D we have $p < \alpha/2$ and $q > 2\pi - \alpha/2$, and therefore

$$\begin{aligned} f(\theta,\xi) &= \alpha/2 - p + \alpha - 2\delta + q - 2\pi + \alpha/2 \\ &= 2\alpha - 2\delta - 2\pi - \theta + \xi/2 + \theta + \pi + \xi/2 \\ &= 2(\alpha - \delta) - \pi + \xi \\ &= 2(a + \delta) - \pi + \xi, \end{aligned}$$

since $\alpha - \delta = a + \delta$. Thus, we have $f(\theta,\xi) \leqslant \pi - \delta^*$ if and only if $\xi \leqslant 2(\pi - a - \delta) - \delta^*$. This is the yellow area in region D in Fig. 4. There are two cases that can occur, as is shown on the bottom left and bottom right of Fig. 4. We postpone the discussion of these cases to the proof of the following lemma, which summarizes our discussions above.

Lemma 5. *Let $\delta^* \in [0, \pi - 2]$, and let s be a shortcut of length $a \in (0, 2]$. Then, the region $\mathfrak{R}(s, \delta^*)$ of s in the cylinder $\mathfrak{S}(\delta^*) = [0, 2\pi] \times [-\delta^*, +\delta^*]$ forms two identical rectangles whose width is exactly $\pi - a - \delta^*$ and whose height is*

$$\begin{cases} 2\delta^* & \text{if } \delta^* \leqslant \delta(a) \\ 2\delta(a) & \text{if } \delta^* > \delta(a) \text{ and } \delta^* \leqslant \pi - a - \delta(a) \\ 2(\pi - a - \delta^*) & \text{otherwise.} \end{cases}$$

Proof. We consider each case separately. Let $\alpha = \alpha(a)$ and $\delta = \delta(a)$.

First assume that $\delta^* \leqslant \delta$. Since $a \leqslant 2$ and $\delta \leqslant \pi/2 - 1$, we have

$$\pi - a - \delta \geqslant \pi/2 - 1 = \delta(2) \geqslant \delta \geqslant \delta^*.$$

Thus, we have $\delta^* \leqslant \pi - a - \delta$ and $2(\pi - a - \delta) - \delta^* \geqslant \delta^*$. This implies that the region $\mathfrak{R}(s, \delta^*)$ contains the whole region D if D is nonempty. In this case, $\mathfrak{R}(s, \delta^*)$ forms two identical rectangles that span the entire height $2\delta^*$ of $\mathfrak{S}(\delta^*)$, as shown in Fig. 4(top right). Each rectangle has width

$$(\pi - \alpha/2 - \delta^*/2 + \delta) - (\alpha/2 + \delta^*/2 - \delta) = \pi - \alpha + 2\delta - \delta^* = \pi - a - \delta^*,$$

since $\mathfrak{R}(s, \delta^*)$ spans the θ-interval $\alpha/2 + \delta^*/2 - \delta \leqslant \theta \leqslant \pi - \alpha/2 - \delta^*/2 + \delta$.

Second, suppose that $\delta^* > \delta$ and $\delta^* \leqslant \pi - a - \delta$. Then, we again have $2(\pi - a - \delta) - \delta^* \geqslant \delta^*$ and the region $\mathfrak{R}(s, \delta^*)$ contains the whole region D if D is nonempty. In this case, $\mathfrak{R}(s, \delta^*)$ forms two rectangles in $\mathfrak{S}(\delta^*)$. One touches the top boundary of $\mathfrak{S}(\delta^*)$ (as shown in Fig. 4(top left, bottom left)), the other one the bottom boundary. The height of the rectangle is exactly 2δ as $\mathfrak{R}(s, \delta^*)$ in region C is delimited by $\xi \geqslant \delta^* - 2\delta$, while the width of the rectangle is $\pi - a - \delta^*$ as above.

Finally, we consider the remaining case where $\delta^* > \pi - a - \delta$. In this case D must be nonempty, as $\delta^* > \pi - a - 2\delta = \pi - \alpha$, and the region $\mathfrak{R}(s, \delta^*)$ in region D is delimited by $\xi \leqslant 2(\pi - a - \delta) - \delta^*$. Since $\delta^* > \pi - a - \delta$, we have a strict inequality $2(\pi - a - \delta) - \delta^* < \delta^*$. Fig. 4(bottom right) illustrates the situation. Observe on one hand that the horizontal width and the vertical height of $\mathfrak{R}(s, \delta^*) \cap D$ (the yellow area in the figure) are exactly $\pi - a - \delta^*$. On the other hand, the width and the height of $\mathfrak{R}(s, \delta^*) \cap C$ (the red area in the figure) is also equal to $\pi - a - \delta^*$ since $(2(\pi - a - \delta) - \delta^*) - (\delta^* - 2\delta) = 2(\pi - a - \delta^*)$. Thus, $\mathfrak{R}(s, \delta^*) \cap (C \cup D)$ is of width $\pi - a - \delta^*$ and thus fits in between $\alpha/2 + \delta^*/2 - \delta \leqslant \theta \leqslant \pi - \alpha/2 - \delta^*/2 + \delta$. So, the region $\mathfrak{R}(s, \delta^*)$ of s again forms two rectangles in $\mathfrak{S}(\delta^*)$. Neither of them touches a boundary of $\mathfrak{S}(\delta^*)$, both have width $\pi - a - \delta^*$ and height $2(\pi - a - \delta^*)$. □

Note that if $\delta^* \leqslant \delta(2) = \pi/2 - 1$, then it always holds that $\delta^* \leqslant \pi - a - \delta(a)$ for any $0 \leqslant a \leqslant 2$ since $\pi - a - \delta(a) \geqslant \pi - 2 - \delta(2) = \delta(2) \geqslant \delta^*$. Hence, the last case of Lemma 5 where $\delta^* > \pi - a - \delta(a)$ only happens when $\delta^* > \delta(2) = \pi/2 - 1$.

Let $\mathfrak{M} = \{\xi = 0\}$ be the middle line of $\mathfrak{S}(\delta^*)$, and let $\mathfrak{B} = \{\xi = \delta^*\}$ be its upper boundary. We will also be interested in the length of the intersections $\mathfrak{M} \cap \mathfrak{R}(s, \delta^*)$ and $\mathfrak{B} \cap \mathfrak{R}(s, \delta^*)$. Note that both $\mathfrak{M}$ and $\mathfrak{B}$ have length 2π. We have the following corollary to Lemma 5:

Corollary 6. *Let $\delta^* \in [0, \pi - 2]$, and let s be a shortcut of length $a \in (0, 2]$. Then*

$$|\mathfrak{M} \cap \mathfrak{R}(s, \delta^*)| = \begin{cases} 2(\pi - a - \delta^*) & \text{if } \delta(a) \geqslant \delta^*/2 \\ 0 & \text{otherwise.} \end{cases}$$

and

$$|\mathfrak{B} \cap \mathfrak{R}(s, \delta^*)| = \begin{cases} 2(\pi - a - \delta^*) & \text{if } \delta(a) \geqslant \delta^* \\ \pi - a - \delta^* & \text{if } \delta(a) < \delta^* \leqslant \pi - a - \delta(a) \\ 0 & \text{otherwise.} \end{cases}$$

3. Up to five shortcuts

In this section we derive the exact value of diam(k) for $k \in \{2, 3, 4, 5\}$ and the unique optimal configuration of shortcuts in each case. The proof is quite easy, comparing the areas of $\mathfrak{R}(s, \delta^*)$ with the area of $\mathfrak{S}(\delta^*)$, if one assumes that the shortest path between any pair of points uses at most one shortcut. Showing that using a combination of shortcuts does not help takes considerable additional effort.

3.1. Using only one shortcut

Again we consider a target diameter of the form $\pi - \delta^*$, with $\delta^* \in [0, \pi - 2]$. By Lemma 5, the region $\mathfrak{R}(s, \delta^*)$ of a shortcut s of length a consists of two rectangles of width $\pi - a - \delta^*$ and height $2\delta(a)$ for $\delta(a) < \delta^*$, and height $2\delta^*$ for $\delta(a) \geqslant \delta^*$. We define a^* such that $\delta(a^*) = \delta^*$, or $a^* = 2$ when $\delta^* > \delta(2)$.

Then the area $A(a, \delta^*)$ of $\mathfrak{R}(s, \delta^*)$ is

$$A(a, \delta^*) = \begin{cases} 4\delta^*(\pi - a - \delta^*) & \text{for} \quad a > a^* \\ 4\delta(a)(\pi - a - \delta^*) & \text{for} \quad a \leqslant a^* \end{cases}$$

Lemma 7. *For fixed $\delta^* \leqslant 0.7$, the function $a \mapsto A(a, \delta^*)$ is increasing for $a \leqslant a^*$ and decreasing for $a \geqslant a^*$. Its maximum value is $A(a^*, \delta^*) = 4\delta^*(\pi - a^* - \delta^*)$.*

Proof. For $a \geqslant a^*$, the function $a \mapsto A(a, \delta^*)$ is a decreasing linear function. To verify that $A(a, \delta^*)$ is increasing for $a \leqslant a^*$, we consider the derivative $\frac{d}{da}(\delta(a)(\pi - a - \delta^*))$ of the function $a \mapsto \delta(a)(\pi - a - \delta^*)$ for any fixed δ^* with $0 < \delta^* \leqslant 0.7$. We have

Table 1
The values a_k^*, δ_k^*, $\pi - \delta_k^*$, and μ_k.

k	a_k^*	δ_k^*	$\operatorname{diam}(S) = \pi - \delta_k^*$	μ_k
2	1.4782	0.0926	3.0490	1.2219
3	1.8435	0.2509	2.8907	1.5943
4	1.9619	0.3943	2.7473	1.7623
5	1.9969	0.5164	2.6252	1.8526
6	2.0000	0.5708	2.5708	1.8828

$$\begin{aligned}\frac{d}{da}(\delta(a)(\pi - a - \delta^*)) &= (\pi - a - \delta^*)\Big(\frac{1}{\sqrt{4-a^2}} - \frac{1}{2}\Big) - \delta(a)\\ &\geqslant (2.4 - a)\Big(\frac{1}{\sqrt{4-a^2}} - \frac{1}{2}\Big) - \delta(a)\\ &= \frac{2.4-a}{\sqrt{4-a^2}} + a - 1.2 - \arcsin\big(\frac{a}{2}\big) = g\big(\frac{a}{2}\big),\end{aligned}$$

where we define

$$g(x) = \frac{1.2 - x}{\sqrt{1-x^2}} + 2x - 1.2 - \arcsin x.$$

We will prove the lemma by showing that $g(x) > 0$ for $0 < x < 1$. Consider again the derivative:

$$g'(x) = \frac{1}{(1-x^2)^{3/2}}\big(x^2 + 1.2x - 2 + 2(1-x^2)^{3/2}\big).$$

Set $h(x) = x^2 + 1.2x - 2 + 2(1-x^2)^{3/2}$. The function $h(x)$ is continuous on the interval $[0, 1]$ and has only a single real root at $x = 0$, so $h(1/2) > 0$ implies $h(x) > 0$ for $0 < x < 1$. This implies that $g'(x) > 0$ for $0 < x < 1$, so $g(0) = 0$ implies $g(x) > 0$ for $0 < x < 1$, completing the proof. □

Let $k \in \{2, 3, 4, 5\}$. Since $a \mapsto a + \delta(a)$ is an increasing function that maps $[0, 2]$ to $[0, \pi/2 + 1]$, there is a unique a_k^* that solves the equation

$$a_k^* + \delta(a_k^*) = \frac{k-1}{k}\pi.$$

We set $\delta_k^* := \delta(a_k^*)$, and will show that this number determines the optimal diameter for k shortcuts. Table 1 shows the numerical values. For completeness, we already include the case $k = 6$ in the table by setting $a_6^* = 2$.

Lemma 8. *For* $k \in \{2, 3, 4, 5\}$ *there is a set S of k shortcuts that achieves* $\operatorname{diam}(S) = \pi - \delta_k^*$.
Assuming that no pair of points uses more than one shortcut, this is optimal and the solution is unique up to rotation.

Proof. By Lemma 5, the region $\mathfrak{R}(s, \delta_k^*)$ of a shortcut s of length $|s| = a_k^*$ consists of two rectangles of height $2\delta_k^*$ and width $\pi - (a_k^* + \delta_k^*) = \pi/k$. Each rectangle covers the entire height of $\mathfrak{S}(\delta^*)$, and by rotating s about the origin we can translate the rectangles anywhere inside $\mathfrak{S}(\delta^*)$. This implies that we can use k such rectangles to cover the range $0 \leqslant \theta \leqslant \pi$. Then for every $(\theta, \xi) \in \mathfrak{S}(\delta^*)$ there is a shortcut s such that $d_s(\theta - \xi/2, \theta + \pi + \xi/2) \leqslant \pi - \delta_k^*$, and $\operatorname{diam}(S) = \pi - \delta_k^*$. Fig. 5 shows the resulting configurations.

Assume now that a set $S = \{s_1, \dots, s_k\}$ of k shortcuts is given with $\operatorname{diam}(S) \leqslant \pi - \delta^*$, where $\delta^* \geqslant \delta_k^*$, and that no pair of points uses more than one shortcut. This implies that the regions $\mathfrak{R}(s_i, \delta^*)$ must entirely cover the strip $\mathfrak{S}(\delta^*)$, and in particular

$$\sum_{i=1}^{k} A(|s_i|, \delta^*) \geqslant 4\delta^*\pi.$$

If we choose a^* such that $\delta(a^*) = \delta^*$, then $a^* \geqslant a_k^*$. By Lemma 7 we have

$$A(|s_i|, \delta^*) \leqslant A(a^*, \delta^*) = 4\delta^*(\pi - a^* - \delta^*).$$

From $kA(a^*, \delta^*) \geqslant 4\delta^*\pi$ we have $k(\pi - a^* - \delta^*) \geqslant \pi$, or $a^* + \delta^* \leqslant \frac{k-1}{k}\pi$, which implies $a^* = a_k^*$ and $\delta^* = \delta_k^*$. But then the regions $\mathfrak{R}(s_i, \delta_k^*)$ must be non-overlapping, and the solution is unique up to rotation. □

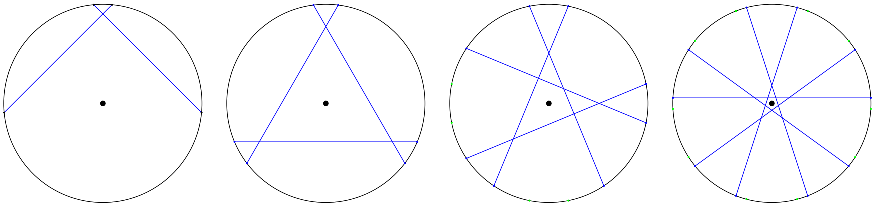

Fig. 5. The optimal shortcut configurations for $k = 2, 3, 4, 5$.

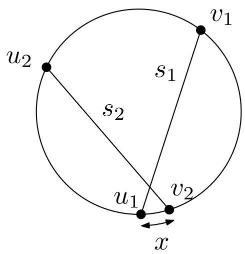

Fig. 6. The two shortcuts must intersect.

It remains to show that the configurations in Fig. 5 are optimal even if combinations of shortcuts can be used. The case $k = 2$ is somewhat special and handled first.

We start by defining $\mu_k \in [0, 2]$ to be such that $\delta(\mu_k) = \delta_k^*/2$. Table 1 shows the numerical values. By Lemma 5, $\mathfrak{R}(s, \delta_k^*)$ intersects the middle line $\mathfrak{M}$ if and only if $|s| \geqslant \mu_k$. In other words, for two antipodal points p and q we can make $d_s(p, q) \leqslant \pi - \delta_k^*$ only if $|s| \geqslant \mu_k$.

3.2. Optimality for two shortcuts

Lemma 9. *If S is a set of two shortcuts that achieves diameter $\operatorname{diam}(S) \leqslant \pi - \delta_2^*$, then S is identical to the configuration of Fig. 5 up to rotation.*

Proof. Let $S = \{s_1, s_2\}$ with $|s_1| \leqslant |s_2|$. Let p and q be the midpoints of the inner and outer umbra of s_2. The shortest path between p and q cannot use s_2 at all by Observation 1, so $d_{s_1}(p, q) \leqslant \pi - \delta_2^*$. This implies $|s_1| \geqslant \mu_2 \approx 1.2219$. Since $\delta_2^* \approx 0.0926 < \mu_2/2$, the interval $[q - \delta_2^*, q + \delta_2^*]$ lies in $U(s_2)$, and so we have $d_{s_1}(p, q') \leqslant \pi - \delta_2^*$ for all $q' \in [q - \delta_2^*, q + \delta_2^*]$. This implies $|s_1| \geqslant a_2^*$.

Next, we observe that $U(s_1) \cap U(s_2) = \emptyset$. Otherwise, Observation 1 applied to an antipodal pair in $U(s_1) \cap U(s_2)$ implies $\operatorname{diam}(S) = \pi$, a contradiction.

The two arcs between the inner and outer umbras of s_1 have length $\pi - |s_1| \leqslant \pi - a_2^*$. The inner umbra $U(s_2)$ has length $|s_2| \geqslant a_2^*$ and lies in one of these arcs. That leaves a gap of at most $\pi - 2a_2^* = 2\delta_2^*$ between the two inner umbras (by definition of a_2^*, we have $a_2^* + \delta_2^* = \pi/2$). Since $\delta(s_2) \geqslant \delta(s_1) \geqslant \delta_2^*$, this implies that the two shortcuts intersect, see Fig. 6.

Let x be the length of the overlap of the arcs of s_1 and s_2, that is, $x = |\widehat{u_1 v_2}|$ in Fig. 6. Any path that uses both s_1 and s_2 has length at least $|s_1| + |s_2| + x \geqslant 2a_2^* + x = \pi - 2\delta_2^* + x$. This is bounded by $\pi - \delta_2^*$ only if $x \leqslant \delta_2^*$. But then the arc $\widehat{v_1 u_2}$ has length at most

$$2\pi - \alpha(s_1) - \alpha(s_2) + x \leqslant 2\pi - 2(a_2^* + 2\delta_2^*) + \delta_2^* = \pi + (\pi - 2a_2^*) - 3\delta_2^* = \pi - \delta_2^*,$$

and there is no reason to use the two shortcuts at all. It follows that there is no pair of points that uses more than one shortcut, and Lemma 8 implies the claim. □

3.3. Antipodal pairs cannot use combinations of shortcuts

The key to the general proof for $3 \leqslant k \leqslant 6$ is the following lemma. We prove it for two separate cases: $k = 3$, and $k \in \{4, 5, 6\}$.

Lemma 10. *Let S be a set of k shortcuts for $k \in \{3, 4, 5, 6\}$ such that $\operatorname{diam}(S) \leqslant \pi - \delta_k^*$. Then there is no antipodal pair of points $p, q \in C$ such that the path of length $d_S(p, q)$ uses more than one shortcut.*

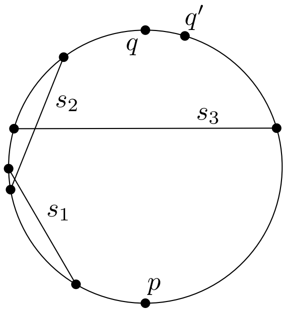

Fig. 7. Proof of Lemma 10 for k=3.

Table 2
Numeric values for σ_k, and λ_k.

k	a_k^*	δ_k^*	μ_k	σ_k	λ_k
4	1.9619	0.3943	1.7623	1.0373	1.7100
5	1.9969	0.5164	1.8526	0.7862	1.8390
6	2.0000	0.5707	1.8828	0.6958	1.8751

Proof of Lemma 10 for $k = 3$. Let $S = \{s_1, s_2, s_3\}$ with $|s_1| \leqslant |s_2| \leqslant |s_3|$. Again assume the opposite to the statement of the lemma, that is, assume that $\operatorname{diam}(S) \leqslant \pi - \delta_3^*$ and that there is an antipodal pair for which the shortest path between the two points uses at least two shortcuts.

We now show several properties of the configuration S.

(i) $|s_1| < 1.45 < \mu_3$: For two shortcuts to be a valid combination their combined length must be at most $\pi - \delta_3^*$, hence, $|s_1| \leqslant (\pi - \delta_3^*)/2 < 1.45 < \mu_3 \approx 1.5943$.
(ii) $U(s_2) \cap U(s_3) = \emptyset$: Otherwise we have an antipodal pair (p, q) with $p, q \in U(s_2) \cap U(s_3)$. Using Observation 1 and $|s_1| < \mu_3$ by (i) gives $d_S(p, q) = d_{s_1}(p, q) > \pi - \delta_3^*$.
(iii) $|s_2| < \pi/2 < \mu_3$: We have $|s_2| + |s_3| < \pi$ and thus $|s_2| < \pi/2 < \mu_3$ by (ii).
(iv) $U(s_1) \cap U(s_3) = \emptyset$: This follows using $|s_2| < \mu_3$ in (ii).
(v) $\delta(s_1) < 0.09$: From $|s_1| < 1.45$ by (i).
(vi) $\delta(s_2) < 0.12$: From $|s_2| < \pi/2$ by (iii).
(vii) $|s_3| > 1.32$: Pick an antipodal pair p, q with p in the deep umbra $\hat{U}(s_2)$. This is always possible by Lemma 4. By Observations 2 and 3, we have $d_S(p, q) = d_{\{s_1, s_3\}}(p, q) \geqslant \pi - 2(\delta(s_1) + \delta(s_3))$. This implies $\delta(s_1) + \delta(s_3) \geqslant \delta_3^*/2$. Therefore $\delta(s_3) \geqslant \delta_3^*/4 > 0.06$, implying that $|s_3| > 1.32$.

Now let p and q be the midpoints of the two umbras of s_3, that is, the inner and outer umbra. Without loss of generality assume that $\overline{pq}$ is vertical with p below q as shown in Fig. 7. Since the umbras of s_1, s_2 are disjoint from $U(s_3)$ and $\delta(s_1) \leqslant \delta(s_2) < 0.12 < |s_3|/2$ by (vi) and (vii), s_1 and s_2 do not cross $\overline{pq}$ and lie either in the left or right semicircle determined by $\overline{pq}$.

Since $|s_1| \leqslant |s_2| < \mu_3$ by (i) and (iii), we have $d_{s_1}(p, q) \geqslant d_{s_2}(p, q) > \pi - \delta_3^*$, and so the pair p, q must use s_1 and s_2 in combination. But this means that s_1 and s_2 lie in the same semicircle of C, let's say the left one as shown in Fig. 7.

Consider now the point $q' = q - \delta_3^* + \varepsilon$, for some small $\varepsilon > 0$. Traveling counter-clockwise from p to q' cannot use any shortcut and has length $\pi - \delta_3^* + \varepsilon > \pi - \delta_3^*$, traveling clockwise using s_1 and s_2 and arguing as in Observation 2 has distance at least $\pi + \delta_3^* - \varepsilon - 2\delta(s_1) - 2\delta(s_2) > \pi + 0.25 - 0.18 - 0.24 - \varepsilon > \pi - 0.25 > \pi - \delta_3^*$, which is a contradiction. $\square$

In order to handle the remaining case $k \in \{4, 5, 6\}$ we need a bound on the lengths of shortcuts that appear in combination.

Lemma 11. *Let $S' \subset S$ be the set of shortcuts used by the shortest path for an antipodal pair (p, q). If $d_{S'}(p, q) \leqslant \pi - \delta_k^*$, where $k \in \{4, 5, 6\}$, then the longest shortcut in S' has length at least λ_k, all others have total length at most σ_k. Here σ_k and λ_k, with $\sigma_k < \lambda_k$, are the two solutions to the equation $\delta(x) + \delta(\pi - \delta_k^* - x) = \delta_k^*/2$ for $x \in [\pi - \delta_k^* - 2, 2]$.*

Proof. The function $x \mapsto \delta(x)$ is increasing and strictly convex on the interval $[0, 2]$. Therefore the function $x \mapsto g(x) = \delta(x) + \delta(\pi - \delta_k^* - x)$ is strictly convex on the interval $[\pi - \delta_k^* - 2, 2]$, it is also symmetric about $x_0 = (\pi - \delta_k^*)/2$. Since $g(x_0) < \delta_k^*/2 < g(2)$, there are exactly two solutions σ_k and λ_k to the equation $g(x) = \delta_k^*/2$. Because of the symmetry of $g(x)$ we have $\sigma_k + \lambda_k = \pi - \delta_k^*$. We compute σ_k and λ_k numerically and list their values in Table 2.

Let $S' = \{s_1, \ldots, s_n\}$ with $|s_1| \geqslant |s_2| \geqslant \cdots \geqslant |s_n|$. If $n = 1$ the statement follows from $\mu_k > \lambda_k$, so assume $n \geqslant 2$.

We use Karamata's theorem. It states that if f is a strictly convex non-decreasing function on an interval I and $x_1 \geqslant x_2 \geqslant \cdots \geqslant x_n$ and $y_1 \geqslant y_2 \geqslant \ldots y_n$ are values in I such that $x_1 + \cdots + x_i \geqslant y_1 + \cdots + y_i$ for all $1 \leqslant i \leqslant n$, then $\sum_{i=1}^n f(x_i) \geqslant \sum_{i=1}^n f(y_i)$, and equality holds only if $x_i = y_i$ for all i.

We apply this with $f(x) = \delta(x)$ on the interval $I = [0, 2]$. We set $y_i = |s_i|$, $x_1 = \lambda_k$, $x_2 = \sigma_k$, and $x_i = 0$ for $i > 2$. Assume for a contradiction that $y_1 = |s_1| < \lambda_k$. Then the conditions of Karamata's theorem are satisfied: $x_1 = \lambda_k > y_1$ and $x_1 + x_2 = \lambda_k + \sigma_k = \pi - \delta_k^* \geqslant d_{S'}(p, q) \geqslant \sum_{i=1}^n |s_i| = \sum_{i=1}^n y_i$. We thus have $\sum_{i=1}^n \delta(s_i) < \delta(\lambda_k) + \delta(\sigma_k) = \delta_k^*/2$. By Observation 2 we then have $d_{S'}(p, q) \geqslant \pi - 2\sum_{i=1}^n \delta(s_i) > \pi - \delta_k^*$, a contradiction.

It follows that $|s_1| \geqslant \lambda_k$. Then $\sum_{i=1}^n |s_i| \leqslant \pi - \delta_k^*$ implies $\sum_{i=2}^n |s_i| \leqslant \pi - \delta_k^* - \lambda_k = \sigma_k$. □

Proof of Lemma 10 for $k \in \{4, 5, 6\}$. Let $S = \{s_1, s_2, \ldots, s_k\}$ with $|s_1| \leqslant |s_2| \leqslant \cdots \leqslant |s_k|$, and assume that some antipodal pair of points needs to use more than one shortcut. By Lemma 11 this means that $|s_1| \leqslant \sigma_k$.

Consider an antipodal pair p, q with $p \in \hat{U}(s_k)$. By Observation 3 it cannot use s_k, and needs either a single shortcut of length at least μ_k, or a shortcut combination whose longest element has length at least λ_k by Lemma 11. This implies that $|s_{k-1}| \geqslant \lambda_k$.

Since $|s_{k-1}| + |s_k| \geqslant 2\lambda_k > \pi + 0.2$, the intersection $U(s_{k-1}) \cap U(s_k)$ has total length at least 0.4. Since it consists of at most four arcs, one arc has length at least 0.1. Let p be the midpoint of this intersection arc, and q be its antipode. Since the interval $[p - 0.05, p + 0.05] \subset U(s_{k-1}) \cap U(s_k)$, Lemma 4 implies that either p or q lies in the deep umbra $\hat{U}(s_{k-1})$, and either p or q lies in $\hat{U}(s_k)$. By Observation 3 the pair (p, q) cannot use either s_{k-1} or s_k. As above we can now conclude that $|s_{k-2}| \geqslant \lambda_k$.

Let us call a shortcut s_i *short* if $|s_i| \leqslant \sigma_k$, and let ℓ be the number of short shortcuts. By the above, $1 \leqslant \ell \leqslant k - 3$. Let $S' = \{s_{\ell+1}, \ldots, s_k\}$ be the set of $k - \ell$ shortcuts that are not short.

We claim that $\mathrm{diam}(S') \leqslant \pi - \delta_k^* + 2\ell\delta(\sigma_k)$. Indeed, for any pair of points $p, q \in C$ there is a path from p to q using S of length at most $\pi - \delta_k^*$. Arguing as in Observation 2, we replace a short shortcut s_i by walking along the circle, and obtain a path γ of length at most $\pi - \delta_k^* + \ell \times 2\delta(\sigma_k)$. By Lemma 11, if p, q is an antipodal pair, then γ uses only one shortcut of S'.

Set $\hat{\delta} := \delta_k^* - 2(k-3)\delta(\sigma_k)$. Since $\ell \leqslant k - 3$, we have $\mathrm{diam}(S') \leqslant \pi - \hat{\delta}$. We consider the strip $\mathfrak{S}(\hat{\delta})$. The middle line $\mathfrak{M}$ of $\mathfrak{S}(\hat{\delta})$ corresponds to antipodal pairs $p, q \in C$. By the argument above, there is $s \in S'$ such that $d_s(p, q) \leqslant \pi - \hat{\delta}$. It follows that the regions $\mathfrak{R}(s, \hat{\delta})$, for $s \in S'$ and $|s| \geqslant \lambda_k$, cover $\mathfrak{M}$ entirely. The width of such a region is at most $\pi - \lambda_k - \hat{\delta}$, and so we must have $m \cdot (\pi - \lambda_k - \hat{\delta}) \geqslant \pi$, where m is the number of shortcuts in S' of length at least λ_k. Calculation shows that $m \geqslant k - 1$. Since $m \leqslant |S'| = k - \ell \leqslant k - 1$, this implies $\ell = 1$, and $|s_2| \geqslant \lambda_k$.

Since $2 \cdot \lambda_k > \pi$, no two shortcuts in S' can be combined, so $\mathfrak{S}(\hat{\delta})$ must be entirely covered by the regions $\mathfrak{R}(s_i, \hat{\delta})$, for $i \in \{2, 3, \ldots, k\}$. The area of $\mathfrak{S}(\hat{\delta})$ is $4\hat{\delta}\pi$. By Lemma 7, the area of $\mathfrak{R}(s_i, \hat{\delta})$ is at most $A(\hat{a}, \hat{\delta}) = 4\hat{\delta}(\pi - \hat{a} - \hat{\delta})$, where $\hat{a}$ is such that $\delta(\hat{a}) = \hat{\delta}$, and so we must have $(k-1) \times (\pi - \hat{a} - \hat{\delta}) \geqslant \pi$. However, calculation shows that this is false, contradicting our assumption that some pair of antipodal points uses more than one shortcut. □

3.4. Optimality of our configurations

It remains to show that the configurations of Fig. 5 are optimal even if combinations of shortcuts can be used. We prove this in the following lemma:

Lemma 12. *Let S be a set of k shortcuts for $k \in \{3, 4, 5\}$ such that* $\mathrm{diam}(S) \leqslant \pi - \delta_k^*$. *Then there is no pair of points $p, q \in C$ such that the path of length $d_S(p, q)$ uses more than one shortcut.*

Proof. By Lemma 10 pairs of antipodal points cannot use more than one shortcut. This implies that the middle line $\mathfrak{M}$ of $\mathfrak{S}(\delta_k^*)$ is covered by the regions $\mathfrak{R}(s_i, \delta_k^*)$. The region $\mathfrak{R}(s_i, \delta_k^*)$ intersects $\mathfrak{M}$ only if $|s_i| \geqslant \mu_k$, so by Corollary 6 $\mathfrak{R}(s_i, \delta_k^*)$ covers at most $2(\pi - \mu_k - \delta_k^*)$ of $\mathfrak{M}$. Calculation shows that $(k-1)(\pi - \mu_k - \delta_k^*) < \pi$, so all k shortcuts have length at least μ_k. Since $2\mu_k > \pi$, this implies that no shortcuts can be combined. □

Combining Lemmas 8, 9, and 12, we obtain our first theorem.

Theorem 13. *For $k \in \{2, 3, 4, 5\}$ there is a set S of k shortcuts that achieves* $\mathrm{diam}(S) = \pi - \delta_k^*$. *This is optimal and the solution is unique up to rotation.*

4. Six and seven shortcuts

The configuration of six shortcuts of length 2 (that is, all shortcuts are diameters of the circle) shown in Fig. 8 achieves diameter $\pi - \delta(2) = \pi/2 + 1$. Unlike the cases $2 \leqslant k \leqslant 5$, this configuration is not unique—it can be perturbed quite a bit without changing the diameter.

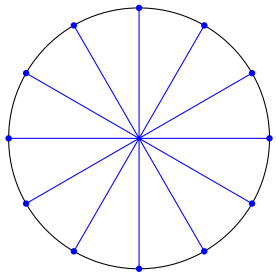

Fig. 8. An optimal configuration of six shortcuts.

It remains to argue that the configuration is indeed optimal, that is, there is no set S of six shortcuts that achieves $\mathrm{diam}(S) < \pi - \delta(2)$. Here, we cannot use a simple area argument as in the case $k < 6$, as the regions of the optimal solution in $\mathfrak{S}(\delta_6^*)$ overlap heavily.

In fact, we can show that even if we allow *seven* shortcuts, there is no set S of shortcuts that achieves $\mathrm{diam}(S) < \pi - \delta(2)$. This implies a collapse between the cases of $k = 6$ and $k = 7$, that is, $\mathrm{diam}(7) = \mathrm{diam}(6) = \pi - \delta(2)$.

Theorem 14. *There is a set S of six shortcuts that achieves* $\mathrm{diam}(S) = \pi - \delta(2) = \pi/2 + 1$. *There is no configuration of six or seven shortcuts that has diameter smaller than* $\pi/2 + 1$. *Therefore, we have* $\mathrm{diam}(7) = \mathrm{diam}(6) = \pi/2 + 1$.

4.1. A short proof for six shortcuts...

Since the proof for seven shortcuts is quite long and rather technical, we first give a short proof for the case of six shortcuts (even though this is of course implied by the proof for seven shortcuts).

Proof of Theorem 14 for six shortcuts. Let $S = \{s_1, \dots, s_6\}$ with $|s_1| \leqslant \cdots \leqslant |s_6|$, and assume that $\mathrm{diam}(S) \leqslant \pi - \delta^*$ with $\delta^* > \delta_6^*$.

By Lemma 10, this implies that no antipodal pair uses more than one shortcut. This means that the middle line $\mathfrak{M}$ of $\mathfrak{S}(\delta^*)$ is covered by the regions $\mathfrak{R}(s_i, \delta^*)$. The region $\mathfrak{R}(s_i, \delta^*)$ intersects $\mathfrak{M}$ only if $|s_i| \geqslant \mu_6$, so by Corollary 6 $\mathfrak{R}(s_i, \delta^*)$ covers at most $2(\pi - \mu_6 - \delta_6^*)$ of $\mathfrak{M}$. Calculation shows that $4(\pi - \mu_6 - \delta_6^*) < \pi$, so at least five shortcuts have length at least μ_6, that is $|s_2| \geqslant \mu_6$.

If no pair of points uses more than one shortcut, then the strip $\mathfrak{S}(\delta^*)$ must be covered by the six regions $\mathfrak{R}(s_i, \delta^*)$. In particular, the upper boundary $\mathfrak{B}$ of $\mathfrak{S}(\delta^*)$ is covered. Since $\delta^* > \delta_6^* = \delta(2) \geqslant \delta(a)$ for $a \in [0, 2]$, by Corollary 6 region $\mathfrak{R}(s_i, \delta^*)$ covers at most $\pi - |s_i| - \delta^*$ of $\mathfrak{B}$. This leads to a contradiction:

$$\sum_{i=1}^{6} (\pi - |s_i| - \delta^*) \leqslant (\pi - \delta_6^*) + 5(\pi - \mu_6 - \delta_6^*) = 6\pi - 6\delta_6^* - 5\mu_6 \approx 6.0106 < 2\pi.$$

It follows that some pair of points uses two shortcuts. Since $2\mu_6 > \pi$, this can only be a combination involving s_1, so we must have $|s_1| + |s_2| \leqslant \pi - \delta^*$. From $|s_2| \geqslant \mu_6$ we get $|s_1| \leqslant \pi - \delta^* - \mu_6$, and so $\delta(s_1) \leqslant \delta(\pi - \delta_6^* - \mu_6) < 0.008$. We set $S' = \{s_2, \dots, s_6\}$, and have $\mathrm{diam}(S') \leqslant \pi - \hat{\delta}$, where $\hat{\delta} = \delta_6^* - 0.016$ by Observation 2. Since no two shortcuts in S' can be combined, the strip $\mathfrak{S}(\hat{\delta})$ of area $4\hat{\delta}\pi$ is covered by the regions $\mathfrak{R}(s_i, \hat{\delta})$, for $2 \leqslant i \leqslant 6$. By Lemma 7, the area of $\mathfrak{R}(s_i, \hat{\delta})$ is at most $A(\hat{a}, \hat{\delta}) = 4\hat{\delta}(\pi - \hat{a} - \hat{\delta})$, where $\hat{a}$ is such that $\delta(\hat{a}) = \hat{\delta}$. However $5(\pi - \hat{a} - \hat{\delta}) \approx 2.9353 < \pi$, another contradiction. □

4.2. ... and a long proof for seven

We let $S = \{s_1, s_2, \dots, s_7\}$ with $|s_1| \leqslant |s_2| \leqslant \cdots \leqslant |s_7|$, set $\delta^* = \delta_6^* + \varepsilon = \delta(2) + \varepsilon$ for some small $\varepsilon > 0$, and assume that $\mathrm{diam}(S) \leqslant \pi - \delta^*$. We will show that this leads to a contradiction.

We will use the following short result.

Lemma 15. *A region $\mathfrak{R}(s, \hat{\delta})$ reaches the middle line of $\mathfrak{S}(\hat{\delta})$ for $\hat{\delta} = \delta(2) - 2\delta(\sigma_6)$ if and only if $|s| \geqslant \lambda_6$.*

Proof. Recall from Lemma 11 that σ_6 and λ_6 satisfy the relation $\delta(\sigma_6) + \delta(\lambda_6) = \delta(2)/2$. Region $\mathfrak{R}(s, \hat{\delta})$ reaches the middle line if and only if $2\delta(s) \geqslant \hat{\delta} = \delta(2) - 2\delta(\sigma_6)$ by Corollary 6 or, equivalently, $\delta(s) \geqslant \delta(2)/2 - \delta(\sigma_6)$. This is equivalent to $|s| \geqslant \lambda_6$. □

We start as in the proof of Lemma 10 for $k \in \{4, 5, 6\}$ and argue that at least three of the shortcuts in S have length at least λ_6, that is, $|s_5| \geqslant \lambda_6$. Let ℓ denote the number of short shortcuts (of length at most σ_6). We have $0 \leqslant \ell \leqslant 4$, and we let $S' = \{s_{\ell+1}, \dots, s_7\}$ be the set of $7-\ell$ shortcuts that are not short. Observe that $\mathrm{diam}(S') \leqslant \pi - \hat{\delta}(\ell)$, where $\hat{\delta}(\ell) = \delta^* - \ell \cdot 2\delta(\sigma_6)$. Here, we are again using the same argument as in the proof of Lemma 10 for $k \in \{4, 5, 6\}$.

Lemma 16. *We have* $|S'| \geqslant 6$, $\mathrm{diam}(S') \leqslant \pi - \hat{\delta}(1) = \pi - (\delta^* - 2\delta(\sigma_6))$, *and* $|s_3| \geqslant \lambda_6$.

Proof. Consider the strip $\mathfrak{S}(\hat{\delta}(4))$ and its middle line $\mathfrak{M}$. A point on the middle line corresponds to an antipodal pair, and by Lemma 11 the shortest path for an antipodal pair can use at most one shortcut of length larger than σ_6. It follows that the regions $\mathfrak{R}(s, \hat{\delta}(4))$ for $s \in S'$ must cover $\mathfrak{M}$. The region $\mathfrak{R}(s, \hat{\delta}(4))$ only reaches $\mathfrak{M}$ if $2\delta(s) \geqslant \hat{\delta}(4)$, which implies $|s| > 1.849$. By Corollary 6, the width of the two rectangles of such a region is at most $\pi - 1.849 - \hat{\delta}(4)$. Since $4 \times (\pi - 1.849 - \hat{\delta}(4)) < \pi$, there must be at least five shortcuts of length at least $1.849 > \sigma_6$, and so $|S'| \geqslant 5$ and therefore $\ell \leqslant 2$. This implies that $\mathrm{diam}(S') \leqslant \pi - \hat{\delta}(2) < \pi - \delta_5^*$. Theorem 13 now implies $|S'| \geqslant 6$. This in turn means $\ell \leqslant 1$ and therefore $\mathrm{diam}(S') \leqslant \pi - \hat{\delta}(1)$. We now redo the above argument: The region $\mathfrak{R}(s, \hat{\delta}(1))$ only reaches $\mathfrak{M}$ if $s \geqslant \lambda_6$ by Lemma 15. The width of the two rectangles of such a region is at most $\pi - \lambda_6 - \hat{\delta}(1)$. Since $4 \times (\pi - \lambda_6 - \hat{\delta}(1)) < \pi$, there must be at least five shortcuts of length at least λ_6, that is, $|s_3| \geqslant \lambda_6$. □

We will need the following lemma about the *six* shortcuts $s_2, s_3, \dots, s_7$:

Lemma 17. *If* $|s_2| > 1.7$ *and* $\mathrm{diam}(\{s_2, s_3, \dots, s_7\}) \leqslant \pi - 0.54$, *then* $|s_2| > 1.999$.

Proof. Since $|s_2| + |s_3| > 1.7 + \lambda_6 > \pi$, no combinations of the shortcuts $s_2, s_3, \dots, s_7$ are possible, and so the six regions $\mathfrak{R}(s_2, 0.54), \dots, \mathfrak{R}(s_7, 0.54)$ must cover the upper boundary $\mathfrak{B}$ of $\mathfrak{S}(0.54)$. Let $\hat{a}$ be such that $\delta(\hat{a}) = 0.54$. Since $\delta(1.999) < 0.54$, we have $\hat{a} > 1.999$. Lemma 5 and Corollary 6 imply the following: If $|s_i| \geqslant \hat{a}$, then $\mathfrak{R}(s_i, 0.54)$ covers *two* segments of $\mathfrak{B}$ of length at most $\pi - 1.999 - 0.54 < 0.603$; if $\lambda_6 \leqslant |s_i| < \hat{a}$, then $\mathfrak{R}(s_i, 0.54)$ covers *one* segment of $\mathfrak{B}$ of length at most $\pi - \lambda_6 - 0.54 < 0.727$; and if $1.7 \leqslant |s_i| < \lambda_6$, then $\mathfrak{R}(s_i, 0.54)$ covers *one* segment of $\mathfrak{B}$ of length at most $\pi - 1.7 - 0.54 < 0.902$.

Assume that $|s_4| < \hat{a}$. Then the total length of the coverage of the upper boundary $\mathfrak{B}$ of $\mathfrak{S}(0.54)$ by the six regions is at most $0.902 + 2 \times 0.727 + 3 \times 2 \times 0.603 < 2\pi$, a contradiction. So we have $|s_4| \geqslant \hat{a}$.

Therefore the five regions $\mathfrak{R}(s_3, 0.54), \dots, \mathfrak{R}(s_7, 0.54)$ consist of ten rectangles of total width at most $2 \times 0.727 + 8 \times 0.603 < 2\pi$. This implies that there must be a θ such that the segment $\{(\theta, \xi) \mid -0.54 \leqslant \xi \leqslant 0.54\}$ is disjoint from these five regions. The segment must therefore be contained in $\mathfrak{R}(s_2, 0.54)$. This is only possible if $\mathfrak{R}(s_2, 0.54)$ covers the entire height of $\mathfrak{S}(0.54)$, or equivalently, if $\delta(s_2) \geqslant 0.54$ by Lemma 5. This implies that $|s_2| \geqslant \hat{a} > 1.999$. □

We now distinguish two cases, based on the length of s_1.

4.2.1. A short shortcut exists

We first assume that $|s_1| \leqslant \sigma_6$, so that $|S'| = 6$. By Lemma 16 we have $\mathrm{diam}(S') \leqslant \pi - \hat{\delta}$, where $\hat{\delta} = \hat{\delta}(1) = \delta^* - 2\delta(\sigma_6)$.

Lemma 18. $|s_2| \geqslant \lambda_6$.

Proof. Assume for a contradiction that $|s_2| < \lambda_6$. By Lemma 15, $\mathfrak{R}(s_2, \hat{\delta})$ does not reach the middle line $\mathfrak{M}$ of strip $\mathfrak{S}(\hat{\delta})$, and so the regions of the remaining five shortcuts $s_3, \dots, s_7$ in S' must cover $\mathfrak{M}$. By Corollary 6 we have

$$\sum_{i=3}^{7} 2(\pi - |s_i| - \hat{\delta}) \geqslant 2\pi \quad \text{or, equivalently,} \quad \sum_{i=3}^{7} |s_i| \leqslant 4\pi - 5\hat{\delta},$$

which implies that the shortest of these five segments has length $|s_3| \leqslant (4\pi - 5\hat{\delta})/5 = \frac{4}{5}\pi - \hat{\delta}$.

On the other hand, the six regions $\mathfrak{R}(s_2, \hat{\delta}), \dots, \mathfrak{R}(s_7, \hat{\delta})$ must cover $\mathfrak{S}(\hat{\delta})$ entirely. The strip $\mathfrak{S}(\hat{\delta})$ has area $4\hat{\delta}\pi \approx 6.9862$. By Lemma 7 and using $\delta(1.999) < \hat{\delta}$, the total area of the six regions is

$$\sum_{i=2}^{7} A(|s_i|, \hat{\delta}) = A(|s_2|, \hat{\delta}) + A(|s_3|, \hat{\delta}) + \sum_{i=4}^{7} A(|s_i|, \hat{\delta})$$
$$\leqslant A(\lambda_6, \hat{\delta}) + A(\frac{4}{5}\pi - \hat{\delta}, \hat{\delta}) + 4A(1.999, \hat{\delta}) \approx 6.9765 < 6.9862 \approx 4\hat{\delta}\pi,$$

a contradiction. □

Since $\lambda_6 > 1.7$, we can now apply Lemma 17, and obtain $|s_2| > 1.999$, implying that *all* six shortcuts in S' have length larger than 1.999.

Finally we return to the full set S with diameter $\mathrm{diam}(S) \leqslant \pi - \delta^*$. We rotate the configuration such that the midpoint of the inner umbra of s_1 is at coordinate zero.

Lemma 19. *The pairs corresponding to configurations (θ, δ^*) on the upper boundary $\mathfrak{B}$ of $\mathfrak{S}(\delta^*)$ with θ in the following set*

$$\left\{\pi - {}^{\delta^*}\!/2\right\} \cup \left[\pi + 0.4, 2\pi - 0.4\right] \cup \left\{{}^{\delta^*}\!/2\right\}$$

cannot use shortcut s_1.

Proof. Recall that the point pair for the configuration (θ, δ^*) consists of $p = \theta - \delta^*/2$ and $q = \theta + \pi + \delta^*/2$. So for $\theta = \pi - \delta^*/2$ we have $p = \pi - \delta^*$ and $q = 0$, while for $\theta = \delta^*/2$ we have $p = 0$ and $q = \pi + \delta^*$. Since 0 lies in the deep umbra of s_1, Observation 3 implies that s_1 cannot be used.

Consider now that $\theta \in [\pi + 0.4, 2\pi - 0.4]$. We have

$$\pi + 0.11 < \pi + 0.4 - \delta^*/2 \leqslant p \leqslant 2\pi - 0.4 - \delta^*/2 < 2\pi - 0.68$$
$$0.68 < 0.4 + \delta^*/2 \leqslant q \leqslant \pi - 0.4 + \delta^*/2 < \pi - 0.11$$

Assume for a contradiction that there is a shortest path γ from p to q that uses s_1. Since s_1 lies on the long arc from p to q of length $\pi + \delta^*$, the path γ must also use another shortcut s_i with $|s_i| > 1.999$. On the other hand, γ must use s_1 and then go along the circle boundary to p or q. Since neither p nor q lie in the interval $(-0.68, 0.68)$, this subpath of γ has length at least

$$|s_1| + 0.68 - |s_1|/2 - \delta(s_1) = 0.68 + |s_1|/2 - \delta(s_1),$$

which is strictly larger than 0.68. So the entire path γ has length at least $1.999 + 0.68 > \pi - \delta(2) > \pi - \delta^*$, a contradiction. □

It follows that the six regions $\mathfrak{R}(s_2, \delta^*), \mathfrak{R}(s_3, \delta^*), \ldots, \mathfrak{R}(s_7, \delta^*)$ must cover the points on the upper boundary $\mathfrak{B}$ with θ in the set above. Since the shortcuts have length at least 1.999 and $\delta^* > \delta(2)$, each region covers a single interval on $\mathfrak{B}$ of length at most $\pi - 1.999 - \delta^* < 0.572$.

The interval $[\pi + 0.4, 2\pi - 0.4]$ has length $\pi - 2 \times 0.4 > 4 \times 0.572$, and therefore requires five regions to be covered. The distance between the isolated point $\pi - \delta^*/2$ and the interval is $0.4 + \delta^*/2 > 0.68$, and the same holds for the distance between $\delta^*/2$ and the interval. Both isolated points thus require a region that covers them and that cannot contribute to the coverage of the interval. It follows that we need seven regions to cover this subset of $\mathfrak{B}$, a contradiction.

4.2.2. No short shortcut

We now assume that $|s_1| > \sigma_6$, that is $|S'| = 7$. Since $|s_3| \geqslant \lambda_6$ and $|s_1| + |s_3| > \sigma_6 + \lambda_6 = \pi - \delta_6^* = \pi - \delta(2) > \pi - \delta^*$, the only possible combination of shortcuts that can be used is the combination of s_1 and s_2.

Lemma 20. $\delta(s_1) + \delta(s_2) < 0.2$

Proof. We first assume that the combination of shortcuts s_1 and s_2 is never used. Then the seven regions $\mathfrak{R}(s_i, \delta^*)$ for $i = 1, 2, \ldots, 7$ cover the upper boundary $\mathfrak{B}$ of $\mathfrak{S}(\delta^*)$. By Corollary 6 we have

$$\sum_{i=1}^{7}(\pi - |s_i| - \delta^*) \geqslant 2\pi \quad \text{or, equivalently,} \quad |s_1| + |s_2| \leqslant 5\pi - 7\delta^* - \sum_{i=3}^{7}|s_i|$$

and using $|s_3| \geqslant \lambda_6$ this gives

$$|s_1| + |s_2| \leqslant 5\pi - 7\delta^* - 5\lambda_6 < 2.34.$$

Since $|s_1| \geqslant \sigma_6$, convexity of δ gives us $\delta(s_1) + \delta(s_2) \leqslant \delta(\sigma_6) + \delta(2.34 - \sigma_6) \approx 0.15 < 0.2$, proving the claim.

It remains to consider the case where for some pair of points the combination of s_1 and s_2 needs to be used (and one of the two shortcuts alone does not suffice). This implies that $|s_1| + |s_2| \leqslant \pi - \delta^* < \pi - \delta(2) = \pi/2 + 1$.

If $|s_1| \geqslant 0.83$, then convexity of δ implies $\delta(s_1) + \delta(s_2) \leqslant \delta(0.83) + \delta(\pi/2 + 1 - 0.83) \approx 0.1986 < 0.2$.

If $|s_1| < 0.83$ and $|s_2| \leqslant 1.7$, then $\delta(s_1) + \delta(s_2) \leqslant \delta(0.83) + \delta(1.7) \approx 0.1789 < 0.2$.

Finally, if $|s_1| < 0.83$ and $|s_2| > 1.7$, then we observe that $\operatorname{diam}(\{s_2, \ldots, s_7\}) \leqslant \pi - \delta^* + 2\delta(s_1) < \pi - \delta^* + 2\delta(0.83) < \pi - 0.54$. We can thus apply Lemma 17 and find that $|s_2| > 1.999$. Since $\sigma_6 + 1.999 > \pi - \delta(2) > \pi - \delta^*$, no combination of s_1 and s_2 is possible and we are back in the first case. □

Let $\zeta = \delta(2) - 0.4 = \pi/2 - 1.4 \approx 0.1708$, and consider the configurations (θ, ζ) in $\mathfrak{S}(\delta^*)$. Let (p, q) be a pair of points on C corresponding to some (θ, ζ) in $\mathfrak{S}(\delta^*)$. The shorter arc $\widehat{pq}$ along C has length $\pi - \zeta = \pi - (\pi/2 - 1.4) = \pi/2 + 1.4$. Thus, any path from p to q using s_1, s_2, or their combination has length at least $\pi - \zeta - 2\delta(s_1) - 2\delta(s_2) > \pi/2 + 1$ since $\delta(s_1) + \delta(s_2) < 0.2$ according to Lemma 20. Since $\pi - \delta^* < \pi - \delta(2) = \pi/2 + 1$, it follows that the pairs corresponding

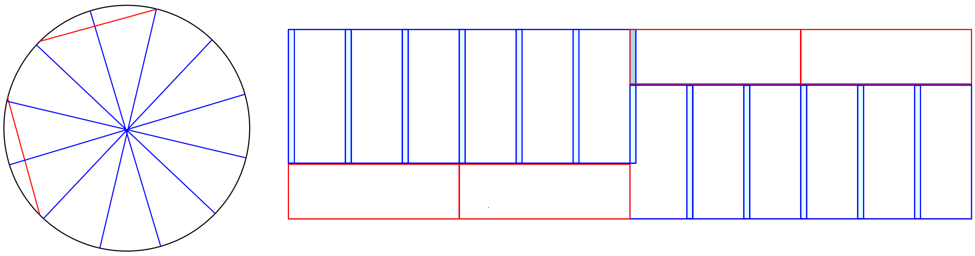

Fig. 9. A shortcut configuration S of 8 shortcuts with $\mathrm{diam}(S) < \mathrm{diam}(6)$, and the corresponding regions in the strip $\mathfrak{S}(\delta^*)$.

to (θ, ζ) cannot make any use of s_1 and s_2. That is, they can only be covered by the five regions $\mathfrak{R}(s_3, \delta^*), \dots, \mathfrak{R}(s_7, \delta^*)$. Since $|s_3| \geqslant \lambda_6$, these regions have width at most $\pi - \lambda_6 - \delta^* < 0.696$. For a region $\mathfrak{R}(s, \delta^*)$ to cover *two* pieces of $\xi = \zeta$ in $\mathfrak{S}(\delta^*)$, we need $2\delta(s) \geqslant \delta^* + \zeta$ by Lemma 5, which implies that $|s| > 1.949$. This in turn means that the width of the region is at most $\pi - 1.949 - \delta^* < 0.622$. It follows that the five regions $\mathfrak{R}(s_3, \delta^*), \dots, \mathfrak{R}(s_7, \delta^*)$ can cover at most $5 \times 2 \times 0.622 < 2\pi$ of the line $\xi = \zeta$ in $\mathfrak{S}(\delta^*)$, a contradiction.

This concludes the proof of Theorem 14 for seven shortcuts.

5. Eight shortcuts

With eight shortcuts we can improve on the diameter, obtaining $\mathrm{diam}(8) < \mathrm{diam}(7) = \mathrm{diam}(6)$. Our construction S consists of six long shortcuts with length $a_1 \approx 1.999870869$ and two short ones with length $a_2 \approx 0.988571799$, placed as in Fig. 9(left), and achieves the diameter $\mathrm{diam}(S) \approx \pi - 0.5822245291 = 2.559368125 < \mathrm{diam}(6)$.

We obtained this construction by maximizing δ^* with constraints $\pi - a_1 - \delta^* \geqslant \pi/6$, $\pi - a_2 - \delta^* \geqslant \pi/2$, and $\delta(a_1) + \delta(a_2) \geqslant \delta^*$. We can thus cover $\mathfrak{S}(\delta^*)$ as seen in the diagram in Fig. 9(right). In particular, we have $\pi - a_2 - \delta^* = \pi/2$ and $\delta(a_1) + \delta(a_2) = \delta^*$, while we have a strict inequality $\pi - a_1 - \delta^* > \pi/6$ in our construction. So, in the strip $\mathfrak{S}(\delta^*)$, the regions slightly overlap.

6. An asymptotically tight bound

In this final section, we show that $\mathrm{diam}(k) = 2 + \Theta(1/k^{2/3})$ as k goes to infinity.

Theorem 21. *To achieve diameter at most* $2 + 1/m$, $\Theta(m^{3/2})$ *shortcuts are both necessary and sufficient.*

Proof. We prove the necessary condition first. Consider two points p, q that form an angle of $\pi - t/m$, for some integer $0 \leqslant t \leqslant \sqrt{m} - 2$. Consider two intervals I_p and I_q, both of arc length $4/m$, with the midpoint of the intervals at p and q, respectively. We claim that if there is no shortcut connecting a point of I_p with a point of I_q, then the distance between p and q is larger than $2 + 1/m$.

If there is no such shortcut, then the shortest path from p to q must visit a point r on the circle not in either interval, see Fig. 10(left). The sum $|pr| + |rq|$ is minimized when r is the point making angle $2/m$ with q, so we have $\alpha(pr) = \pi - (t+2)/m$ and $\alpha(rq) = 2/m$.

This gives us

$$|qr| = 2\sin\frac{2}{2m} = 2\sin\frac{1}{m} \geqslant \frac{2}{m} - \frac{2}{3!}\frac{1}{m^3} > \frac{2}{m} - \frac{1}{3m} = \frac{5}{3m},$$

$$|pr| = 2\sin(\frac{\pi}{2} - \frac{t+2}{2m}) = 2\cos\frac{t+2}{2m} \geqslant 2\cos\frac{\sqrt{m}}{2m} = 2\cos\frac{1}{2\sqrt{m}} \geqslant 2 - \frac{1}{4m},$$

and so $|pr| + |rq| > 2 + 1/m$.

We now subdivide C into $\Theta(m)$ intervals of length at least $6/m$. Consider a pair of intervals I, J at arc distance at least $\pi - 1/\sqrt{m}$. Then there are points $p \in I$ and $q \in J$ with $I_p \subset I$ and $I_q \subset J$ and p, q forming an angle of the form $\pi - t/m$ for an integer $0 \leqslant t \leqslant \sqrt{m} - 2$. It follows that there must be some shortcut connecting I and J. Since there are $\Theta(m^{3/2})$ such pairs of intervals, we must have at least $\Omega(m^{3/2})$ shortcuts.

We now turn to the sufficient condition, and construct a set of $\Theta(m^{3/2})$ shortcuts that give a diameter of $2 + 1/m$.

We start by placing $4\pi m$ points uniformly around the circle, and connect each pair that makes an angle larger than $\pi - 4\sqrt{1/m}$, as shown in Fig. 10(center). This creates $\Theta(m^{3/2})$ shortcuts and ensures that for points p, q with angle larger than $\pi - 4\sqrt{1/m}$ the distance between p and q is bounded by $2 + 1/m$.

It remains to add shortcuts to decrease the distance of point pairs p, q that form an arc between 2 and $\pi - 4\sqrt{1/m}$. For each integer t with $4\sqrt{m} < t < 2m$ we will create a set of shortcuts of arc length $\pi - t/m$, see Fig. 10(right). These shortcuts will be used for pairs p, q forming an arc between $\pi - t/m$ and $\pi - (t-1)/m$.

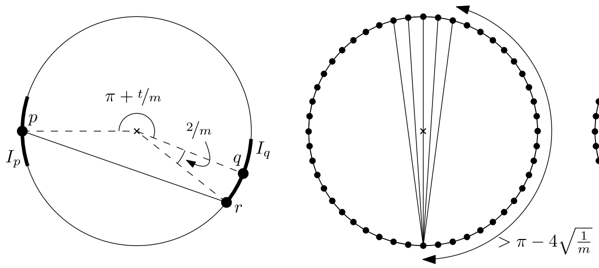

Fig. 10. (left) If there is no shortcut between I_p and I_q then the shortest path from p to q must visit a point r on the circle not in either interval. (center) Shortcut between every pair that makes an angle larger than $\pi - 4\sqrt{1/m}$. (right) Adding shortcuts of arc length $\pi - t/m$.

Let us fix such a value t, and consider a shortcut s of arc length $\pi - t/m$. Then the length of the shortcut is

$$|s| = 2\sin\frac{\pi - t/m}{2} = 2\cos\frac{t}{2m}.$$

Using the bound $\cos x \leqslant 1 - \frac{x^2}{2} + \frac{x^4}{24} \leqslant 1 - (\frac{1}{2} - \frac{1}{24})x^2 = 1 - \frac{11}{24}x^2$ for $x < 1$, we have

$$|s| \leqslant 2 - 2\frac{11}{24}\frac{t^2}{4m^2} = 2 - \frac{11}{48}\frac{t^2}{m^2} < 2 - \frac{1}{6}\frac{t^2}{m^2} = 2 - 2\Delta,$$

where we define $\Delta = \frac{1}{12}(\frac{t}{m})^2$. Since $t > 4\sqrt{m}$ we have $\Delta > \frac{16}{12}\frac{1}{m} > \frac{1}{m}$.

We repeat shortcuts of this length every arc interval of length Δ. Consider now a pair of points p, q forming an angle in the interval $\pi - t/m$ to $\pi - (t-1)/m$. We can go from p to q by first going to the nearest shortcut along an arc of length at most Δ, then following the shortcut of length at most $2 - 2\Delta$, and finally going backwards by at most Δ, or forward by at most $1/m < \Delta$. It follows that the distance between p and q is at most $2 - 2\Delta + 2\Delta = 2$.

The number of shortcuts of length $\pi - t/m$ is $2\pi/\Delta$, and so the total number of shortcuts of this type is

$$\sum_{t=4\sqrt{m}+1}^{2m} 24\pi\frac{m^2}{t^2} = 24\pi m^2 \sum_{t=4\sqrt{m}+1}^{2m} \frac{1}{t^2} \leqslant 24\pi m^2 \int_{4\sqrt{m}}^{\infty} \frac{1}{x^2}\,dx = 6\pi m^{3/2}.$$

This completes the proof. □

7. Conclusions

We have given exact bounds on the diameter for up to seven shortcuts. In all cases, the shortcuts are of equal length. For $k = 8$, however, our upper bound construction uses shortcuts of two different lengths. On the other hand, it is not difficult to see that eight shortcuts of equal length cannot even achieve a slightly better diameter than diam(6). In general, what is the diameter achievable with k shortcuts of equal length?

We have shown that for $k = 0$ and $k = 6$ we have $\text{diam}(k) = \text{diam}(k+1)$. Are there any other values of k for which this holds?

Finally, in all our constructions, including the one for large k, we never use combinations of shortcuts: the shortest path for any pair of points uses at most one shortcut. Is it true that combinations of shortcuts never help, for any k? Meanwhile, one could make it a requirement and ask: What is the best diameter achievable with k shortcuts, under the restriction that no path can use more than one shortcut? This problem then reduces to the problem of covering the strip $\mathfrak{S}(\delta^*)$ by regions $\mathfrak{R}(s, \delta^*)$, and may be more tractable than the general form.

Appendix A. Calculations

```
Source code at: http://github.com/otfried/circle-shortcuts

Lemma 4:
=========

delta(delta(2)) = 0.00402

Section 3.1:
===========
```

```
Table 1:
k  a*     d*     pi - d*
2  1.4782 0.0926 3.0490
3  1.8435 0.2509 2.8907
4  1.9619 0.3943 2.7473
5  1.9969 0.5164 2.6252
6  2.0000 0.5708 2.5708

Computing mu_k:
mu_2 = 1.2219
mu_3 = 1.5943
mu_4 = 1.7623
mu_5 = 1.8526
mu_6 = 1.8828

Lemma 10:
=========

Lemma 10 for k = 3

a* = 1.8435, d* = 0.2509, mu = 1.5943
(i) (pi - d*)/2 = 1.4454
(v) delta(1.45) = 0.0860
(vi) delta(pi/2) = 0.1179
(vii) a such that delta(a) = 0.06 = 1.3150

Table 2:
k & a*     & d*     & mu_k   & sigma_k & lambda_k
4 & 1.9619 & 0.3943 & 1.7623 & 1.0373  & 1.7100 \\
5 & 1.9969 & 0.5164 & 1.8526 & 0.7862  & 1.8390 \\
6 & 2.0000 & 0.5708 & 1.8828 & 0.6957  & 1.8751 \\

Lemma 10 for k in {4, 5, 6}:

Showing that l = 1:
k = 4:
  d* = 0.3943, sigma = 1.0373, lambda = 1.7100
  delta^ = 0.3411, w = (pi - lambda - d^) = 1.0906
  (k-2) w = 2.1811 < pi
k = 5:
  d* = 0.5164, sigma = 0.7862, lambda = 1.8390
  delta^ = 0.4728, w = (pi - lambda - d^) = 0.8298
  (k-2) w = 2.4894 < pi
k = 6:
  d* = 0.5708, sigma = 0.6957, lambda = 1.8751
  delta^ = 0.5262, w = (pi - lambda - d^) = 0.7403
  (k-2) w = 2.9610 < pi

The final contradiction of Lemma 10:
k = 4:
  delta^ = 0.3411, a^ = 1.9304, w = pi - a^ - delta^ = 0.8701
  (k-1) w = 2.6103 < pi
k = 5:
  delta^ = 0.4728, a^ = 1.9893, w = pi - a^ - delta^ = 0.6795
  (k-1) w = 2.7178 < pi
k = 6:
  delta^ = 0.5262, a^ = 1.9979, w = pi - a^ - delta^ = 0.6174
  (k-1) w = 3.0872 < pi

Proof of Lemma 12:
==================
```

```
k=3: mu=1.5943, w = pi - mu - d* = 1.2964 => (k-1)w = 2.5928 < pi
k=4: mu=1.7623, w = pi - mu - d* = 0.9850 => (k-1)w = 2.9549 < pi
k=5: mu=1.8526, w = pi - mu - d* = 0.7726 => (k-1)w = 3.0902 < pi

Proof of Theorem 14 for k = 6:
==============================

mu=1.8828, d*=0.5708, w = pi - mu - d* = 0.6880
 4 w = 2.7518 < pi
 pi - d* + 5 w = 6.0106 < 2pi
 delta(pi - d* - mu) = 0.0072
 d^ = 0.5548, a^ = 1.9997
 5 (pi - a^ - d^) = 2.9353 < pi

Seven shortcuts, Proof of Lemma 16:
===================================

d* = 0.5708, sigma6 = 0.6957, delta(sigma6) = 0.0074
d^(4) = 0.5114, 2 delta(1.849) = 0.5104
  w = pi - 1.849 - d^(4) = 0.7812
  4 * w = 3.1248 < pi
d^(2) = 0.5411 > 0.5164 = d5*
d^(1) = 0.5559
  w = pi - lambda6 - d^(1) = 0.7105
  4 * w = 2.8422 < pi

Seven shortcuts, Proof of Lemma 17:
===================================

1.7 + lambda6 = 3.5751 > pi
delta(1.999) = 0.5397 < 0.54
  w1 = pi - 1.999 - 0.54 = 0.6026
  w2 = pi - lambda6 - 0.54 = 0.7265
  w3 = pi - 1.7 - 0.54 = 0.9016
  w3 + 2 * w2 + 6 * w1 = 5.9701 < 2pi
  2 * w2 + 8 * w1 = 6.2737 < 2pi

Seven shortcuts, Proof of Lemma 18:
===================================

d^ = d* - 2 * delta(sigma6) = 0.5559
s3 <= 0.8 * pi - d^ = 1.9573
A(lambda6, d^) = 4 * delta(lambda6) * (pi - lambda6 - d^) = 0.7900
A(s3, d^) <= 4 * delta(1.9573) * (pi - 1.9573 - d^) = 0.9681
delta(1.999) = 0.53967 < d^
A(1.999, d^) < 4 * d^ * (pi - 1.999 - d^) = 1.3046
0.7900 + 0.9681 + 4 * 1.3046 = 6.9765 < 6.9862 = 4 * d^ * pi

Seven shortcuts, Proof of Lemma 19:
===================================

0.4 - d*/2 = 0.1146, 0.4 + d*/2 = 0.6854

Seven shortcuts, a short shortcut exists
========================================

pi - 1.999 - d* = 0.5718
pi - 2 * 0.4 = 2.3416 > 2.2872 = 4 * 0.5718
0.4 + d*/2 = 0.6854 > 0.5718
```

```
Seven shortcuts, Proof of Lemma 20:
===================================

s1 + s2 <= 5 * pi - 7 * ds - 5 * lambda6 = 2.3369 < 2.34
delta(sigma6) + delta(2.34 - sigma6) = 0.1505 < 0.2
delta(0.83) + delta(pi/2 + 1 - 0.83) = 0.1986 < 0.2
delta(0.83) + delta(1.7) = 0.1789 < 0.2
1.999 + sigma6 = 2.6947 > 2.5708 = pi - d*

Seven shortcuts, no short shortcut
==================================

zeta = pi/2 - 1.4 = 0.1708
pi - lambda6 - d* = 0.6957
2 * delta(1.949) = 0.7400 < 0.7416 = d* + zeta
pi - 1.949 - d* = 0.6218
 5 * 2 * 0.622 = 6.2200 < 2pi
```

References

[1] P. Erdős, A. Rényi, On a problem in the theory of graphs, Publ. Math. Inst. Hung. Acad. Sci. 7 (1962) 623–641 (in Hungarian).
[2] P. Erdős, A. Rényi, V.T. Sós, On a problem of graph theory, Studia Sci. Math. Hung. 1 (1966) 215–235.
[3] F.R.K. Chung, M.R. Garey, Diameter bounds for altered graphs, J. Graph Theory 8 (1984) 511–534.
[4] A.A. Schoone, H.L. Bodlaender, J. van Leeuwen, Diameter increase caused by edge deletion, J. Graph Theory 11 (1987) 409–427.
[5] F.R.K. Chung, Diameters of graphs: old problems and new results, Congr. Numer. 60 (1987) 295–317.
[6] D. Bilò, L. Gualà, G. Proietti, Improved approximability and non-approximability results for graph diameter decreasing problems, Theor. Comput. Sci. 417 (2012) 12–22.
[7] F. Frati, S. Gaspers, J. Gudmundsson, L. Mathieson, Augmenting graphs to minimize the diameter, Algorithmica 72 (2015) 995–1010.
[8] S. Kapoor, M. Sarwat, Bounded-diameter minimum-cost graph problems, Theory Comput. Syst. 41 (2007) 779–794.
[9] C.L. Li, S.T. McCormick, D. Simchi-Levi, On the minimum-cardinality-bounded-diameter and the bounded-cardinality-minimum-diameter edge addition problems, Oper. Res. Lett. 11 (1992) 303–308.
[10] U. Große, J. Gudmundsson, C. Knauer, M. Smid, F. Stehn, Fast algorithms for diameter-optimally augmenting paths, in: 42nd International Colloquium Automata, Languages, and Programming (ICALP), in: Lect. Notes Comput. Sci., vol. 9134, Springer, 2015, pp. 678–688.
[11] H. Wang, An improved algorithm for diameter-optimally augmenting paths in a metric space, arXiv:1608.04456, 2016.
[12] B. Yang, Euclidean chains and their shortcuts, Theor. Comput. Sci. 497 (2013) 55–67, https://doi.org/10.1016/j.tcs.2012.03.021.
[13] J.-L.D. Carufel, C. Grimm, A. Maheshwari, M. Smid, Minimizing the continuous diameter when augmenting paths and cycles with shortcuts, in: 15th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2016), in: Leibniz International Proceedings in Informatics (LIPIcs), vol. 53, 2016, pp. 27:1–27:14.
[14] J. Cáceres, D. Garijo, A. González, A. Márquez, M.L. Puertas, P. Ribeiro, Shortcut sets for plane Euclidean networks (Extended abstract), Electron. Notes Discrete Math. 54 (2016) 163–168, https://doi.org/10.1016/j.endm.2016.09.029.