Shortcuts for the Circle*

Sang Won Bae†1, Mark de Berg‡2, Otfried Cheong§3, Joachim Gudmundsson¶4, and Christos Levcopoulos5

1 **Kyonggi University, Suwon, Republic of Korea**
swbae@kgu.ac.kr
2 **TU Eindhoven, Eindhoven, The Netherlands**
mdberg@win.tue.nl
3 **KAIST, Daejon, Republic of Korea**
otfried@kaist.edu
4 **University of Sydney, Sydney, Australia**
joachim.gudmundsson@gmail.com
5 **Lund University, Sweden**
christos@cs.lth.se

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 $2 + \Theta(1/k^{\frac{2}{3}})$ for any k.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Computational geometry, graph augmentation problem, circle, shortcut, diameter

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2017.9

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.

---

* A full version of this extended abstract is available on the ArXiv [1], https://arxiv.org/abs/1612.02412

† SWB is supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (2015R1D1A1A01057220).
‡ MdB is supported by the Netherlands' Organisation for Scientific Research (NWO) under project no. 024.002.003.
§ OC is supported by NRF grant 2011-0030044 (SRC-GAIA) funded by the government of Korea.
¶ JG is supported under Australian Research Council's Discovery Projects funding scheme (project number DP150101134)

28th International Symposium on Algorithms and Computation (ISAAC 2017).
Editors: Yoshio Okamoto and Takeshi Tokuyama; Article No. 9; pp. 9:1–9:13
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

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 [7, 8].

Chung and Garey [6] 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 [13].)

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 [2, 9, 11, 12, 13]. 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 [14].

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 [15] 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. [4] 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. [3] 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 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 $\operatorname{diam}(7) = \operatorname{diam}(6)$.

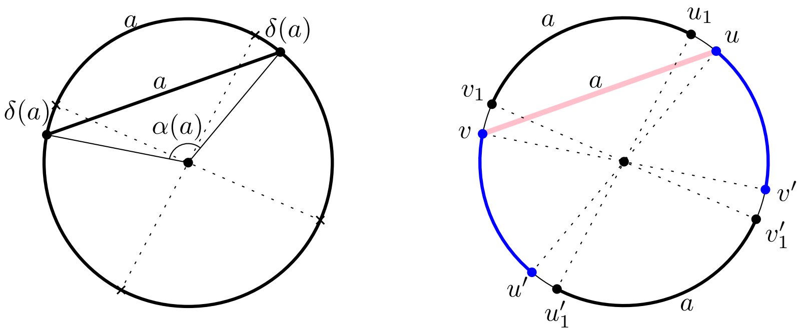

Figure 1 For a shortcut s of length $a = |s|$, (left) $\alpha(a) = a + 2\delta(a)$ and (right) the umbra $U(s)$ (consisting of two arcs of length a in thick black) and radiance (in thick blue).

- We have $\mathrm{diam}(8) < \mathrm{diam}(7)$.
- We show that $\mathrm{diam}(k) = 2 + \Theta(1/k^{\frac{2}{3}})$.

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`.

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 Figure 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 Figure 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 $q \notin U(s_m)$. If γ traverses s_i from its endpoint u_i to its other endpoint v_i, then $v_i \notin U(s_{i+1})$ and $u_{i+1} \notin U(s_i)$ for $i = 1, \ldots, m - 1$.*

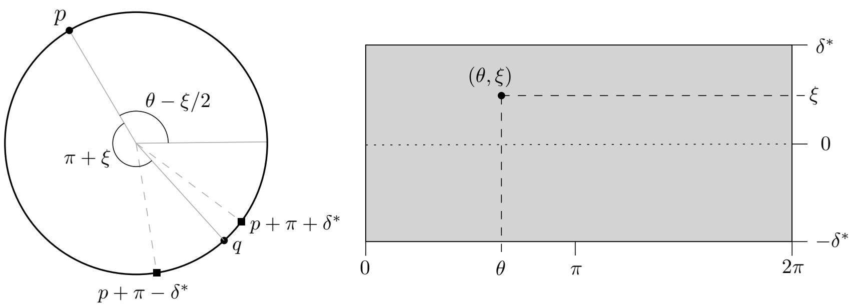

Figure 2 (left) (θ, ξ) corresponds to the pair of points $p = \theta - \xi/2$ and $q = \theta + \pi + \xi/2$. (right) $\mathfrak{S}(\delta^*)$ represents all pair of points $p = \theta - \xi/2$ and $q = \theta + \pi + \xi/2$ with $-\delta^* \leqslant \xi \leqslant \delta^*$.

(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, $\operatorname{diam}(1) = \operatorname{diam}(0) = \pi$.

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 $2\delta(s_i)$.

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 Figure 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 $d_s(\theta - \xi/2, \theta + \pi + \xi/2) \leqslant \pi - \delta^*$. In the following, 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 $-\alpha/2$ 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 topology above, it suffices to understand $\mathfrak{R}(s, \delta^*)$

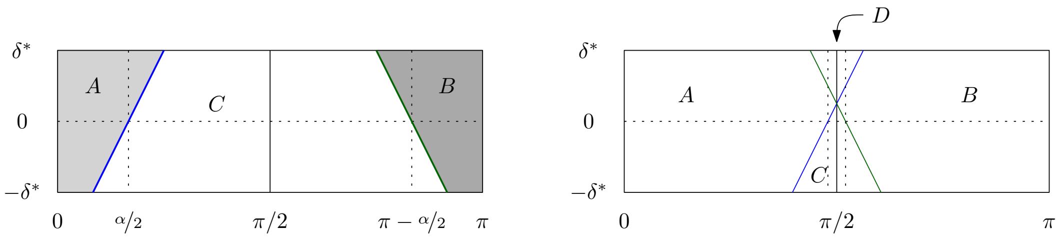

Figure 3 The regions A, B, C, D.

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 light gray region above the blue line in Figure 3(left). Second, we have $q > 2\pi - \alpha/2$ if and only if $\xi > 2\pi - \alpha - 2\theta$. This is the dark gray region above the green line in Figure 3(left). If the two regions do not intersect then we get three regions as shown in Figure 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 Figure 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 Figure 4.

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 Figure 4.

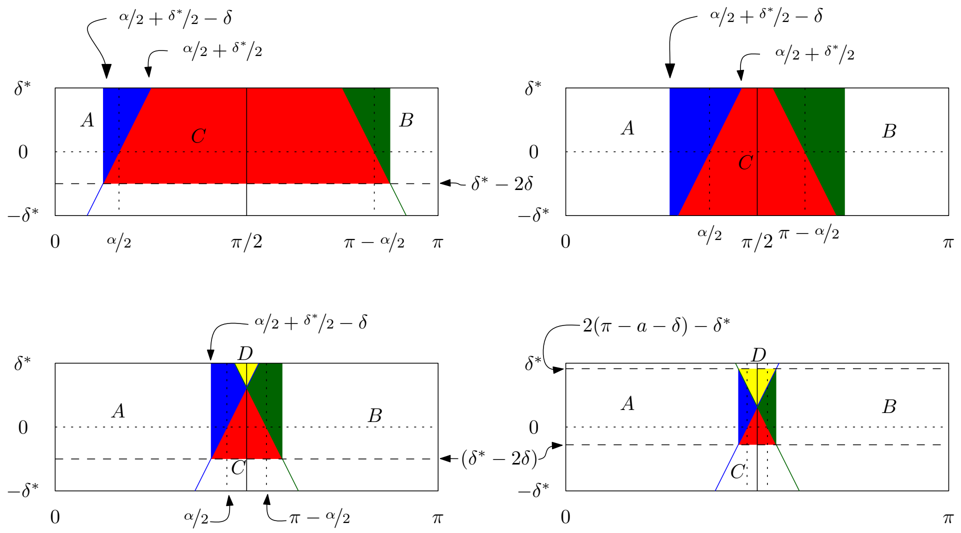

Figure 4 The region $\mathfrak{R}(s, \delta^*)$ in four different situations.

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 Figure 4.

When $\alpha > \pi - \delta^*$ regions A and B intersect in region D, as shown in Figure 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 Figure 4. There are two cases that can occur, as is shown on the bottom left and bottom right of Figure 4. The discussion of these cases can be found in the proof of the following lemma, given in the full paper [1].

▶ **Lemma 3.** *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*

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

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 3 where $\delta^* > \pi - a - \delta(a)$ only happens when $\delta^* > \delta(2) = \pi/2 - 1$.

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

k	a_k^*	δ_k^*	$\mathrm{diam}(S) = \pi - \delta_k^*$
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

We will also be interested in the length of the intersection of $\mathfrak{R}(s, \delta^*)$ with the middle line $\mathfrak{M} = \{\xi = 0\}$ of $\mathfrak{S}(\delta^*)$. Note that $\mathfrak{M}$ has length 2π. We have the following corollary to Lemma 3:

▶ **Corollary 4.** *Let $\delta^* \in [0, \pi - 2]$, and let s be a shortcut of length $a \in (0, 2]$. Then*

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

3 Up to five shortcuts

In this section we derive the exact value of $\mathrm{diam}(k)$ for $k \in \{2, 3, 4, 5\}$, and show 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.

Using only one shortcut. Again we consider a target diameter of the form $\pi - \delta^*$, with $\delta^* \in [0, \pi - 2]$. By Lemma 3, 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 5.** *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 $A(a^*, \delta^*) = 4\delta^*(\pi - a^* - \delta^*)$.*

The proof can be found in the full paper [1].

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

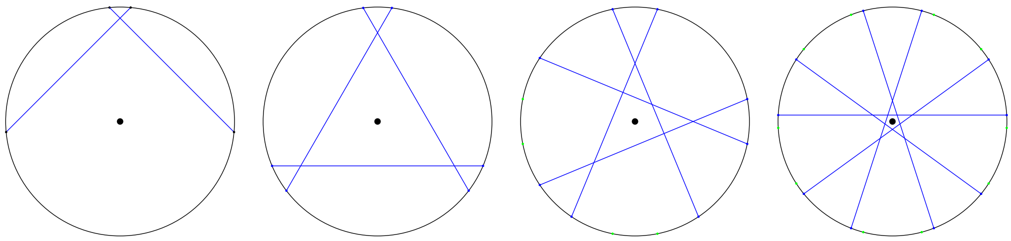

Figure 5 The optimal shortcut configurations for $k = 2, 3, 4, 5$.

Proof. By Lemma 3, 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 $\operatorname{diam}(S) = \pi - \delta_k^*$. Figure 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 5 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. ◀

Shortcuts cannot be combined. It remains to show that the configurations in Figure 5 are optimal even if combinations of shortcuts can be used. We start by defining $\mu_k \in [0, 2]$ to be such that $\delta(\mu_k) = \delta_k^*/2$. By Lemma 3, $\mathfrak{R}(s, \delta_k^*)$ 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 have $d_s(p, q) \leqslant \pi - \delta_k^*$ only if $|s| \geqslant \mu_k$.

We handle the somewhat special case $k = 2$ first.

▶ **Lemma 7.** *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 Figure 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 $|s_1| \geqslant a_2^*$.

We next 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

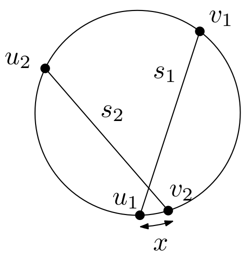

Figure 6 The two shortcuts must intersect.

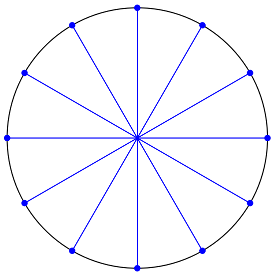

Figure 7 An optimal configuration of six shortcuts.

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 Figure 6.

Let x be the length of overlap of the arcs of s_1 and s_2, that is, $x = |\widehat{u_1 v_2}|$ in Figure 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 6 implies the claim. ◀

For $3 \leqslant k \leqslant 6$, the key insight is the following lemma, proven in the full paper [1].

▶ **Lemma 8.** *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.*

▶ **Lemma 9.** *Let S be a set of k shortcuts for $k \in \{3, 4, 5\}$ such that* $\operatorname{diam}(S) \leqslant \pi - \delta_k^*$. *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 8 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 4 $\mathfrak{R}(s_i, \delta_k^*)$ 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 6, 7, and 9, we obtain our first theorem.

▶ **Theorem 10.** *For $k \in \{2, 3, 4, 5\}$ there is a set S of k shortcuts that achieves* $\operatorname{diam}(S) = \pi - \delta_k^*$. *This is optimal and the solution is unique up to rotation.*

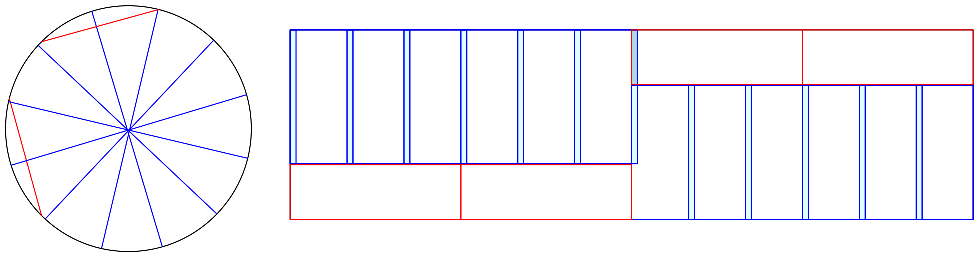

Figure 8 A shortcut configuration S of 8 shortcuts with $\text{diam}(S) < \text{diam}(6)$, and the corresponding regions in the strip $\mathfrak{S}(\delta^*)$.

4 Six and seven shortcuts

The configuration of six shortcuts of length 2 (that is, all shortcuts are diameters of the circle) shown in Figure 7 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.

It remains to argue that the configuration is indeed optimal, that is, there is no set S of six shortcuts that achieves $\text{diam}(S) < \pi - \delta(2)$. 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 $\text{diam}(S) < \pi - \delta(2)$. This implies a collapse between the cases of $k = 6$ and $k = 7$, that is, $\text{diam}(7) = \text{diam}(6) = \pi - \delta(2)$. The proof is quite long and rather technical, and can be found in the full paper [1].

▶ **Theorem 11.** *There is a set S of six shortcuts that achieves* $\text{diam}(S) = \pi - \delta(2) = \pi/2 + 1$. *There is no configuration of six or seven shortcuts that has diameter smaller than* $\pi/2 + 1$. *Therefore, we have* $\text{diam}(7) = \text{diam}(6) = \pi/2 + 1$.

5 Eight shortcuts

With eight shortcuts we can improve on the diameter, obtaining $\text{diam}(8) < \text{diam}(7) = \text{diam}(6)$. 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 Figure 8(left), and achieves the diameter $\text{diam}(S) \approx \pi - 0.5822245291 = 2.559368125 < \text{diam}(6)$.

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 Figure 8(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 $\text{diam}(k) = 2 + \Theta(1/k^{\frac{2}{3}})$ as k goes to infinity.

▶ **Theorem 12.** *To achieve diameter at most* $2 + 1/m$, $\Theta(m^{\frac{3}{2}})$ *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 $4/m$, and centered at p and q, respectively. We claim that if there is no shortcut

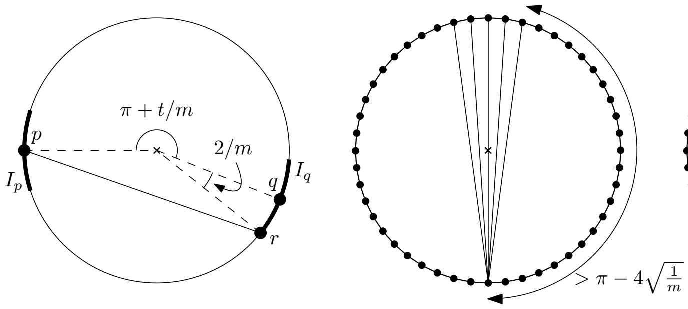

Figure 9 (left) If there is no shortcut between I_p and I_q then the shortest path from p to q must visit a point r on the circle not in either interval. (center) Shortcut between every pair that makes an angle larger than $\pi - 4\sqrt{1/m}$. (right) Adding shortcuts of arc length $\pi - t/m$.

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 Figure 9(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 Figure 9(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 Figure 9(right). These shortcuts will be used for pairs p, q forming an arc between $\pi - t/m$ and $\pi - (t-1)/m$.

Let us fix such a 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 $\mathrm{diam}(k) = \mathrm{diam}(k+1)$. Are there any other values of k for which this holds?

References

1. S. W. Bae, M. de Berg, O. Cheong, J. Gudmundsson, and C. Levcopoulos. Shortcuts for the Circle, 2016. `arXiv:1612.02412`.
2. D. Bilò, L. Gualà, and G. Proietti. Improved approximability and non-approximability results for graph diameter decreasing problems. *Theoretical Computer Science*, 417:12–22, 2012.
3. J. Cáceres, D. Garijo, A. González, A. Márquez, M. L. Puertas, and P. Ribeiro. Shortcut sets for plane Euclidean networks (extended abstract). *Electronic Notes in Discrete Mathematics*, 54:163 – 168, 2016. `doi:10.1016/j.endm.2016.09.029`.
4. J.-L. De Carufel, C. Grimm, A. Maheshwari, and M. Smid. Minimizing the Continuous Diameter when Augmenting Paths and Cycles with Shortcuts. In *15th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2016)*, volume 53 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 27:1–27:14, 2016. `doi:10.4230/LIPIcs.SWAT.2016.1`.
5. F. R. K. Chung. Diameters of graphs: old problems and new results. *Congressus Numerantium*, 60:295–317, 1987.
6. F. R. K. Chung and M. R. Garey. Diameter bounds for altered graphs. *Journal of Graph Theory*, 8:511–534, 1984.
7. P. Erdös and A. Rényi. On a problem in the theory of graphs (in Hungarian). *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 7:623–641, 1962.
8. P. Erdös, A. Rényi, and V. T. Sós. On a problem of graph theory. *Studia Scientarium Mathematicarum Hungar*, 1:215–235, 1966.
9. F. Frati, S. Gaspers, J. Gudmundsson, and L. Mathieson. Augmenting graphs to minimize the diameter. *Algorithmica*, 72(4):995–1010, 2015.

10 U. Große, J. Gudmundsson, C. Knauer, M. Smid, and F. Stehn. Fast algorithms for diameter-optimally augmenting paths. In *42nd International Colloquium Automata, Languages, and Programming (ICALP)*, pages 678–688. Springer, 2015.

11 S. Kapoor and M. Sarwat. Bounded-diameter minimum-cost graph problems. *Theory of Computing Systems*, 41(4):779—794, 2007.

12 C. L. Li, S. T. McCormick, and D. Simchi-Levi. On the minimum-cardinality-bounded-diameter and the bounded-cardinality-minimum-diameter edge addition problems. *Operation Research Letters*, 11(5):303—-308, 1992.

13 A. A. Schoone, H. L. Bodlaender, and J. van Leeuwen. Diameter increase caused by edge deletion. *Journal of Graph Theory*, 11:409–427, 1987.

14 H. Wang. An improved algorithm for diameter-optimally augmenting paths in a metric space, 2016. `arXiv:1608.04456`.

15 B. Yang. Euclidean chains and their shortcuts. *Theoretical Computer Science*, 497:55 – 67, 2013. `doi:10.1016/j.tcs.2012.03.021`.