Minimizing the Continuous Diameter when Augmenting Paths and Cycles with Shortcuts*

Jean-Lou De Carufel[1], Carsten Grimm[2], Anil Maheshwari[3], and Michiel Smid[4]

1 **School of Electrical Engineering and Computer Science, University of Ottawa, Ottawa, Canada**
2 **School of Computer Science, Carleton University, Ottawa, Canada and Institut für Simulation und Graphik, Otto-von-Guericke-Universität Magdeburg, Magdeburg, Germany**
`carsten.grimm@ovgu.de`
3 **School of Computer Science, Carleton University, Ottawa, Canada**
4 **School of Computer Science, Carleton University, Ottawa, Canada**

Abstract

We seek to augment a geometric network in the Euclidean plane with shortcuts to minimize its continuous diameter, i.e., the largest network distance between any two points on the augmented network. Unlike in the discrete setting where a shortcut connects two vertices and the diameter is measured between vertices, we take all points along the edges of the network into account when placing a shortcut and when measuring distances in the augmented network.

We study this network augmentation problem for paths and cycles. For paths, we determine an optimal shortcut in linear time. For cycles, we show that a single shortcut never decreases the continuous diameter and that two shortcuts always suffice to reduce the continuous diameter. Furthermore, we characterize optimal pairs of shortcuts for convex and non-convex cycles. Finally, we develop a linear time algorithm that produces an optimal pair of shortcuts for convex cycles. Apart from the algorithms, our results extend to rectifiable curves.

Our work reveals some of the underlying challenges that must be overcome when addressing the discrete version of this network augmentation problem, where we minimize the discrete diameter of a network with shortcuts that connect only vertices.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, G.2.2 Graph Theory, I.3.5 Computational Geometry and Object Modeling

Keywords and phrases Network Augmentation, Shortcuts, Diameter, Paths, Cycles

Digital Object Identifier 10.4230/LIPIcs.SWAT.2016.27

1 Introduction

The minimum-diameter network augmentation problem is concerned with minimizing the largest distance between two vertices of an edge-weighted graph by introducing new edges as shortcuts. We study this problem from a new perspective in a continuous and geometric setting where the network is a geometric graph embedded into the Euclidean plane, the weight of a shortcut is the Euclidean distance of its endpoints, and shortcuts can be introduced between any two points along the network that may be vertices or points along edges.

* This work was partially supported by NSERC.

15th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2016).
Editor: Rasmus Pagh; Article No. 27; pp. 27:1–27:14
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

As a sample application, consider a network of highways where we measure the distance between two locations in terms of the travel time. An urban engineer might want to improve the worst-case travel time along a highway or along a ring road by introducing shortcuts. Our work advises where these shortcuts should be built. For example, we show where to find the best shortcut for a highway and we show that a ring road requires two shortcuts.

1.1 Preliminaries

A *network* is an undirected graph that is embedded into the Euclidean plane and whose edges are weighted with their Euclidean length. For our algorithms we focus on networks with straight line edges, whereas the remaining results require rectifiable curves as edges. We say a point p lies on a network G and write $p \in G$ when there is an edge e of G such that p is a point along the embedding of e. A point p on an edge e of length l subdivides e into two sub-edges lengths $(1-\lambda) \cdot l$ and $\lambda \cdot l$ for some value $\lambda \in [0,1]$. We represent the points on G in terms of their relative position (expressed by λ) along their containing edge.

The *network distance* between two points p and q on a network G is the length of a weighted shortest path from p to q in G. We denote the network distance between p and q by $d_G(p,q)$ and we omit the subscript when the network is understood. The largest network distance between any two points on G is the *continuous diameter* of G, denoted by $\operatorname{diam}(G)$, i.e., $\operatorname{diam}(G) = \max_{p,q \in G} d_G(p,q)$. The term *continuous* distinguishes this notion from the *discrete diameter* that measures the largest network distance between any two vertices.

We denote the Euclidean distance between p and q by $|pq|$. A line segment pq, with $p, q \in G$ is a *shortcut* for G when $|pq| < d_G(p,q)$. We augment a network G with a shortcut pq as follows. We introduce new vertices at p and at q in G, subdividing their containing edges, and we add an edge from p to q of length $|pq|$. We do not introduce any vertices at any crossings between pq and other edges of G. The resulting network is denoted by $G + pq$. We seek to minimize the continuous diameter of a network by introducing shortcuts.

When G is a path or cycle, $|G|$ denotes its length. A cycle C is convex, when C forms a convex polygon with non-empty interior, i.e., a convex cycle cannot be confined to a line.

1.2 Related Work

In the discrete abstract setting, we consider an abstract graph G with unit weights and ask whether we can decrease the discrete diameter of G to at most D by adding at most k edges. For any fixed $D \geq 2$, this problem is NP-hard [2, 7, 9], has parametric complexity W[2]-hard [4, 5], and remains NP-hard even if G is a tree [2]. For an overview of the approximation algorithms in terms of both D and k refer, for instance, to Frati *et al.* [4].

In the discrete geometric setting, we consider a geometric graph, where a shortcut connects two vertices. Große *et al.* [6] are the first to consider diameter minimization in this setting. They determine a shortcut that minimizes the discrete diameter of a path with n vertices in $O(n \log^3 n)$ time. The spanning ratio of a geometric network, i.e., the largest ratio between the network distance and the Euclidean distance of any two points, has been considered as target function for edge augmentation, as well. For instance, Farshi *et al.* [3] compute a shortcut that minimizes the spanning ratio in $O(n^4)$ time while Luo and Wulff-Nilsen [8] compute a shortcut that maximizes the spanning ratio in $O(n^3)$ time.

1.3 Structure and Results

Our results concern networks that are paths, cycles, and convex cycles. Figures 1 and 2 illustrate examples of optimal shortcuts for paths and cycles. In Section 2, we develop an

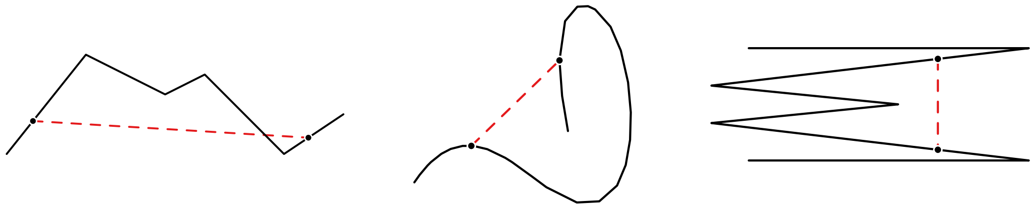

Figure 1 Examples for paths with an optimal shortcut.

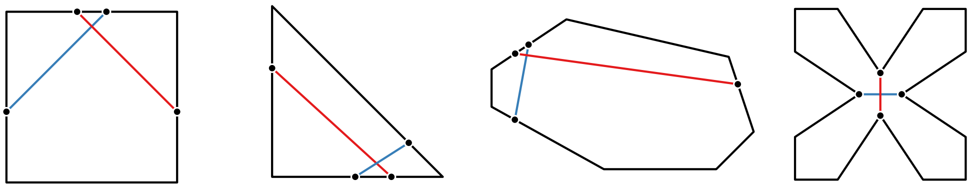

Figure 2 Examples for cycles with optimal pairs of shortcuts.

algorithm that produces an optimal shortcut for a path with n vertices in $O(n)$ time. In Section 3, we show that for cycles a single shortcut never suffices to reduce the diameter and that two shortcuts always suffice. We characterize pairs of optimal shortcuts for convex and non-convex cycles. Based on this characterization, we develop an algorithm in Section 4 that determines an optimal pair of shortcuts for a convex cycle with n vertices in $O(n)$ time.

The full version of this paper [1] contains all proofs that were omitted in this version.

2 Shortcuts for Paths

Consider a polygonal path P in the plane with n vertices. We seek a shortcut pq for a path P that minimizes the continuous diameter of the augmented path $P + pq$, i.e.,

$$\mathrm{diam}(P + pq) = \min_{a,b \in P} \mathrm{diam}(P + ab) = \min_{a,b \in P} \max_{u,v \in P+ab} d_{P+ab}(u, v) \ .$$

For this section, let s and e be the endpoints of P and let p be closer to s than q along P, i.e., $d(s, p) < d(s, q)$, as illustrated in Figure 3. For $a, b \in P$, let $P[a, b]$ denote the sub-path from a to b along P, and let $C(p, q)$ be the simple cycle in $P + pq$.

▶ **Lemma 1.** *Let pq be a shortcut for P. Every continuous diametral path in $P + pq$ contains an endpoint of P, except when the shortcut connects the endpoints of P.*

According to Lemma 1, we have the following three candidates for continuous diametral paths in the augmented network $P + pq$, two of which are illustrated in Figure 4.

1. The path $U(p, q)$ from s to e via the shortcut pq,
2. the path $S(p, q)$ from s to the farthest point from s on $C(p, q)$, and
3. the path $E(p, q)$ from e to the farthest point from e on $C(p, q)$.

Let $\bar{p}$ be the farthest point from p on $C(p, q)$, and let $\bar{q}$ be the farthest point from q on $C(p, q)$. Furthermore, let $\delta(p, q) := \frac{d(p,q)-|pq|}{2}$ denote the slack between p and $\bar{q}$ (and symmetrically between $\bar{p}$ and q) along $C(p, q)$. With this notation, we have

$$d(p, \bar{p}) = d(q, \bar{q}) = \frac{|C(p, q)|}{2} = \frac{d(p, q) + |pq|}{2} = \frac{d(p, q) - |pq|}{2} + |pq| = \delta(p, q) + |pq| \ ,$$

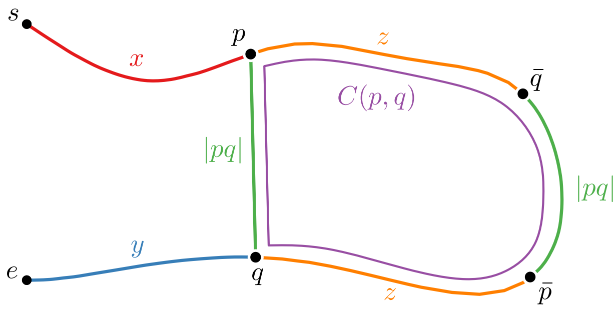

Figure 3 Augmenting a path P with a shortcut pq. The shortcut creates a cycle $C(p,q)$ with the sub-path from p to q along P. The farthest point from p on this cycle is $\bar{p}$ and $\bar{q}$ is farthest from q on $C(p,q)$. The distance $d(\bar{q},\bar{p})$ between $\bar{q}$ and $\bar{p}$ along P matches the Euclidean distance between p and q, because of the following. When we move a point g from p to q along the shortcut pq, then the farthest point $\bar{g}$ from g along $C(p,q)$ moves from $\bar{p}$ to $\bar{q}$ traveling the same distance as g, i.e., $|pq|$.

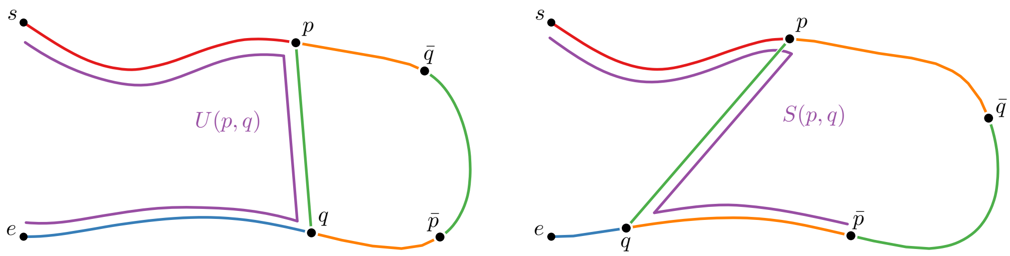

(a) The path $U(p,q)$ from s to e.

(b) The path $S(p,q)$ from s to the cycle.

Figure 4 Two candidate diametral paths in $P+pq$, namely the shortest path connecting s and e in a and a path from s via p to the farthest point from p on the cycle $C(p,q)$ in b. For the latter, there is a second path $S'(p,q)$ of the same length traversing $C(p,q)$ in the other direction.

and we can express the lengths of $U(p,q)$, $S(p,q)$, and $E(p,q)$ as follows.

$$|U(p,q)| = d(s,p) + |pq| + d(q,e)$$
$$|S(p,q)| = d(s,p) + d(p,\bar{p}) = d(s,p) + |pq| + \delta(p,q)$$
$$|E(p,q)| = d(e,q) + d(q,\bar{q}) = d(e,q) + |pq| + \delta(p,q)$$

The following lemma characterizes which of the paths $U(p,q)$, $S(p,q)$, and $E(p,q)$ determine the diameter of $P+pq$. Notice that these cases overlap, for instance, $E(p,q)$ and $S(p,q)$ are both continuous diametral when $d(s,p) = d(e,q) \leq \delta(p,q)$.

▶ **Lemma 2.** *Let pq be a shortcut for a path P. Let $x = d(s,p)$, $y = d(e,q)$, and $z = \delta(p,q)$.*
- *The path $U(p,q)$ is continuous diametral if and only if $z = \min(x,y,z)$.*
- *The path $S(p,q)$ is continuous diametral if and only if $y = \min(x,y,z)$.*
- *The path $E(p,q)$ is continuous diametral if and only if $x = \min(x,y,z)$.*

▶ **Lemma 3.** *For every path P, there is an optimal shortcut pq such that $S(p,q)$ and $E(p,q)$ are continuous diametral paths in $P+pq$, i.e.,* $\operatorname{diam}(P+pq) = |S(p,q)| = |E(p,q)|$.

According to Lemmas 2 and 3, we can restrict our search for an optimal shortcut to those shortcuts satisfying $d(s,p) = d(e,q) \leq \delta(p,q)$. For $x \in [0, |P|/2]$, let $p(x)$ and $q(x)$ be the points on P such that $x = d(s,p(x))$ and $x = d(e,q(x))$, and let $D(x) = |p(x)q(x)|$. Notice

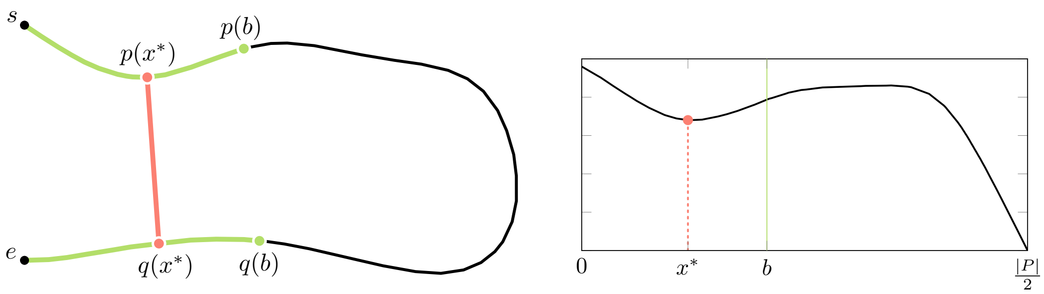

(a) A path P with its optimal shortcut $p(x^*)q(x^*)$. **(b)** A plot of $D(x) = |p(x)q(x)|$.

Figure 5 The optimal shortcut for the path in a with the function $D(x)$ plotted in b.

that $d(p(x), q(x)) = |P| - 2x$. Using this notation, we phrase our problem as

$$\min\ x + \frac{d(p(x), q(x)) + |p(x)q(x)|}{2} = \min\ x + \frac{|P| - 2x + |p(x)q(x)|}{2} = \min\ \frac{|P| + D(x)}{2}$$
$$\text{s.t. } x \leq \delta(p(x), q(x)) = \frac{d(p(x), q(x)) - |p(x)q(x)|}{2} = \frac{|P| - 2x - D(x)}{2}\ ,$$

which simplifies to minimizing $D(x)$ such that $4x + D(x) \leq |P|$.

▶ **Lemma 4.** *The function $B(x) = 4x + D(x)$ is strictly increasing on $[0, |P|/2]$.*

The following theorem describes an optimal shortcut and is illustrated in Figure 5.

▶ **Theorem 5.** *Let P be a path and let b be the unique value in $[0, |P|/2]$ with $B(b) = |P|$. Suppose D has a global minimum in the interval $[0, b]$ at x^*, i.e., $D(x^*) = \min_{x \in [0,b]} D(x)$. Then the shortcut $p(x^*)q(x^*)$ achieves the minimum continuous diameter for P.*

▶ **Lemma 6.** *Let P be a path with n vertices. Then $D^2(x)$ is a continuous function whose graph consists of at most n parabolic arcs or line segments.*

▶ **Corollary 7.** *Given a path P with n vertices, we can compute a shortcut for P achieving the minimal continuous diameter in $O(n)$ time.*

Proof. Let $x_{\pi(1)} \leq x_{\pi(2)} \leq \cdots \leq x_{\pi(n)}$ be the values in $[0, |P|/2]$ where $p(x_{\pi(i)})$ or $q(x_{\pi(i)})$ coincides with the i-th vertex of P for each $i = 1, 2, \ldots, n$.

We compute the minimum of the parabolic arc of D^2 on each interval $[x_{\pi(i)}, x_{\pi(i+1)}]$ for $i = 1, 2, \ldots, n$ until we arrive at k with $B(x_{\pi(k)}) < |P|$ and $B(x_{\pi(k+1)}) \geq |P|$. We then compute b by solving the quadratic equation $D^2(b) = (|P| - 4b)^2$ and, finally, compute the minimum of $D^2(x)$ on $[x_{\pi(k)}, b]$. The lowest minima of the encountered parabolic arcs is the global minimum of D on $[0, b]$, which reveals the position of an optimal shortcut according to Theorem 5. Altogether, the running time is $\Theta(k + 1) = O(n)$, since we obtain $x_{\pi(1)}, x_{\pi(2)}, \ldots, x_{\pi(k+1)}$ by merging the vertices by their distances from s or from e. ◀

▶ Remark. Our result on the location of an optimal shortcut from Theorem 5 also holds for rectifiable curves in the plane. However, obtaining an optimal shortcut for such curves depends on our ability to calculate b and a global minimum of $D(x)$ in the interval $[0, b]$.

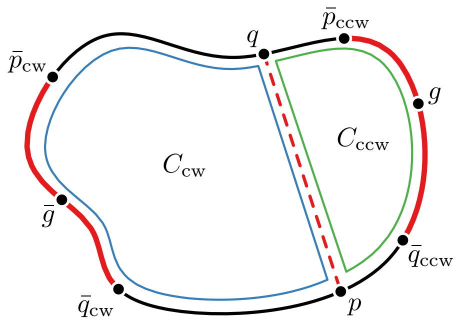

Figure 6 The unaffected regions (solid red) for a shortcut pq (dashed red) to a cycle. The point $\bar{x}_{\text{y}}$ denotes the farthest points from x along the cycle C_{y} for $x \in \{p, q\}$ and $\text{y} \in \{\text{cw}, \text{ccw}\}$. Any point g along the clockwise path from $\bar{p}_{\text{ccw}}$ to $\bar{q}_{\text{ccw}}$ has their farthest point $\bar{g}$ on the clockwise path from $\bar{q}_{\text{cw}}$ to $\bar{p}_{\text{cw}}$ and vice versa. The distance between g and $\bar{g}$ is unaffected by the addition of pq to C.

3 Shortcuts for Cycles

Consider a polygonal cycle C in the plane that may have crossings. For any two points p and q along C that may be vertices or points along edges of C, let $d_{\text{ccw}}(p, q)$ and $d_{\text{cw}}(p, q)$ be their counter-clockwise and clockwise distance along C, respectively. Let $d(p, q) = \min(d_{\text{ccw}}(p, q), d_{\text{cw}}(p, q))$ denote the network distance between p and q along C. We seek to minimize the continuous diameter by augmenting C with shortcuts.

▶ **Lemma 8.** *Adding a single shortcut pq to a polygonal cycle C never decreases the continuous diameter, i.e.,* $\text{diam}(C) = \text{diam}(C + pq)$ *for all* $p, q \in C$.

Proof Sketch. Consider any shortcut pq to a cycle C. Let C_{ccw} be the cycle consisting of pq and the counter-clockwise path from p to q along C, as illustrated in Figure 6. Let $\bar{p}_{\text{ccw}}$ and $\bar{q}_{\text{ccw}}$ be the farthest points from p and from q on C_{ccw}, respectively. Since $\bar{p}_{\text{ccw}}$ and $\bar{q}_{\text{ccw}}$ are antipodal from p and q in C_{ccw}, we have $d(\bar{q}_{\text{ccw}}, \bar{p}_{\text{ccw}}) = |pq|$ and $d(\bar{p}_{\text{ccw}}, q) = d(p, \bar{q}_{\text{ccw}})$.

Consider a point g along the clockwise path from $\bar{p}_{\text{ccw}}$ to $\bar{q}_{\text{ccw}}$ and let $\bar{g} \in C$ be the farthest point from g with respect to C. We can show that $d_C(g, \bar{g}) = d_{C+pq}(g, \bar{g})$. ◀

According to Lemma 8, some points preserve their farthest distance in C when adding a single shortcut pq to C. The points that are unaffected by pq in this sense form the *unaffected region* of pq that consists of the clockwise path from $\bar{p}_{\text{ccw}}$ to $\bar{q}_{\text{ccw}}$ and the clockwise path from $\bar{q}_{\text{cw}}$ to $\bar{p}_{\text{cw}}$, as illustrated in Figure 6. Conversely, every point on C outside of the unaffected region uses pq as a shortcut to their farthest point in $C + pq$.

Consequently, we have to add at least two shortcuts pq and rs in order to decrease the continuous diameter of the augmented cycle $C + pq + rs$. Figure 2 illustrates examples of optimal pairs of shortcuts for cycles. We call a pair of shortcuts pq and rs *useful* when $\text{diam}(C) > \text{diam}(C + rs + pq)$, and we call pq and rs *useless*, otherwise. A pair of shortcuts pq and rs is useful if and only if their unaffected regions are disjoint.

We call a polygonal cycle C *degenerate* when it consists of two congruent line segments of length $|C|/2$. Any number of shortcuts cannot decrease the diameter of a degenerate cycle, since the endpoints of its line segment remain at the same distance.

▶ **Theorem 9.** *For every non-degenerate cycle C, there exists a pair of shortcuts pq and rs that decrease the continuous diameter, i.e.,* $\text{diam}(C) > \text{diam}(C + pq + rs)$.

(a) Alternating Configuration **(b)** Consecutive Configuration

Figure 7 The two cases for adding two shortcuts pq and rs to a cycle C. The endpoints of the shortcuts appear in alternating cyclic order p, r, q, and s, as shown in a, or in consecutive cyclic order p, q, r, and s, as shown in b. The two cases overlap when q coincides with r.

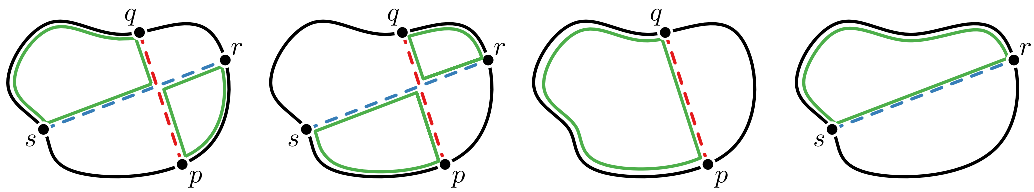

(a) The bowtie ($\bowtie$). **(b)** The hourglass (⧖). **(c)** The red split (⦸). **(d)** The blue split ($\oslash$).

Figure 8 The candidate diametral cycles, except C, for shortcuts in the alternating configuration.

Proof Sketch. Suppose there exist three points p, q, and s on C with $d(p, q) = d(q, s) = |C|/4$ such that pq and qs are shortcuts, i.e., $|pq| < d(p, q)$ and $|qs| < d(q, s)$. We can argue that pq and qs are useful. Conversely, we can show that C is degenerate when at least one of pq and qs is not a shortcut for every three points p, q, and s on C with $d(p, q) = d(q, s) = |C|/4$. ◀

3.1 Alternating vs. Consecutive

When placing two shortcuts pq and rs on a cycle C, we distinguish whether their endpoints appear in alternating order or in consecutive order along the cycle, as illustrated in Figure 7.

We show that there is always an optimal pair of shortcuts in the alternating configuration by studying the cycles created by the insertion of the shortcuts. A cycle in $C + pq + rs$ is *diametral* when it contains a diametral pair. Each configuration has five candidates for diametral cycles: two that use both shortcuts, two that use one of the shortcuts, and one (C) that does not use any shortcut. Figures 8 and 9 illustrate the candidates for diametral cycles in each configuration, except for C itself. To distinguish the cycles using one shortcut, we color pq red and rs blue and we refer to the longer cycle in $C + pq + rs$ using the red shortcut pq as the *red split* and we refer to the longer cycle using the blue shortcut rs as the *blue split*. If C happens to be diametral in $C + pq + rs$, then our pair of shortcuts is useless.

▶ **Lemma 10.** *Two shortcuts pq and rs in alternating configuration are useful if and only if $|pq| + |rs| < d_{ccw}(r, q) + d_{ccw}(s, p)$ and $|pq| + |rs| < d_{ccw}(p, r) + d_{ccw}(q, s)$.*

Proof. Suppose pq and rs are useless, i.e., the unaffected regions of pq and rs overlap. In the alternating configuration, this overlap occurs along the bowtie or along the hourglass. Since these cases are symmetric, we consider only the former in the following.

An overlap on the bowtie manifests along the clockwise path from $\bar{r}_{\text{ccw}}$ to $\bar{p}_{\text{cw}}$ with a mirrored overlap along the clockwise path from $\bar{s}_{\text{cw}}$ to $\bar{q}_{\text{ccw}}$, as illustrated in Figure 10. This means the sum of the lengths of the counter-clockwise paths from r to $\bar{r}_{\text{ccw}}$ and from $\bar{p}_{\text{cw}}$ to p is

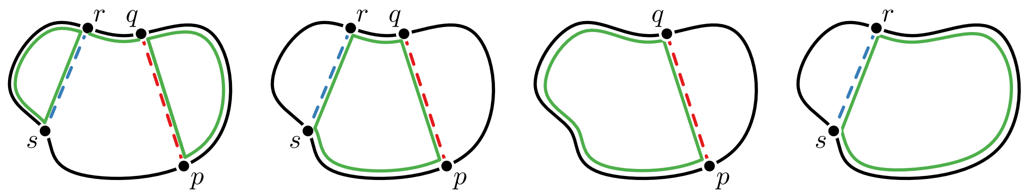

(a) The handset. **(b)** The base station. **(c)** The red split. **(d)** The blue split.

Figure 9 The candidate diametral cycles, except C, for shortcuts in the consecutive configuration, depicted for $d(q,r) \leq d(s,p)$. Even though the handset a is no simple cycle, it might still contain a diametral pair. Observe that the base station b is only listed for the sake of completeness: by the triangle inequality, this cycle is never longer than the split cycles and, therefore, never diametral.

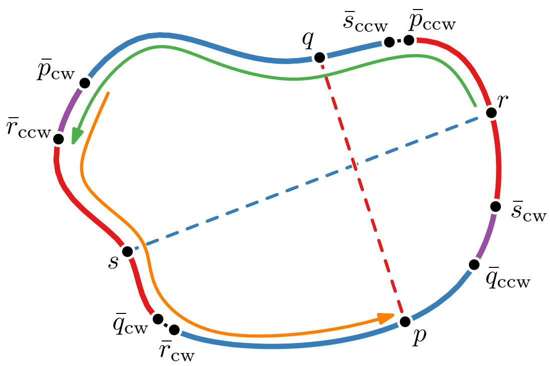

Figure 10 A pair of useless shortcuts whose unaffected regions have an overlap (purple) along the bowtie, i.e., the points s, $\bar{r}_{\text{ccw}}$, $\bar{p}_{\text{cw}}$, and q appear clockwise in this order along the cycle.

at least the length of the counter-clockwise path from r to p, i.e., $d_{\text{ccw}}(\bar{p}_{\text{cw}}, p) + d_{\text{ccw}}(r, \bar{r}_{\text{ccw}}) \geq d_{\text{ccw}}(r,p)$. This is equivalent to $|pq| + |rs| \geq d_{\text{ccw}}(r,q) + d_{\text{ccw}}(s,p)$, since

$$d_{\text{ccw}}(q,p) + |pq| + d_{\text{ccw}}(r,s) + |rs| = 2d_{\text{ccw}}(\bar{p}_{\text{cw}}, p) + 2d_{\text{ccw}}(r, \bar{r}_{\text{ccw}}) \geq 2d_{\text{ccw}}(r,p)$$
$$\iff |pq| + |rs| \geq \underbrace{d_{\text{ccw}}(r,p) - d_{\text{ccw}}(q,p)}_{=d_{\text{ccw}}(r,q)} + \underbrace{d_{\text{ccw}}(r,p) - d_{\text{ccw}}(r,s)}_{=d_{\text{ccw}}(s,p)} .$$

Analogously, we derive that $|pq| + |rs| \geq d_{\text{ccw}}(p,r) + d_{\text{ccw}}(q,s)$ holds if and only if there is an overlap along the hourglass. Consequently, the shortcuts pq and rs are useful if and only if $|pq| + |rs| < d_{\text{ccw}}(r,q) + d_{\text{ccw}}(s,p)$ and $|pq| + |rs| < d_{\text{ccw}}(p,r) + d_{\text{ccw}}(q,s)$. ◀

▶ **Lemma 11.** *Consider two consecutive shortcuts pq and rs with $d_{ccw}(q,r) \leq d_{ccw}(s,p)$. Then pq and rs are useful if and only if $|pq| + |rs| < d_{ccw}(s,p) - d_{ccw}(q,r)$.*

▶ **Theorem 12.** *Let pq and rs be a pair of shortcuts for a cycle C in consecutive configuration. There exists a pair $p'q'$ and $r's'$ of shortcuts in the alternating configuration that are at least as good as pq and rs, i.e.,* $\operatorname{diam}(C + p'q' + r's') \leq \operatorname{diam}(C + pq + rs)$.

Proof Sketch. Suppose pq and rs are useful shortcuts in the consecutive configuration. Assume, without loss of generality, $d_{\text{ccw}}(q,r) \leq d_{\text{ccw}}(s,p)$ and $d_{\text{ccw}}(p,q) \leq d_{\text{ccw}}(r,s)$. We consider the shortcuts $p'q' = pr$ and $r's' = rs$, which are illustrated in Figure 11 and lie in the intersection of the alternating and consecutive case. We argue that pr and rs are useful shortcuts and that each candidate diametral cycle in $C + pq + rs$ has a one-to-one correspondence to a candidate diametral cycle in $C + pr + rs$ of smaller or equal length. ◀

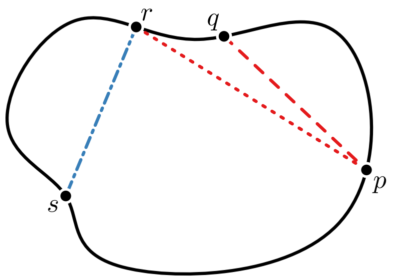

Figure 11 Replacing consecutive shortcuts pq and rs with alternating $p'q' = pr$ and $r's' = rs$.

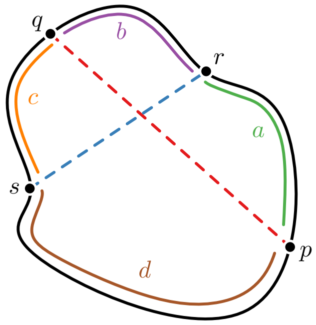

Figure 12 The sections of a cycle with a pair of alternating shortcuts.

3.2 Balancing Diametral Cycles

We show that every cycle has an optimal pair of alternating shortcuts where the bowtie and the hourglass are both diametral and we show that every convex cycle has an optimal pair of shortcuts where both split cycles are diametral, as well. We obtain these results by applying a sequence of operations – some of which are shown in Figure 14 – that each slide the shortcuts along the cycle in a way that reduces or maintains the continuous diameter and brings the candidate diametral cycles closer to the desired balance. The last two operations only reduce the diameter for convex cycles as the shortcuts might get stuck at reflex vertices, which leads to our characterization of optimal shortcuts for convex and non-convex cycles.

Let pq and rs be two alternating shortcuts and let $a = d_{\text{ccw}}(p,r)$, $b = d_{\text{ccw}}(r,q)$, $c = d_{\text{ccw}}(q,s)$, and $d = d_{\text{ccw}}(s,p)$. We assume that the red split contains s and the blue split contains p, i.e., $a + b \leq c + d$ and $b + c \leq a + d$, as in Figure 12. We abbreviate the lengths of the bowtie ($\bowtie$), the hourglass ($\hourglass$), the red split ($\obslash$), and the blue split ($\oslash$) as follows.

$$\bowtie := a + c + |pq| + |rs| \quad \obslash := c + d + |pq| \quad \hourglass := b + d + |pq| + |rs| \quad \oslash := a + d + |rs|$$

▶ **Lemma 13.** *For each relation $\sim \in \{<, =, >\}$, we have*

$$\begin{aligned} \bowtie \sim \hourglass &\iff a + c \sim b + d & \obslash \sim \oslash &\iff c + |pq| \sim a + |rs| \\ \bowtie \sim \obslash &\iff a + |rs| \sim d & \hourglass \sim \obslash &\iff b + |rs| \sim c \\ \bowtie \sim \oslash &\iff c + |pq| \sim d & \hourglass \sim \oslash &\iff b + |pq| \sim a \end{aligned}$$

and pq and rs are useful if and only if $|pq| + |rs| < a + c$ and $|pq| + |rs| < b + d$.

Proof. The claims follow from the definitions of $\bowtie$, $\hourglass$, $\obslash$, and $\oslash$. ◀

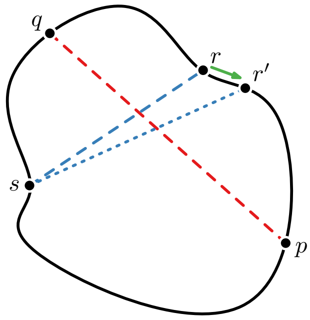

Figure 13 Shrinking the blue split.

▶ **Lemma 14.** *Consider a pair of useful alternating shortcuts where one of the split cycles evenly divides the cycle. Then this split cycle must have length at most* $\bowtie$ *or at most* $\hourglass$*.*

Proof. Assume that we have a pair of useful shortcuts where $\obslash$ divides the cycle evenly, i.e., $a + b = c + d$, and where $\bowtie < \obslash$ and $\hourglass < \obslash$. Then $a + |rs| < d$ and $b + |rs| < c$, by Lemma 13, which yields the contradiction $|rs| < 0$, as $a + |rs| < d = a + b - c < a - |rs|$. ◀

▶ **Lemma 15.** *There exists an optimal pair of shortcuts in alternating configuration such that none of the split cycles is the only diametral cycle.*

Proof Sketch. Suppose pq and rs are useful and $\oslash$ is the only diametral cycle in $C + pq + rs$, i.e., $\bowtie, \hourglass, \obslash < \oslash$. Then we have $b + |pq| < a$ and $c + |pq| - |rs| < a$. We move r clockwise until we arrive at some r' where $b' + |pq| = a'$ or $c + |pq| - |r's| = a'$, as in Figure 13.

Since the blue split is diametral, it cannot divide the cycle evenly, by Lemma 14. Therefore, changing r to r' changes the candidate diametral cycles as follows.

- The blue split shrinks or remains the same, i.e., $\oslash \geq \oslash'$.
- The red split remains the same, i.e., $\obslash = \obslash'$.
- The bowtie changes as the blue split, i.e., $\bowtie \geq \bowtie'$ and $\bowtie' < \oslash'$.
- The hourglass remains the same or increases, i.e., $\hourglass \leq \hourglass'$.
- The hourglass increases when the blue split remains the same, i.e, $\oslash - \hourglass > \oslash' - \hourglass'$.

Consequently, $\mathrm{diam}(C + pq + r's) \leq \mathrm{diam}(C + pq + rs)$ and $\oslash' = \hourglass'$ or $\oslash' = \obslash'$, which implies our claim, provided that pq and $r's$ are useful shortcuts. ◀

▶ **Lemma 16.** *There exists a pair of optimal shortcuts with* $\bowtie = \hourglass$*.*

Proof Sketch. Suppose pq and rs are useful shortcuts with $\bowtie \neq \hourglass$. We balance $\bowtie$ and $\hourglass$ using the following operations, as illustrated in Figure 14. They maintain or decrease the continuous diameter while decreasing the difference between bowtie and hourglass.

1. As long as neither split cycle divides the cycle C evenly, we shrink the larger split cycle in a way that decreases the difference of bowtie and hourglass:
   a. When $\bowtie < \hourglass$ and $\obslash \leq \oslash$, we move s counter-clockwise.
   b. When $\bowtie < \hourglass$ and $\obslash > \oslash$, we move p clockwise.
   c. When $\bowtie > \hourglass$ and $\obslash \leq \oslash$, we move r clockwise.
   d. When $\bowtie > \hourglass$ and $\obslash > \oslash$, we move q counter-clockwise.
2. Once a split cycle evenly divides the cycle, we move the endpoints of the corresponding shortcut in the direction that decreases the difference between bowtie and hourglass:

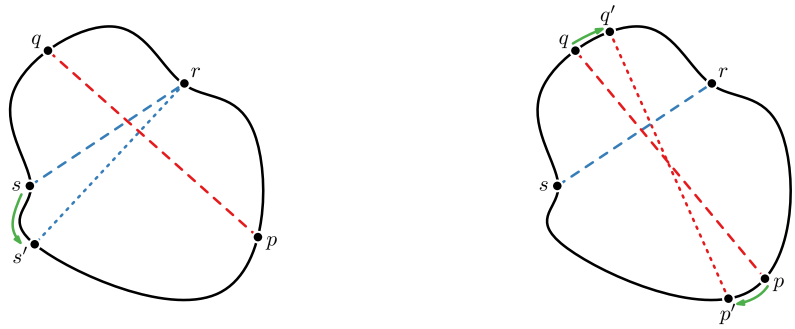

(a) Operation 1.a. **(b)** Operation 2.a.

Figure 14 Some of the operations that are used to balance the candidate diametral cycles.

a. When $\bowtie < ⧗$ and $\obslash$ evenly divides the cycle, we move p and q clockwise.
b. When $\bowtie > ⧗$ and $\obslash$ evenly divides the cycle, we move p and q counter-clockwise.
c. When $\bowtie < ⧗$ and $\oslash$ evenly divides the cycle, we move s and r counter-clockwise.
d. When $\bowtie > ⧗$ and $\oslash$ evenly divides the cycle, we move s and r clockwise.

For each operation, we argue that pq and rs remain useful shortcuts and that the diameter never increases while the difference between hourglass and bowtie always decreases. ◀

▶ **Corollary 17.** *There exists a pair of optimal shortcuts that is in the alternating configuration such that none of the split cycles is the only diametral cycle and such that the bowtie and the hourglass have the same length.*

Proof. By Lemma 15, we have a pair of optimal shortcuts pq and rs for a cycle C where none of the splits is the only diametral cycle. The Operations 1.a to 1.d from Lemma 16 shrink the larger split cycle at the same rate as they shrink the larger of bowtie and hourglass. Thus, we do not create a sole diametral cycle by applying these operations. Furthermore, the Operations 2.a to 2.d rotate an even split that cannot become diametral by Lemma 14. By applying Lemma 16, we obtain a pair of optimal shortcuts $p'q'$ and $r's'$ with at least two diametral cycles and where bowtie and hourglass have the same length. ◀

▶ **Theorem 18.** *For every non-degenerate cycle, there exists an optimal pair of shortcuts such that the hourglass and the bowtie are both diametral.*

Proof. Let C be a non-degenerate cycle. By Corollary 17, there is a pair of optimal shortcuts pq and rs where neither split cycle is the only diametral cycle and where $\bowtie = ⧗$.

Suppose that $\bowtie$ and $⧗$ are not diametral. The cycle C cannot be diametral, since pq and rs are useful. This means a split is diametral, i.e., $\obslash > \bowtie = ⧗$ or $\oslash > \bowtie = ⧗$. Since neither $\obslash$ nor $\oslash$ is the only diametral cycle, we have $\obslash = \oslash > \bowtie = ⧗$.

We shrink the splits by simultaneously moving p and r clockwise while moving q and s counter-clockwise, as illustrated in Figure 15. By moving pq and rs at appropriate speeds, we ensure that this operation maintains both the balance between the split cycles and the balance between bowtie and hourglass, i.e., $\obslash' = \oslash'$ and $\bowtie' = ⧗'$. This decreases the continuous diameter, provided that the line segments $p'q'$ and $r's'$ remain useful shortcuts.

Assume, for the sake of a contradiction, that $p'q'$ is not a shortcut, i.e., $|p'q'| = d(p'q')$. By Lemma 14 we cannot pass through an even red split during our operation. Thus, we have $d(p', q') = d_{\text{ccw}}(p', q')$, i.e., the line segment $p'q'$ contains pq contradicting the choice of pq as shortcut. Therefore, $p'q'$ is a shortcut. Symmetrically, we can argue that $r's'$ is a shortcut.

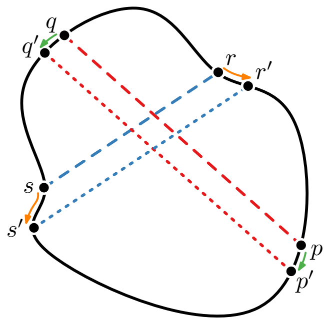

Figure 15 Shifting the shortcuts to shrink the split cycles while maintaining $⦸ = \oslash$ and $\bowtie = ⧖$.

We argue that $p'q'$ and $r's'$ remain useful. From $\bowtie' = ⧖' \leq ⦸' = \oslash'$, we obtain $|p'q'| \leq d' - c'$, $|p'q'| \leq a' - b'$, and $|r's'| \leq c' - b'$, by Lemma 13. Together with $b' > 0$ and $a' + c' = b' + d'$, we derive that $p'q'$ and $r's'$ are useful, because $|p'q'| + |r's'| \leq d' - c' + c' - b' = d' - b' < d' + b' = a' + c'$. This means $p'q'$ and $r's'$ are useful shortcuts with $\bowtie' = ⧖' = ⦸' = \oslash'$ and $\operatorname{diam}(C + pq + rs) > \operatorname{diam}(C + p'q' + r's')$ contradicting the optimality of pq and rs. Therefore, there are optimal shortcuts where the hourglass and the bowtie are diametral. ◀

▶ **Theorem 19.** *For every convex cycle, there exists an optimal pair of alternating shortcuts such that the hourglass, the bowtie, and the splits are diametral, i.e.,* $\bowtie = ⧖ = ⦸ = \oslash$.

Proof Sketch. According to Theorem 18, there are optimal shortcuts pq and rs with $\bowtie = ⧖ \geq ⦸$ and $\bowtie = ⧖ \geq \oslash$. Suppose we have $\bowtie = ⧖ > ⦸$ or $\bowtie = ⧖ > \oslash$.

First, since C is convex we can increase each split in a way that shrinks its shortcut. Second, we grow the smaller split until both splits are equal. Third, we grow both splits at the same rate until they are equal to bowtie and hourglass. ◀

▶ **Corollary 20.** *For every non-degenerate cycle, there exists an optimal pair of shortcuts such that the hourglass and the bowtie are diametral and such that each split cycle is diametral or the shortcut of the split has at least one endpoint at a reflex vertex.*

▶ **Corollary 21.** *For every convex cycle, there exists an optimal pair of shortcuts with* $a + b \leq c + d$ *and* $b + c \leq a + d$ *such that the following holds.*

$$d = \frac{|C|}{4} + \frac{|pq| + |rs|}{2} = \operatorname{diam}(C + pq + rs) \qquad a = \frac{|C|}{4} + \frac{|pq| - |rs|}{2}$$
$$b = \frac{|C|}{4} - \frac{|pq| + |rs|}{2} = \operatorname{diam}(C) - \operatorname{diam}(C + pq + rs) \qquad c = \frac{|C|}{4} + \frac{|rs| - |pq|}{2}$$

Surprisingly, this means we can read the new continuous diameter of $C + pq + rs$ from d and we can read the benefit of adding the shortcuts pq and rs to C from b.

4 A Linear-Time Algorithm for Convex Cycles

For convex cycles, we restrict our search to the pairs of shortcuts satisfying $\bowtie = ⧖ = ⦸ = \oslash$, due to Theorem 19. We proceed as follows. First, we pick some point p on the cycle C and compute three points q, r, and s such that $\bowtie = ⧖ = ⦸ = \oslash$ – regardless of whether pq and rs are shortcuts. We show that the points q, r, and s exist and are unique for every point p along C. Once we have balanced p, q, r, and s, we slide p along C maintaining the balance by moving q, r, and s appropriately. We show that q, r, and s move in the same direction as

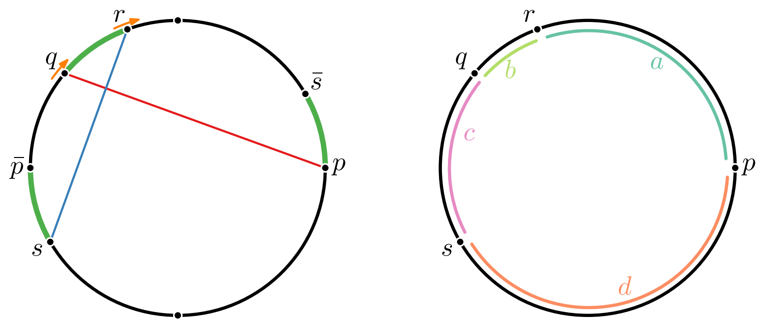

Figure 16 Locating r and s to balanced the splits for fixed p and s.

p while preserving their order along C. Thus, each endpoint traverses each edge of the n edges of C at most once throughout this process, which therefore takes $O(n)$ time.

For the remainder of this section, we only focus on convex cycles with non-empty interior. Consider a cycle C and a fixed point p on C. We say a triple of points q, r, and s is *in balanced configuration with* p when the points p, r, q, and s appear counter-clockwise in this order along C, $d_{\text{ccw}}(p,q) \leq |C|/2$, $d_{\text{ccw}}(r,s) \leq |C|/2$, and $\bowtie = ⧖ = ⦸ = \oslash$.

▶ **Theorem 22.** *Consider a convex cycle C and a point p on C. There exists a unique triple q, r, s of points on C that are in balanced configuration with p.*

Proof Sketch. Suppose we place s at some position on C with $|C|/4 \leq d_{\text{ccw}}(s,p) \leq |C|/2$. The points r and q must have fixed distance $d(r,q) = |C|/2 - d_{\text{ccw}}(s,p)$ to ensure $\bowtie = ⧖$. Suppose we slide q and r along the clockwise path from $\bar{p}$ to $\bar{s}$ while maintaining $d(q,r) = d(\bar{p},s)$, as in Figure 16. When q and r are close to s, we have $⦸ < \oslash$ and when q and r are close to p, we have $⦸ > \oslash$. By the intermediate value theorem, there exist positions for q and r such that $\bowtie = ⧖$ and $⦸ = \oslash$ and these positions are unique, since C is convex.

Suppose we slide s from $\bar{p}$ towards p while maintaining $\bowtie = ⧖$ and $⦸ = \oslash$. When $d(s,p) = |C|/2$, we end up with $\bowtie = ⧖ < ⦸ = \oslash$ and when $d(s,p) = |C|/4$, we end up with $\bowtie = ⧖ > ⦸ = \oslash$. By the intermediate value theorem, there exist positions for s, q, and r such that $\bowtie = ⧖ = ⦸ = \oslash$. We find these positions with two nested binary searches. ◀

▶ **Lemma 23.** *Consider a convex cycle C. Suppose p moves counter-clockwise along C. Then any three points in balanced configuration with p are moving counter-clockwise, as well.*

▶ **Theorem 24.** *Consider a convex cycle C with n vertices. We can compute an optimal pair of shortcuts for C in $O(n)$ time.*

Proof. We pick an arbitrary point p along some edge e_p of C and identify the edges e_q, e_r, and e_s containing the points q, r, and s that form a balanced configuration with p, as described in the proof sketch of Theorem 22. We find a (locally) optimal pair of shortcuts p^*q^* and r^*s^* with whose endpoints lie on the edges e_p, e_q, e_r, e_s by minimizing $d = \text{diam}(C + pq + rs)$ subject to $a + b \leq |C|/2$ and $b + c \leq |C|/2$, and the constraints stated in Corollary 21 that ensure $\bowtie = ⧖ = ⦸ = \oslash$. Then, we identify the four edges that would host p, q, r, and s next, if p were to move counter-clockwise: for each endpoint $x \in \{p,q,r,s\}$, we calculate how far the other endpoints would move under the assumption that x is the first point to hit a vertex. Theorem 22 guarantees that this calculation has a unique solution. Since all points move in the same direction as p, an edge e will never host an endpoint x in any subsequent step, once x has left e. Therefore, the entire process takes $O(n)$ time. Since we encounter every four points in balanced configuration, we also encounter an optimal pair of shortcuts. ◀

5 Conclusion and Future Work

Our work reveals some of the underlying challenges that must be overcome when addressing the discrete version of the network augmentation problem, where we minimize the discrete diameter of a network with shortcuts that connect only vertices. We shall investigate to what extend we can translate to the discrete setting. For instance, we would like to know when and how well the optimal continuous shortcuts approximate the optimal discrete shortcuts.

By Corollary 20, we can determine an optimal pair of shortcuts for non-convex cycles with r reflex vertices in $O(rn^3)$ time: we compute the best shortcuts satisfying $\bowtie = ⧖ = ⦸ = \oslash$ and we check all possible triples of edges that might contain the other endpoints when one shortcut is stuck at a reflex vertex. We seek to improve this naïve approach by generalizing our sliding-sweep algorithm for convex cycles to non-convex cycles. In addition the shortcuts with $\bowtie = ⧖ = ⦸ = \oslash$ (which may be non-optimal), we would have to keep track of each locally optimal shortcuts with one endpoint at a reflex vertex. However, some properties, such as the uniqueness of shortcuts in balance, break down for non-convex cycles.

As the next natural step after paths and cycles, we shall study minimizing the continuous diameter of trees, uni-cyclic networks, and so forth by introducing shortcuts.

References

1. Jean-Lou De Carufel, Carsten Grimm, Anil Maheshwari, and Michiel Smid. Minimizing the continuous diameter when augmenting paths and cycles with shortcuts. This is the full version of this paper., 2015. `arXiv:1512.02257`.
2. Victor Chepoi and Yann Vaxès. Augmenting trees to meet biconnectivity and diameter constraints. *Algorithmica*, 33(2):243–262, 2002. `doi:10.1007/s00453-001-0113-8`.
3. Mohammad Farshi, Panos Giannopoulos, and Joachim Gudmundsson. Improving the stretch factor of a geometric network by edge augmentation. *SIAM Journal on Computing*, 38(1):226–240, 2008. `doi:10.1137/050635675`.
4. Fabrizio Frati, Serge Gaspers, Joachim Gudmundsson, and Luke Mathieson. Augmenting graphs to minimize the diameter. *Algorithmica*, 72(4):995–1010, 2015. `doi:10.1007/s00453-014-9886-4`.
5. Yong Gao, Donovan R. Hare, and James Nastos. The parametric complexity of graph diameter augmentation. *Discrete Applied Mathematics*, 161(10-11):1626–1631, 2013. `doi:10.1016/j.dam.2013.01.016`.
6. Ulrike Große, Joachim Gudmundsson, Christian Knauer, Michiel Smid, and Fabian Stehn. Fast algorithms for diameter-optimally augmenting paths. In *42nd International Colloquium on Automata, Languages, and Programming (ICALP 2015)*, pages 678–688, 2015. `doi:10.1007/978-3-662-47672-7_55`.
7. Chung-Lun Li, S. Thomas McCormick, and David Simchi-Levi. On the minimum-cardinality-bounded-diameter and the bounded-cardinality-minimum-diameter edge addition problems. *Operations Research Letters*, 11(5):303–308, 1992. `doi:10.1016/0167-6377(92)90007-P`.
8. Jun Luo and Christian Wulff-Nilsen. Computing best and worst shortcuts of graphs embedded in metric spaces. In *19th International Symposium on Algorithms and Computation (ISAAC 2008)*, pages 764–775, 2008. `doi:10.1007/978-3-540-92182-0_67`.
9. Anneke A. Schoone, Hans L. Bodlaender, and Jan van Leeuwen. Diameter increase caused by edge deletion. *Journal of Graph Theory*, 11(3):409–427, 1987. `doi:10.1002/jgt.3190110315`.