

Optimal Data Structures for Farthest-Point Queries in Cactus Networks*

Prosenjit Bose[†] Jean-Lou De Carufel[†] Carsten Grimm^{†‡§}
Anil Maheshwari[†] Michiel Smid[†]

August 25, 2014

Abstract

Consider the continuum of points on the edges of a network, i.e., a connected, undirected graph with positive edge weights. We measure the distance between these points in terms of the weighted shortest path distance, called the *network distance*. Within this metric space, we study farthest points and farthest distances. We introduce optimal data structures supporting queries for the farthest distance and the farthest points on trees, cycles, uni-cyclic networks, and cactus networks.

1 Introduction

Consider the continuum of points on the edges of a network, i.e., a graph with positive edge weights. We measure the distance between these points in terms of the weighted shortest path distance, called the *network distance*. Within this metric space, we study farthest points and farthest distances.

Decisions where to place facilities are often complex and involve optimization of multiple criteria. Our data structures enable decision makers to quickly compare farthest distances from potential locations, which may constitute an essential factor, e.g., impacting emergency response times for a possible location of a new hospital. Furthermore, our results provide a *heat-map* of farthest distances illuminating the aspect of centrality of the network at hand and, thereby, serve as a visual aid for decision makers.

We introduce optimal data structures supporting queries for the farthest distance and the farthest points on trees, cycles, uni-cyclic networks, and cactus networks. We begin with data structures for simple networks and then use them as building blocks for more complex networks. With this modular approach we can easily extend our results as new building blocks become available.

The remainder of this section is organized as follows. In [Section 1.1](#), we set the stage for this work by making our notions of networks, points along networks, and network distance precise. In [Section 1.2](#), we summarize related work. In [Section 1.3](#), we outline our main contributions and the structure of this work.

*This research has been partially funded by NSERC and by a fellowship from the German Academic Exchange Service (DAAD).

A preliminary version of this work was presented at the 25th Canadian Conference on Computational Geometry [3] and was part of the Diplomarbeit (Master's thesis) of the third author [9].

[†]School of Computer Science, Carleton University

[‡]Institut für Simulation und Graphik, Fakultät für Informatik, Otto-von-Guericke-Universität Magdeburg

[§]Corresponding author carsten.grimm@ovgu.de

28 **1.1 Preliminaries and Problem Definition**

29 We call a simple, finite, undirected graph with positive edge weights a *network*. Unless stated otherwise, we
 30 consider only connected networks. Let $G = (V, E)$ be a network with n vertices and m edges, where V is
 31 the set of vertices and E is the set of edges. We write uv to denote an edge with endpoints $u, v \in V$ and
 32 we write w_{uv} to denote its weight. A point p on edge uv subdivides uv into two sub-edges up and pv with
 33 $w_{up} = \lambda w_{uv}$ and $w_{pv} = (1 - \lambda)w_{uv}$, where λ is the real number in $[0, 1]$ for which $p = \lambda u + (1 - \lambda)v$. We
 34 write $p \in uv$ when p is on edge uv and $p \in G$ when p is on some edge of G .




Figure 1: In the network shown in (a), the network distance from $p = \frac{1}{4}u + \frac{3}{4}v$ to $q = \frac{1}{2}s + \frac{1}{2}t$ is $d(p, q) = 10$. This distance is, for instance, achieved along the shortest path (blue) from p to q depicted in (b). The eccentricity of p is $\text{ecc}(p) = 12$ and the farthest point from p lies along the edge xs , as indicated by the shortest path tree from p (orange) and its extension (orange \cup green) illustrated in (c).

35 As illustrated in Figure 1, we measure distance between points $p, q \in G$ in terms of the weighted length
 36 of a shortest path from p to q in G , denoted by $d_G(p, q)$. We say that p and q have *network distance* $d_G(p, q)$.
 37 The points on G and the network distance form a metric space. Within this metric space, we study farthest
 38 points and farthest distances. We call the largest network distance from some point p on G the *eccentricity*
 39 of p and denote it by $\text{ecc}_G(p)$, i.e., $\text{ecc}_G(p) = \max_{q \in G} d_G(p, q)$. A point \bar{p} on G is farthest from p if and only
 40 if $d_G(p, \bar{p}) = \text{ecc}_G(p)$. We omit the subscript G whenever the network is understood from the context.

41 We extend the definition of shortest path trees from graphs to networks. When we say we *cut* an edge
 42 ab at a point $p \in ab$ we mean that we introduce two new vertices x and y and replace the edge ab with two
 43 edges ax and yb of weight $w_{ax} = w_{ap}$ and $w_{yb} = w_{pb}$. If p coincides with an endpoint of ab , then one of
 44 ax and yb has weight zero and is omitted. Let s be some point on a network G , and let T_s be the shortest
 45 path tree of s in G . We split each non- T_s edge ab at the farthest point from s on ab and add the resulting
 46 edges ax and yb to T_s . The resulting tree is called the *extended shortest path tree* [25] of s in G ; this tree
 47 encapsulates both the eccentricity of s in G and the farthest points from s in G , as illustrated in Figure 1c.

48 We aim to construct data structures for a fixed network G supporting the following queries. Given a
 49 point p on G , what is the eccentricity of p ? What is the set of farthest points from p in G ? We refer to the
 50 former as an *eccentricity query* and to the latter as a *farthest-point query*. Both queries consist of the query
 51 point p represented by the edge uv containing p and the value $\lambda \in [0, 1]$ such that $p = \lambda u + (1 - \lambda)v$.

52 We study trees, cycles, uni-cyclic networks, and cactus networks. A *uni-cyclic network* is a network with
 53 exactly one simple cycle. A *cactus network* is a network in which no two simple cycles share an edge or, in
 54 other words, a network where each edge is contained in at most one simple cycle.

1.2 Related Work

Our data structures implicitly represent (generalized) farthest-point network Voronoi diagrams, where the sites are the entire continuum of points along a network [4]. Usually, Voronoi diagrams are defined with respect to a finite set of sites representing points of interest in some metric space [1, 24]. The existing research on Voronoi diagrams on networks covers a wide range of queries including queries for the closest [12, 28], the farthest [7, 25], and the k -th nearest neighbors [8, 21, 31] among a finite set of sites. We refer the interested reader to, for instance, Okabe et al. [24], Okabe and Sugihara [26], Okabe and Suzuki [27], and Taniar and Rahayu [31] for guides through the vast literature on network Voronoi diagrams.

Network Voronoi diagrams relate to center problems from location analysis [12, 27]. Let $G = (V, E)$ be a network. A *center* [11] of G is a vertex v that minimizes the network distance to any other vertex, i.e., $\max_{u \in V} d(u, v) = \min_{x \in V} \max_{u \in V} d(u, x)$. An *absolute center* a generalizes a center in that it may be placed anywhere along edges of the network, i.e., $\max_{u \in V} d(u, a) = \min_{p \in G} \max_{u \in V} d(u, p)$. A *continuous absolute center* is a point c with minimal eccentricity, i.e., $\max_{p \in G} d(p, c) = \min_{q \in G} \max_{p \in G} d(p, q) = \min_{q \in G} \text{ecc}(q)$.

Our results draw three ingredients from the literature on center problems: First, the absolute center [14] plays a crucial role when querying for farthest points in a tree network, as we shall see in Section 2.1. Second, the decomposition of cactus networks into *blocks*, *branches*, and *hinges* [2, 5, 17, 23] helps us exploit their tree structure [10, 16, 19] in Section 3. Third, viewing the network from different *perspectives* [2, 13, 20, 23] allows us to process branches and blocks independently, as explained in Sections 2.3 and 3.

Conversely, our constructions solve some center problems *en passant*: For example, we obtain the absolute center of a uni-cyclic network [13] as a by-product when building our data structure for queries in uni-cyclic networks. Furthermore, while it was known how to locate a single continuous absolute center in a cactus network in linear time [2], we produce the entire set of these centers in $O(n)$ time improving the $O(m^2 \log n)$ bound for general networks with n vertices and m edges [15].

A comprehensive summary of the literature about center problems is beyond the scope of this work. We refer readers interested in this sub-field of location analysis to a wealth of surveys [18, 22, 29, 33, 34], as well as to treatments of more recent results in the books by Kincaid [19], Shi [30], and Tansel [32].

1.3 Structure and Results

We introduce data structures supporting eccentricity queries and farthest-points queries for trees, cycles, uni-cyclic networks, and cactus networks. As summarized in Table 1, these data structures improve our results for general networks [4] and achieve optimal query times, sizes, and construction times.

Type	Eccentricity Query	Farthest-Point Query	Size	Construction Time
Tree	$O(1)$	$O(k)$	$O(n)$	$O(n)$
Cycle	$O(1)$	$O(\log n)$	$O(n)$	$O(n)$
Uni-Cyclic	$O(\log n)$	$O(k + \log n)$	$O(n)$	$O(n)$
Cactus	$O(\log n)$	$O(k + \log n)$	$O(n)$	$O(n)$
General [4]	$O(\log n)$	$O(k + \log n)$	$O(m^2)$	$O(m^2 \log n)$

Table 1: The traits of our data structures for queries in different types of networks, where n is the number of vertices, m is the number of edges, and k is the number of reported farthest points.

85 The remainder of this work is organized as follows. In [Section 2](#), we introduce data structures for trees,
 86 cycles, and uni-cyclic networks. In [Section 3](#), we construct data structures supporting eccentricity queries
 87 and farthest-point queries on cactus networks. Our approach is to reduce a cactus network to smaller
 88 networks having a sufficiently simple structure such that the data structures and query algorithms of
 89 [Section 2](#) can be applied. In [Section 4](#), we discuss directions for future research on closing the gap between
 90 general networks and cactus networks in [Table 1](#).

91 2 Trees, Cycles, and Uni-Cyclic Networks

92 In this section, we introduce data structures for farthest-point queries in trees and cycles. We then combine
 93 these two structures to support queries in uni-cyclic networks, i.e., networks with exactly one simple cycle.
 94 These data structures have linear size and construction times while providing optimal query times, and
 95 they serve as building blocks for our data structure for cactus networks in [Section 3](#).

96 2.1 Trees

97 The layout of farthest points on a tree hinges on the position of the absolute center. This point subdivides
 98 a tree into sub-trees where all points in a given sub-tree have their farthest points in other sub-trees.
 99 Conversely, each sub-tree has a set of leaves that are farthest from the absolute center and these leaves will
 100 be farthest from points in other sub-trees.

101 On tree networks, every farthest point is a leaf and the point c whose farthest leaves are closest is an
 102 *absolute center* [[11](#), [20](#)]. In other words, an absolute center on a tree T is a point with minimal eccentricity,
 103 i.e., $\text{ecc}(c) = \min_{g \in T} \text{ecc}(g)$. We say that two leaves l and l' of T are *most distant* when they realize the
 104 maximum distance between two points on T , i.e., $d(l, l') = \max_{a, b \in T} d(a, b)$.

105 **Theorem 1** (Handler [[14](#)]). *Every tree has exactly one absolute center midway along any path connecting two*
 106 *most distant leaves and we can locate the absolute center of a tree with n vertices in $O(n)$ time.*

107 Handler [[14](#)] determines the absolute center of a tree with two rounds of breadth-first-search: The first
 108 breadth-first-search starts from an arbitrary leaf l . With this search, we determine a farthest leaf \bar{l} from l .
 109 We then start the second breadth-first-search from \bar{l} to determine a farthest leaf \hat{l} from \bar{l} . Handler [[14](#)] shows
 110 that \bar{l} and \hat{l} are most distant leaves and that the absolute center is located midway on the path from \bar{l} to \hat{l} .

111 We split a tree T at its absolute center c into sub-trees as follows. When c lies on an edge uv with
 112 $u \neq c \neq v$, we split T into two sub-trees: the sub-tree T_1 containing the sub-edge uc , and the sub-tree T_2
 113 containing the sub-edge cv . When c lies on a vertex with neighbors v_1, v_2, \dots, v_r , we split T into r sub-trees
 114 T_1, T_2, \dots, T_r , where sub-tree T_i contains the sub-edge cv_i . [Figures 2](#) and [3](#) exemplify splitting trees at their
 115 absolute center and illustrate the following lemma that relates absolute centers with farthest points.

116 **Lemma 2.** *Let T be a tree with absolute center c , let T_i be one of the sub-trees obtained by splitting T at c , and*
 117 *let p be a point on T_i with $p \neq c$. The farthest distance from p in T is $\text{ecc}(p) = d(p, c) + \text{ecc}(c)$ and the farthest*
 118 *points from p are precisely the farthest leaves from c outside of T_i , i.e., for every leaf l we have*

$$\text{ecc}(p) = d(p, l) \iff l \notin T_i \wedge \text{ecc}(c) = d(c, l) .$$




Figure 2: A tree network T where all edges have unit weight unless indicated otherwise. The absolute center c splits T into two sub-trees T_1 (orange) and T_2 (blue). The farthest distance from c is $\text{ecc}(c) = 11.5$ and c has three farthest leaves: \bar{c}_1 and \bar{c}_2 in T_1 ; and \bar{c}_3 in T_2 . According to [Lemma 2](#), \bar{c}_3 is farthest from every point on sub-tree T_1 , whereas \bar{c}_1 and \bar{c}_2 are farthest from every point on sub-tree T_2 .




Figure 3: A tree T where all edges have unit weight. The absolute center c is located at a vertex and splits T into four sub-trees T_1 (blue), T_2 (orange), T_3 (green), and T_4 (red). The farthest distance from c is $\text{ecc}(c) = 6$ and c has seven farthest leaves: \bar{c}_1, \bar{c}_2 , and \bar{c}_3 in T_1 ; \bar{c}_4 in T_2 ; \bar{c}_5, \bar{c}_6 , and \bar{c}_7 in T_3 ; and none in T_4 . Let $L_1 = \{\bar{c}_1, \bar{c}_2, \bar{c}_3\}$, $L_2 = \{\bar{c}_4\}$, $L_3 = \{\bar{c}_5, \bar{c}_6, \bar{c}_7\}$, and $L_4 = \emptyset$. According to [Lemma 2](#), the leaves in $L_2 \cup L_3 \cup L_4$ are farthest from all points on T_1 , the leaves in $L_1 \cup L_3 \cup L_4$ are farthest from all points on T_2 , and the leaves in $L_1 \cup L_2 \cup L_4$ are farthest points from all points on T_3 . All points on the red sub-tree T_4 share their farthest points with the absolute center c , since $L_4 = \emptyset$.

119 *Proof of Lemma 2.* We show that every path from p to a farthest leaf \bar{p} from p passes through c .

120 Let l and \bar{l} be two most distant leaves of T . According to [Theorem 1](#), c is located midway along the
 121 path from l to \bar{l} with $\text{ecc}(c) = d(c, l) = d(c, \bar{l})$. When splitting T at c , the leaves l and \bar{l} end up in different
 122 sub-trees. Assume, without loss of generality, that l and p lie in different sub-trees, which implies that c lies
 on a path from p to l , i.e., $d(p, l) = d(p, c) + d(c, l)$.




Figure 4: The impossible layout of the paths in a tree T where a point p on T has a farthest point \bar{p} such that the shortest path from p to \bar{p} avoids the absolute center c of T . According to [Theorem 1](#), the absolute center c is located midway on a path between two most distant leaves l and \bar{l} . The path from p to one of these leaves— l in this case—passes through c .

123

124 Assume, for the sake of a contradiction, that p has a farthest leaf \bar{p} and the path in T from p to \bar{p} avoids c ,
 125 i.e., \bar{p} is in the same sub-tree as p . Let q be the meeting point of the path from c to \bar{p} with the path from p to
 126 \bar{p} . [Figure 4](#) shows this (impossible) constellation in which we have

$$d(p, q) + d(q, \bar{p}) = d(p, \bar{p}) = \text{ecc}(p) \geq d(p, l) = d(p, q) + d(q, l) ,$$

127 which implies $d(q, \bar{p}) \geq d(q, l)$, i.e., l is no farther away from q than \bar{p} . Since q lies on the path from p to \bar{p}
 128 and this path avoids c , we have $q \neq c$ and, thus, $d(c, q) > 0$. We arrive at a contradiction via

$$d(c, \bar{p}) = d(c, q) + d(q, \bar{p}) > d(q, \bar{p}) \geq d(q, l) = d(q, c) + d(c, l) > d(c, l) ,$$

129 which implies that \bar{p} is farther away from c than the farthest leaf l , i.e., $d(c, \bar{p}) > d(c, l) = \text{ecc}(c) \geq d(c, \bar{p})$.
 130 Therefore, every path from p to any farthest leaf must contain the absolute center c of T .

131 The farthest distance we can travel along a path from p through c is $d(p, c) + d(c, l) = d(p, c) + \text{ecc}(c)$.
 132 Conversely, every farthest leaf from c outside of T_i achieves this distance and, thus, is farthest from p . \square

133 We perform eccentricity queries on tree networks as follows. Let T be a tree with absolute center c .
 134 Consider an eccentricity query from a point p on edge uv where u is closer to c than v , i.e., $d(c, u) < d(c, v)$.
 135 The shortest path from p to c leads through u , i.e., $d(p, c) = w_{pu} + d(u, c)$, and we have $\text{ecc}(p) = d(p, c) +$
 136 $\text{ecc}(c) = w_{pu} + d(u, c) + \text{ecc}(c)$. Thus, we can determine the eccentricity of p in constant time, provided that
 137 we know the eccentricity of c and the network distance from c to every vertex of T .

138 We perform farthest-point queries on T as follows. Let T_1, T_2, \dots, T_r be the sub-trees obtained by
 139 splitting T at c . For $i = 1, 2, \dots, r$, let L_i denote the set of farthest leaves from c in sub-tree T_i , i.e.,
 140 $L_i := \{l \in T_i \mid \text{ecc}(c) = d(c, l)\}$. For a farthest point query from a point p on sub-tree T_i with $p \neq c$, we
 141 report all leaves in each L_j with $j \neq i$ as farthest points of p . For a farthest-point query from the absolute
 142 center c , we report all leaves in L_i for all $i = 1, 2, \dots, r$. For a query point with k farthest points, this takes
 143 $O(k)$ time, provided that we know the sets of farthest leaves L_1, L_2, \dots, L_r and provided that we can identify
 144 the sub-tree among T_1, T_2, \dots, T_r containing the query point in constant time.

145 The description of eccentricity queries and farthest-point queries on trees suggests which auxiliary data
 146 should be pre-computed. For a given tree T with n vertices, we first locate the absolute center c of T in
 147 $O(n)$ time using Handler's Algorithm [14]. Using a breadth-first-search from c , we perform three tasks: we
 148 compute the distances from c to every vertex of T , we label each edge with an index indicating its sub-tree,
 149 and we determine the farthest leaves in each sub-tree, which yields the sets L_1, L_2, \dots, L_r . Altogether, we
 150 spend $O(n)$ time to obtain our data structure, which is summarized in the following theorem.

151 **Theorem 3.** *Let T be a tree network with n vertices. There is a data structure with $O(n)$ construction time
 152 supporting eccentricity queries on T in constant time and farthest-point queries on T in $O(k)$ time, where k is
 153 the number of reported farthest points.*

154 By storing the lengths of the lists L_1, L_2, \dots, L_r , the data structure from **Theorem 3** also allows us to count
 155 the number of farthest-points from any query point in a tree network in constant time.

156 2.2 Cycles

157 Let C be a cycle network and let w_C be the sum of all edge weights of C . Each point p on C has exactly one
 158 farthest point \bar{p} located on the opposite side of C with $\text{ecc}(p) = d(p, \bar{p}) = w_C/2$. Supporting eccentricity
 159 queries on C amounts to calculating and storing the value $w_C/2$.

160 To support farthest-point queries, we subdivide C at each farthest point of a vertex and store a pointer
 161 from each vertex to its farthest point and vice versa. To compute this subdivision, we first locate the farthest
 162 point \bar{v} for some initial vertex v by walking a distance of $w_C/2$ from v along C . As illustrated in **Figure 5**,
 163 we then sweep a point p from position $p = v$ to position $p = \bar{v}$ along C while maintaining the farthest point
 164 \bar{p} . During this sweep we subdivide C at p whenever \bar{p} hits a vertex and at \bar{p} whenever p hits a vertex. The
 entire sweep takes linear time, thus, the resulting data structure occupies linear space.




Figure 5: A sweep along cycle C starting from $p = v_1$ and the resulting subdivision of C . For instance, any point on sub-edge $v_5\bar{v}_2$ has its farthest point on sub-edge \bar{v}_5v_2 .

165 Using the subdivided network, we answer farthest-point queries as follows. For a query point p on edge
 166 uv of C , we first locate the sub-edge ab containing p using binary search. This takes $O(\log n)$ time for a cycle
 167 with n vertices, since we subdivide uv at most n times. Let \bar{a} and \bar{b} be farthest from a and b , respectively.
 168 The farthest point \bar{p} from p is located on sub-edge $\bar{a}\bar{b}$ at distance w_{ap} from \bar{a} .
 169

170 **Lemma 4.** *Let C be a cycle network with n vertices. There is a data structure with construction time $O(n)$
 171 supporting eccentricity queries on C in constant time and farthest-point queries on C in $O(\log n)$ time.*

172 **2.3 Uni-Cyclic Networks**

173 A network with exactly one simple cycle is called *uni-cyclic* [13]. As illustrated in Figure 6, every uni-cyclic
 174 network U consists of a cycle C with trees T_1, T_2, \dots, T_l attached to C at vertices v_1, v_2, \dots, v_l , respectively.
 We refer to the trees T_1, T_2, \dots, T_l as the *branches* of U and to the vertex v_i as the *hinge* of branch T_i .


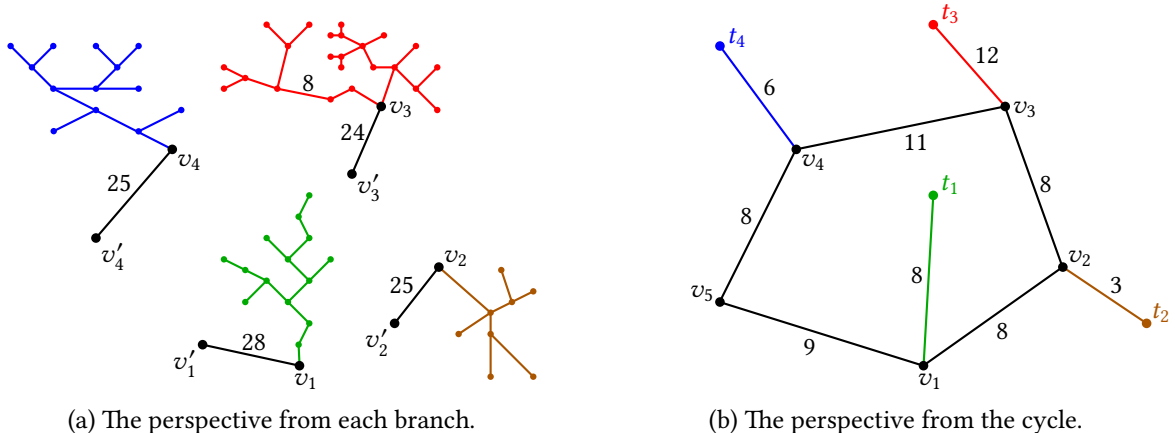


Figure 6: A uni-cyclic network with four branches. Edges have weight one whenever no weight is indicated.

175 Our data structure for uni-cyclic networks consists of two adjoined data structures: one for queries
 176 from the branches and one for queries from the cycle. These data structures are based on the following
 177 contractions. For each branch T with hinge v in a uni-cyclic network U , we represent all paths in U starting
 178 at hinge v and leading out of T by an edge vv' (to a dummy vertex v') of weight $w_{vv'} = ecc_{U \setminus T}(v)$. We call
 179 the tree consisting of T and the edge vv' the *perspective* of T from U denoted by $per_U(T)$. For the cycle C ,
 180 we represent the paths starting from hinge v leading into a branch T by an edge vt (to a dummy vertex t)
 181 of weight $w_{vt} = ecc_T(v)$. We call the resulting network the *perspective* of U from C denoted by $per_U(C)$.
 182 Figure 7 summarizes the different perspectives for the network from Figure 6.



(a) The perspective from each branch. (b) The perspective from the cycle.
 Figure 7: The network from Figure 6 viewed from the branches (a) and from the cycle (b).

184 The data structure for queries from the branches T_1, T_2, \dots, T_l consists of the tree data structures for
 185 $\text{per}(T_1), \text{per}(T_2), \dots, \text{per}(T_l)$ from Section 2.1. Consider a point p on branch T_i . An eccentricity query from p
 186 in $\text{per}(T_i)$ yields the eccentricity of p in U . A farthest-point query from p in $\text{per}(T_i)$ reports the farthest
 187 points from p that are located in T_i and this query reports t_i whenever p has farthest points outside of T_i .
 188 As we shall explain later, we obtain the farthest points from p outside of T_i using a query from the hinge of
 189 T_i in our data structure for the cycle perspective $\text{per}(C)$.

190 The data structure for queries from the cycle C consists of two components: the first component reports
 191 the farthest points from C on C itself, using the data structure for cycles from Section 2.2 on C ignoring
 192 the remainder of U . The second component reports which branches among T_1, T_2, \dots, T_l , if any, contain
 193 farthest points by supporting queries for the farthest vertices among t_1, t_2, \dots, t_l from any query point on
 194 C in the cycle perspective $\text{per}(C)$. We call this type of query a *farthest-branch query*.

195 2.3.1 Farthest-Branch Queries

196 We consider the perspective $\text{per}_U(C)$ for the cycle C of a uni-cyclic network. The vertices t_1, t_2, \dots, t_l
 197 represent the compressed branches of U in $\text{per}_U(C)$ and are connected to the hinges v_1, v_2, \dots, v_l , respectively.
 198 Furthermore, let \bar{v}_i denote the farthest point on C from hinge v_i , i.e., $d_C(v_i, \bar{v}_i) = \text{ecc}_C(v_i)$.

199 We call a vertex t_i *relevant* if there exists a point p on C that has t_i as a farthest vertex among t_1, t_2, \dots, t_l ,
 200 and we call t_i *irrelevant* otherwise. Knowing which of t_1, t_2, \dots, t_l are relevant will enable us to perform
 201 farthest-branch queries, i.e., report all branches containing farthest points from a query point on C .

202 **Lemma 5.** *Vertex t_i is relevant if and only if t_i is farthest from \bar{v}_i among t_1, t_2, \dots, t_l .*

203 *Proof.* We show both directions via indirection.

When t_i is irrelevant no point on C —including \bar{v}_i —has t_i as a farthest vertex among t_1, t_2, \dots, t_l .




Figure 8: The constellation of t_i and t_j where v_j lies on the clockwise path from v_i to \bar{v}_i .

204 Conversely, let there be some vertex t_j that is farther away from \bar{v}_i than t_i , i.e., $d(\bar{v}_i, t_i) < d(\bar{v}_i, t_j)$. As
 illustrated in Figure 8, v_j lies either on the clockwise or counter-clockwise path from t_i to \bar{v}_i . Since either
 path is a shortest path, we have $d(t_i, \bar{v}_i) = d(t_i, v_j) + d(v_j, \bar{v}_i)$. Thus, t_j is farther from v_j than t_i , as

$$d(t_i, v_j) = d(t_i, \bar{v}_i) - d(\bar{v}_i, v_j) < d(t_j, \bar{v}_i) - d(\bar{v}_i, v_j) = d(t_j, v_j) .$$

205 Now t_i is irrelevant, because t_j is farther away from any point p on C than t_i , since

$$d(p, t_i) \leq d(p, v_j) + d(v_j, t_i) < d(p, v_j) + d(v_j, t_j) = d(p, t_j) . \quad \square$$

206 According to [Lemma 5](#), a vertex t_i is irrelevant when there is some other vertex t_j with $d(t_i, \bar{v}_i) < d(t_j, \bar{v}_i)$.
 207 In this case, we say that t_i is *dominated* by t_j and write $t_i < t_j$. [Algorithm 1](#) below computes the relevant
 208 vertices using dominance. We begin with a circular list containing the vertices t_1, t_2, \dots, t_l in the order as
 209 the hinges v_1, v_2, \dots, v_l appear along the cycle C . We traverse the list in counterclockwise order and delete
 210 vertices whenever they are dominated by their successor (SUCC) or their predecessor (PRED). We mark a
 211 vertex t as processed if we can neither remove t nor any of its neighbors based on this criteria.

Algorithm 1: Determining the relevant vertices

input : The vertices t_1, t_2, \dots, t_l stored in a circular list.
output: The relevant vertices among t_1, t_2, \dots, t_l .

```

1 Mark each  $t_1, t_2, \dots, t_l$  as unprocessed;
2  $t \leftarrow t_1$ ;
3 while  $t$  is unprocessed do
4   if  $t < \text{PRED}(t)$  or  $t < \text{SUCC}(t)$  then
5      $t \leftarrow \text{SUCC}(t)$ ;
6     DELETE(PRED( $t$ ));
7   else if PRED( $t$ ) <  $t$  then DELETE(PRED( $t$ ));
8   else if SUCC( $t$ ) <  $t$  then DELETE(SUCC( $t$ ));
9   else /*  $t \not< \text{PRED}(t) \not< t \not< \text{SUCC}(t) \not< t$  */
10    Mark  $t$  as processed;
11     $t \leftarrow \text{SUCC}(t)$ ;
12  end
13 end

```

212 **Invariant 6.** *The following invariants hold whenever [Algorithm 1](#) marks a vertex as processed in Line 10.*

- 213 (i) *Every marked vertex t dominates none of its current neighbors, i.e., $\text{PRED}(t) \not< t$ and $\text{SUCC}(t) \not< t$.*
 214 (ii) *Every marked vertex t is dominated by none of its current neighbors, i.e., $t \not< \text{PRED}(t)$ and $t \not< \text{SUCC}(t)$.*

215 *Proof.* Assume, for the sake of a contradiction, that there is a marked vertex t_i with a neighbor t_j such that
 216 $t_i < t_j$ or $t_j < t_i$ when Line 10 is executed. Assume, without loss of generality, that t_j has been marked after
 217 t_i or that t_j has never been marked so far. When we marked t_i as processed, t_i and t_j were not neighbors.
 218 When t_j became a neighbor of t_i , the previous neighbor t_k of t_i was either deleted in Line 6 with t_j assuming
 219 the role of variable t in Line 5 or t_k was deleted in Lines 8 with $t = t_j$. In both cases variable t would be at
 220 t_j with t_i as direct neighbor and we would have deleted t_i in Line 7 or in Line 8 before marking any other
 221 vertex as processed. Therefore, we have never marked t_i , which contradicts our assumption. \square

222 As a consequence of [Invariant 6](#), no vertex dominates any neighbor when [Algorithm 1](#) terminates.


223 We are now ready to prove the correctness of [Algorithm 1](#). We pre-compute the distances from v_1 to all
 224 other vertices along C while constructing $\text{per}(C)$. This allows us to compare distances along C in constant
 225 time and, thus, enables us to determine in constant time whether one vertex dominates another.

226 **Theorem 7.** *Let S be the perspective of a uni-cyclic network U from its cycle C , and let t_1, t_2, \dots, t_l be the
 227 vertices of S representing the l branches of U given in clockwise order. [Algorithm 1](#) computes all relevant
 228 vertices among t_1, t_2, \dots, t_l in $O(l)$ time, provided that checking for dominance takes constant time.*

229 *Proof.* In each iteration of the while-loop of [Algorithm 1](#), we either delete some vertex or we mark the
 230 vertex stored in t as processed ensuring that it will never assume the role of t again. Therefore, [Algorithm 1](#)
 231 terminates in $O(l)$ steps, provided that checking for dominance is a constant time operation.

232 [Algorithm 1](#) never deletes a relevant vertex, since we only delete vertices that are dominated and, thus,
 233 irrelevant. Consider the circular list of those vertices that remain after [Algorithm 1](#) terminates. We assume,
 234 for the sake of a contradiction, that this list contains irrelevant vertices. Let t be a relevant vertex and let
 235 t_{cw} be the first irrelevant vertex in clockwise direction from t , and let t_{ccw} be the first irrelevant vertex
 236 in counterclockwise direction from t . Since no vertex in the final list dominates any of its neighbors by
 237 [Invariant 6](#), the final list contains at least four vertices with at least one vertex between t and t_{cw} and at
 238 least one vertex between t and t_{ccw} . The argumentation below remains valid when t_{cw} and t_{ccw} coincide.

239 Let \bar{v} , \bar{v}_{cw} and \bar{v}_{ccw} be the farthest point on C from t , t_{cw} , and t_{ccw} , respectively. We distinguish the two
 240 cases illustrated in [Figure 9](#), based on the relative positions of t , \bar{v} , and \bar{v}_{ccw} .



(a) The point \bar{v}_{ccw} lies clockwise between t to \bar{v} . (b) The point \bar{v}_{ccw} lies counterclockwise between t to \bar{v} .

Figure 9: The cyclic order of the branch representing vertices t , t_{ccw} , t_{cw} , t' , t'' , and their corresponding antipodal points in the two cases from the proof of [Theorem 7](#).

241 Consider the case when the point \bar{v}_{ccw} lies on the clockwise path from t to \bar{v} , as illustrated in [Figure 9a](#).
 242 Let $t' := \text{PRED}(t_{\text{ccw}})$ be the clockwise neighbor of t_{ccw} and let \bar{v}' be the farthest point from t_{ccw} on C . The
 243 vertices \bar{v}_{ccw} , \bar{v}' , and \bar{v} appear clockwise in this order along C , which implies that \bar{v}_{ccw} lies on a shortest
 244 path from t to \bar{v}' , i.e., $d(t, \bar{v}') = d(t, \bar{v}_{\text{ccw}}) + d(\bar{v}_{\text{ccw}}, \bar{v}')$. Furthermore, we have $t_{\text{ccw}} < t$ and $t' \not\prec t$, since t_{ccw}
 245 was the first dominated vertex in clockwise direction from t . Together, this yields

$$d(t_{\text{ccw}}, \bar{v}_{\text{ccw}}) \stackrel{t_{\text{ccw}} < t}{<} d(t, \bar{v}_{\text{ccw}}) = d(t, \bar{v}') - d(\bar{v}', \bar{v}_{\text{ccw}}) \stackrel{t' \not\prec t}{\leq} d(t', \bar{v}') - d(\bar{v}', \bar{v}_{\text{ccw}}) = d(t', \bar{v}_{\text{ccw}}) ,$$

246 which implies $d(t_{\text{ccw}}, \bar{v}_{\text{ccw}}) < d(t', \bar{v}_{\text{ccw}})$, i.e., $t_{\text{ccw}} < t'$. This contradicts [Invariant 6](#).

247 Consider the case when \bar{v}_{ccw} lies on the counterclockwise path from t to \bar{v} , as illustrated in [Figure 9b](#).
 248 This implies that the vertices t , t_{cw} , t_{ccw} , and \bar{v} appear clockwise in this order (with potentially $t_{\text{cw}} = t_{\text{ccw}}$).
 249 Therefore, \bar{v}_{cw} also lies on the counterclockwise path from t to \bar{v} . Similarly, to the above, we derive the
 250 contradiction $t_{\text{cw}} < t''$ where $t'' := \text{SUCC}(t_{\text{cw}})$ is the counterclockwise neighbor of t_{cw} . Therefore, t_{cw} (and
 251 thus t_{ccw}) cannot exist and there are no irrelevant vertices in the circular list produced by [Algorithm 1](#). \square

252 We have a circular list of the relevant vertices among t_1, t_2, \dots, t_l . Using this list, we pre-process S to
 253 support farthest-branch queries as follows. We pick any relevant vertex t_i and traverse C counterclockwise
 254 starting from \bar{v}_i keeping track of the farthest branch. Since t_i and its counterclockwise successor $t_j =$
 255 $\text{succ}(t_i)$ are relevant, \bar{v}_i has t_i as its farthest branch and \bar{v}_j has t_j as its farthest branch. At some point p_{ij}
 256 between \bar{v}_i and \bar{v}_j the farthest branch changes from t_i to t_j ; we subdivide C at p_{ij} and store t_i as farthest
 257 branch for the (sub)edges from \bar{v}_i to p_{ij} . We continue subdividing C in this fashion into at most l chains,
 258 one for each relevant vertex. Using a binary search on these chains, we answer farthest-branch queries in
 259 $O(\log l)$ time. Figure 10 illustrates an example of this subdivision for the network from Figure 6. As in all
 260 figures in this section, edges have unit weight unless indicated otherwise.

261 **Theorem 8.** Let U be a uni-cyclic network with n vertices and cycle C of length l . There is a data structure
 262 with construction time $O(n)$ supporting farthest-branch queries from any query point on C in $O(\log l)$ time.

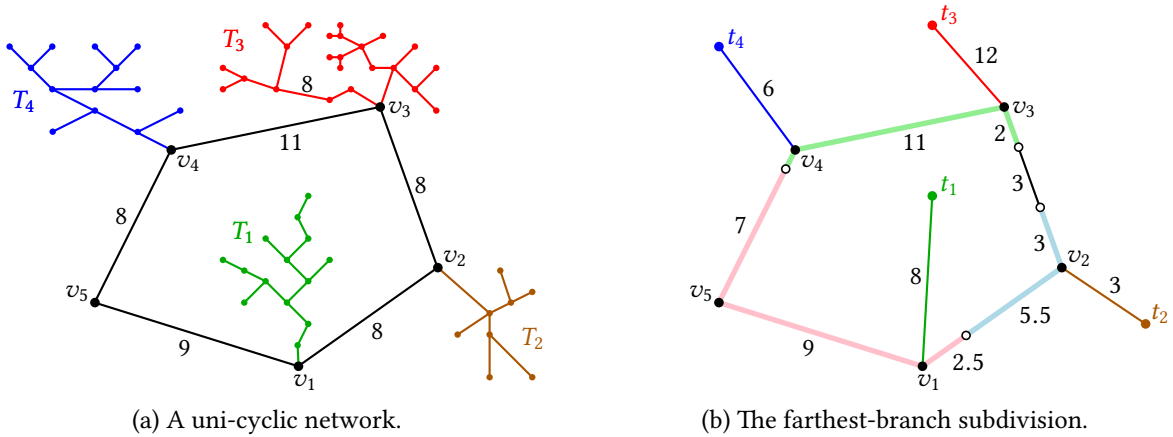


Figure 10: The farthest-branch subdivision (b) of the uni-cyclic network from (a). The sub-edges in (b) are shaded in a colour matching their farthest branch. On black sub-edges, no branch contains farthest points—instead, the farthest point lies on the opposite side of the cycle. In this example, no points on the cycle have farthest points in branch T_2 , i.e., t_2 is irrelevant.

263 2.3.2 Queries in Uni-Cyclic Networks

264 We perform an eccentricity query from a point q on U as follows. When q lies on some branch T , we use
 265 the tree data structures for the perspective $T' = \text{per}(T)$ from T , as $\text{ecc}_{T'}(q) = \text{ecc}_U(q)$ by construction of
 266 T' . When q lies on the cycle C , we first compute the distance $d_B := \max_{i=1}^l d(q, t_i)$ from q to the farthest
 267 vertex among t_1, t_2, \dots, t_l in the perspective $S = \text{per}(C)$ from C , using our data structure for farthest-branch
 268 queries in S . Then, we compute the farthest distance $\text{ecc}_C(q)$ from q on C , using the cycle data structure
 269 for C . The greater distance of the two is the eccentricity of q in S and, thus, the eccentricity of q in U , i.e.,
 270 $\text{ecc}_U(q) = \text{ecc}_S(q) = \max\{d_B, \text{ecc}_C(q)\}$. This way, we answer eccentricity queries from branches in constant
 271 time and eccentricity queries from the cycle in $O(\log l)$ time.

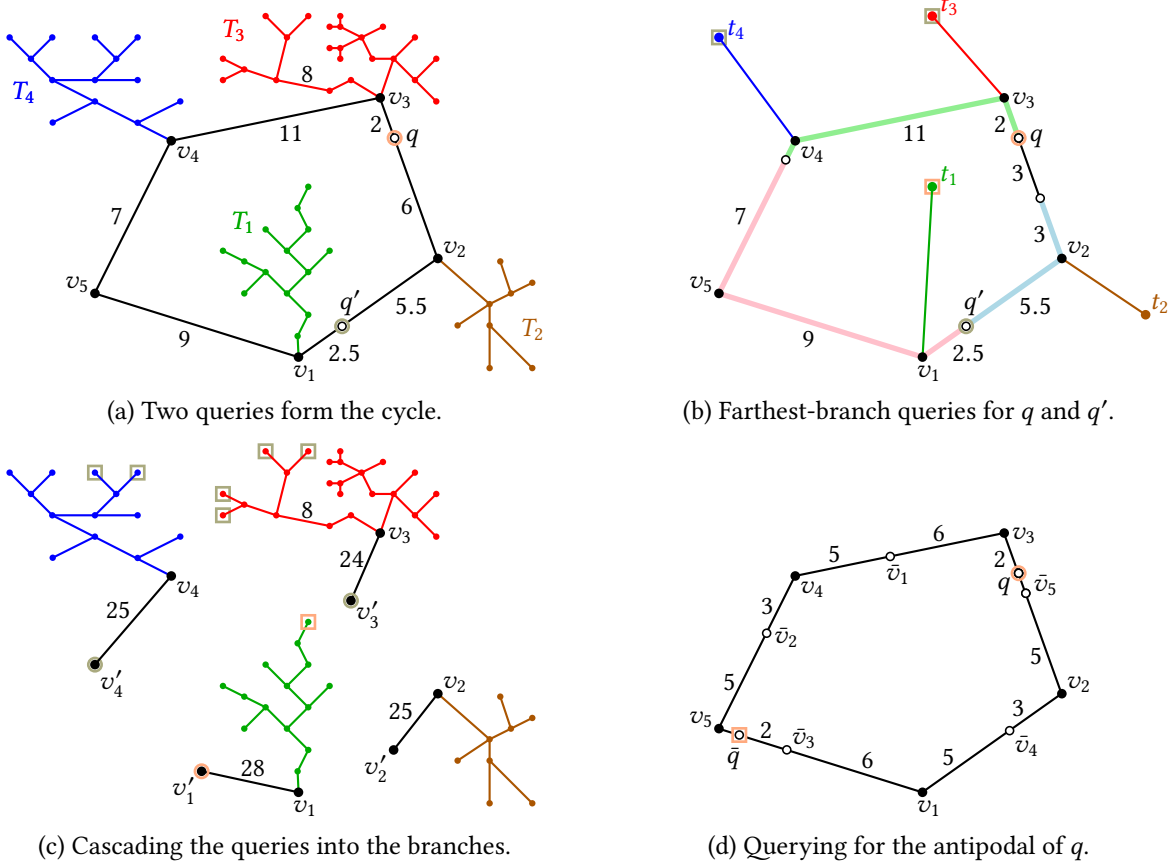


Figure 11: The farthest-point queries from the points q (salmon) and q' (ocker) on the cycle of a uni-cyclic network (a). We perform a farthest-branch query for both query points (b) and then cascade the queries into the perspectives from the branches (c) and into the cycle (d) as needed.

272 We perform a farthest-point query from a point q on the cycle C of U as illustrated in Figure 11. We
 273 first perform a farthest branch query from q in S and a farthest point query from q in C . We report the
 274 antipodal of q on C if it is farthest from q in U and then cascade the query in each reported branch. More
 275 precisely, if T was a branch containing farthest points from q then perform a farthest-point query from
 276 v' in the perspective $T' = \text{per}_U(T)$ of U from T , where v' is the vertex representing the exterior of T . The
 277 farthest leaves from v' in T' are also farthest from any point outside of T that has farthest points in T .

278 We perform a farthest-point query from a point q on a branch T of U with hinge v as follows. We begin
 279 with a farthest point query in the perspective $T' = \text{per}(T)$ of U from T . This reports all farthest points from
 280 q in T and, potentially, the point v' representing the exterior of T . When cascading the query from T' into
 281 the perspective $S = \text{per}(C)$ of U from the cycle C , we need to perform a farthest-point query in S from the
 282 vertex t representing T . However, we have no structure to support this query directly. Instead, we query
 283 from v in S , which leads to a *good case* and a *bad case*.

284 In the good case, shown in Figure 12, v has some farthest point in S other than t , which means that the
 285 farthest points from t are the farthest points from v (potentially excluding t itself). In the bad case, shown
 286 in Figure 13, t is the only farthest point from v in S . Fortunately, the bad case can only appear for exactly
 287 one branch of U , as t being the only farthest vertex from v implies that t is farthest from all other vertices of
 288 S as well. Therefore, we can deal with the bad case by computing the farthest points from v in the network
 289 $S - vt$ (during the preprocessing phase) and storing the result with v . We use this only when cascading a
 290 query from T into S , i.e., when we know that there are farthest points from the query point q outside of T .
 291 The following theorem summarizes our data structure for uni-cyclic networks.

292 **Theorem 9.** *Let U be a uni-cyclic network with n vertices and a cycle of length l . There is a data structure*
 293 *with size and construction time $O(n)$ supporting eccentricity queries from the branches of U in $O(1)$ time,*
 294 *eccentricity queries from the cycle in $O(\log l)$ time, farthest-point queries from branches in $O(k)$ time, and*
 295 *farthest-point queries from the cycle in $O(k + \log l)$ time, where k is the number of reported farthest points.*

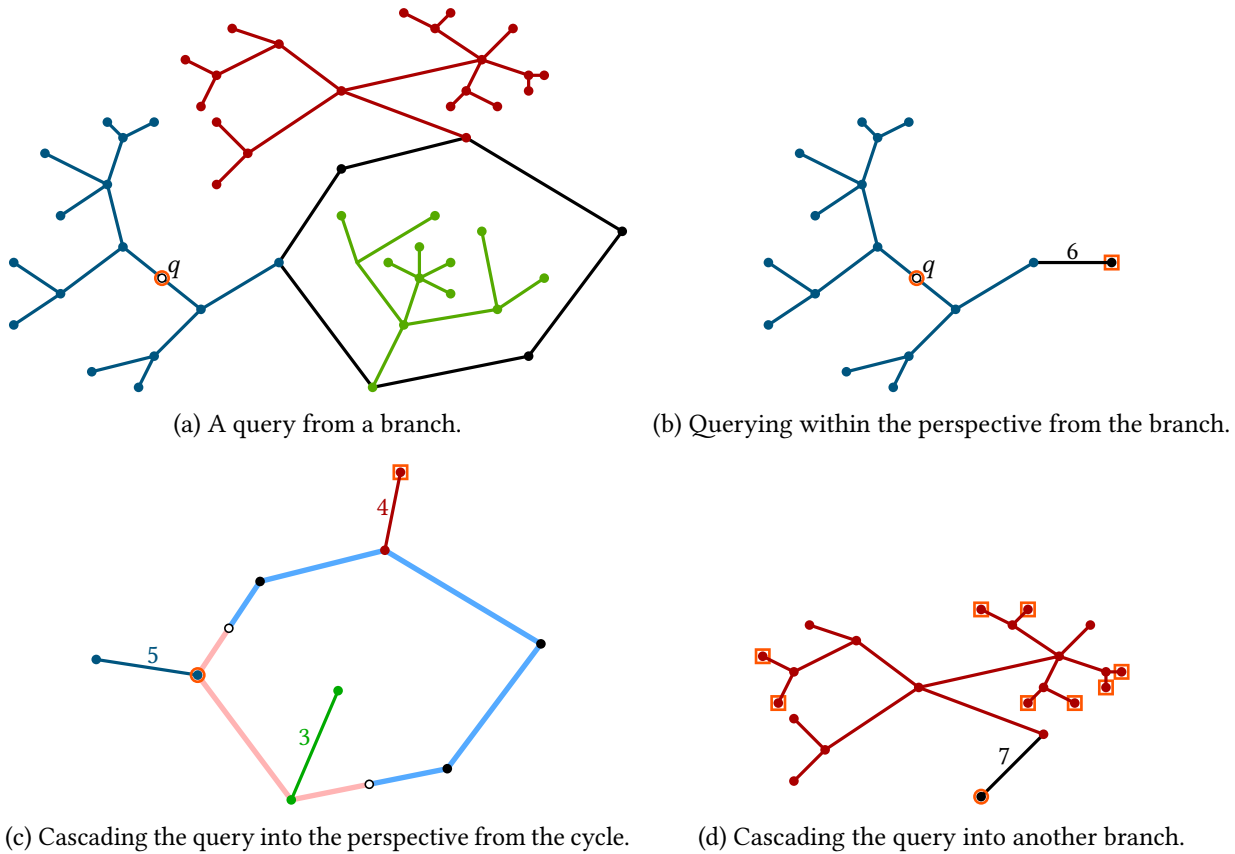


Figure 12: A farthest-point query from a branch of a unicyclic network (a). We perform a farthest-point query in the perspective from the branch containing the query point (b). We cascade the query into the perspective from the cycle (c) and then into another branch as needed (d).

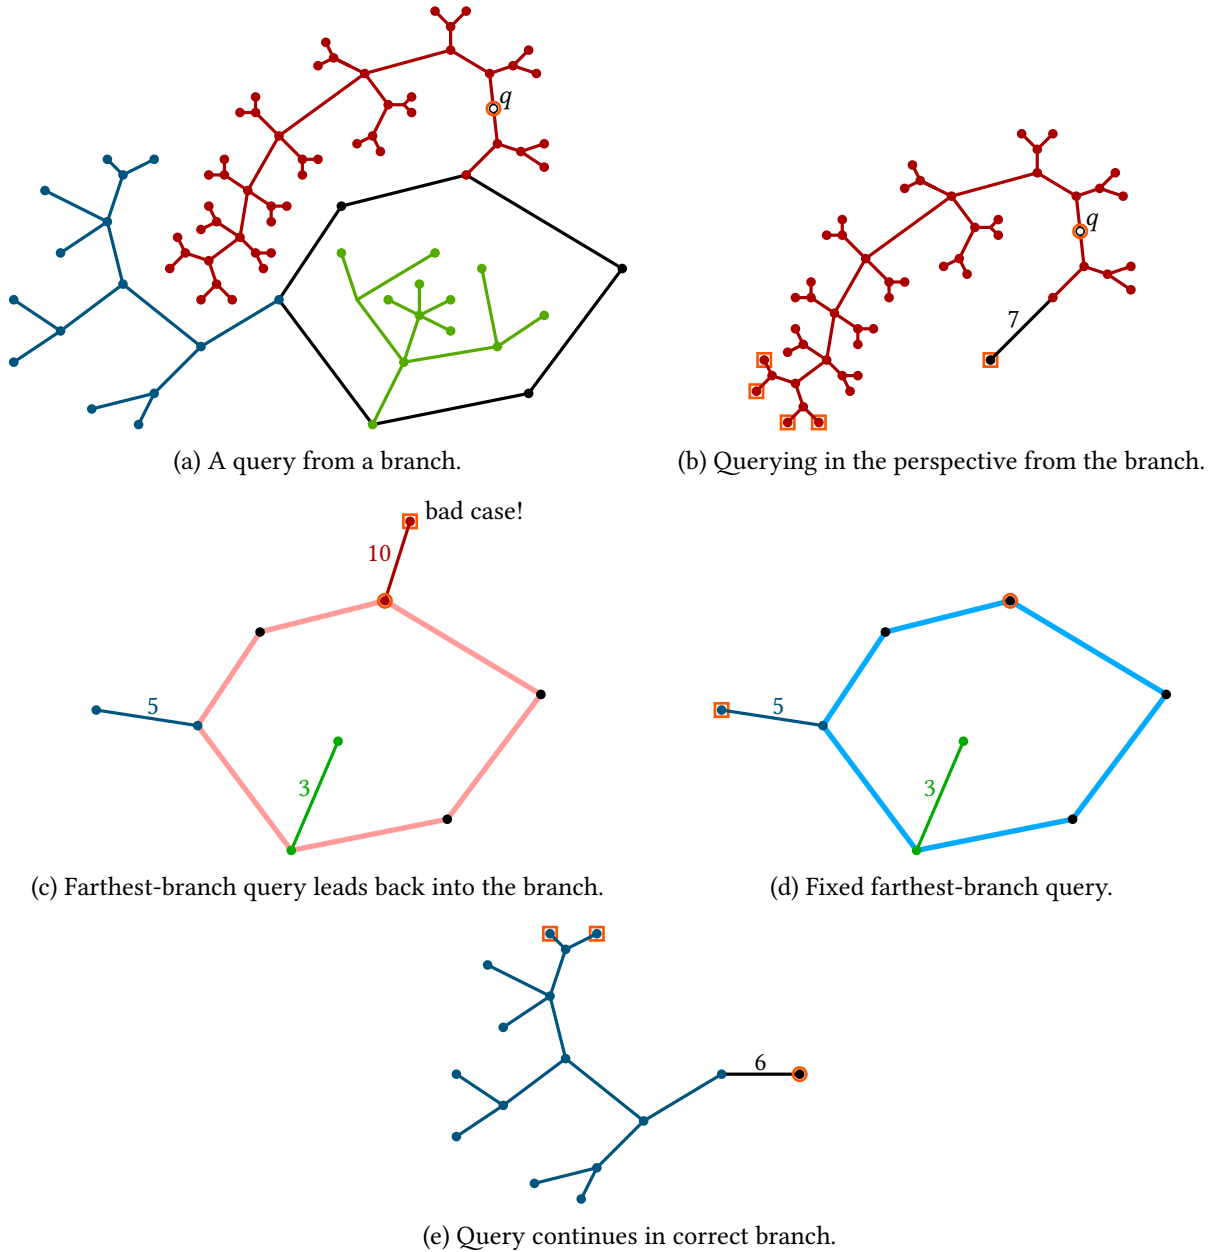


Figure 13: The bad case for a farthest-point query from a branch T (red) of a unicyclic network (a). We perform a farthest-point query in the perspective from the branch containing the query point (b). We cascade the query into the perspective from the cycle (c), however, this query sends us back into the red branch, since the red branch is farthest on the entire cycle. To fix this, we remove the edge representing the red branch (d) and repeat the query in the resulting network. This time, we obtain the correct branch where we cascade the query (e) to determine the farthest points in this (blue) branch. Observe that the fixed network (d) is irrelevant for queries from other branches which will use the normal perspective from the cycle (c).

296 **3 Cactus Networks**

297 In this section, we construct a data structure supporting eccentricity queries and farthest-point queries on
 298 cactus networks. Recall that a cactus network is a network in which no two simple cycles share an edge.

299 The following notions prove useful when describing cactus networks; examples are shown in Figure 14.
 300 A *cut-vertex* is a vertex whose removal increases the number of connected components, a *block* is a maximal
 301 connected sub-graph without cut-vertices and with at least three vertices, and a *branch* is a (maximal) tree
 302 that remains when removing the edges of all blocks and all resulting isolated vertices. We treat blocks and
 303 branches in a similar fashion, so we refer to a sub-network B as a *bag* when B is either a block or a branch.
 304 Decomposing a network with n vertices into its blocks and branches takes $O(n)$ time [17]. A *hinge* is a
 305 vertex contained in more than one bag. The *tree structure* of a cactus network G , denoted by $T(G)$, is the tree
 whose vertices are the hinges and bags of G where a hinge h is connected to a bag B when h lies in B [16].

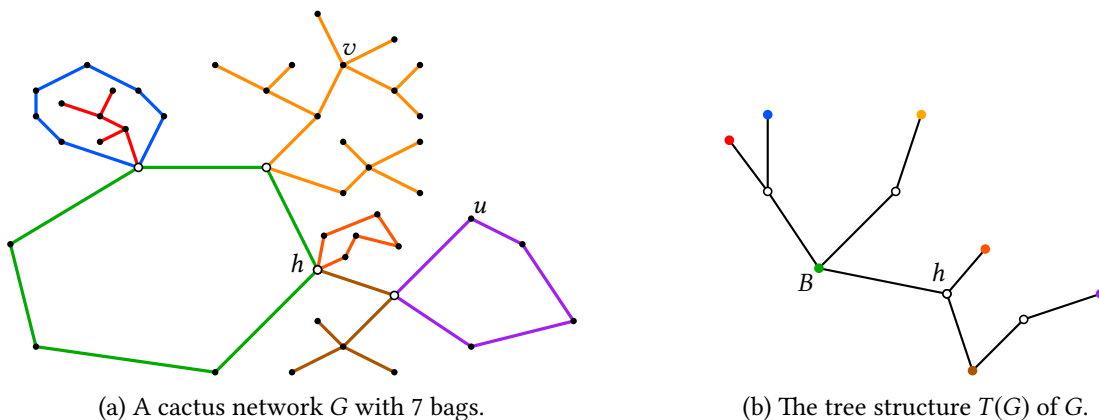


Figure 14: A cactus network (a) together with its tree structure (b). The blocks and branches are indicated in colours and the hinges connecting these bags are marked as empty discs. For example, the green edges form a block and the red edges form a branch. Vertex h is a hinge, because h is a cut-vertex contained in more than one bag; vertex v is a cut-vertex but not a hinge, since v is only contained in the yellow branch; and vertex u is not a cut-vertex and, thus, also not a hinge.

306 Let B be a bag containing hinge h . The component containing h after removing all edges of B is called
 307 the *bag-cut* of B at h , denoted by $\text{bcut}(B, h)$. The component containing h after removing all edges of
 308 $\text{bcut}(B, h)$ is called the *co-bag-cut* of B at h , denoted by $\text{co-bcut}(B, h)$. Figure 15 gives examples of bag-cuts
 309 and co-bag-cuts for the cactus network from Figure 14.

310 We view an (undirected) edge from B to h in the tree structure as two (directed) arcs $B \rightarrow h$ and $h \rightarrow B$.
 311 As illustrated in Figures 15c and 15d, we associate $\text{bcut}(B, h)$ with $B \rightarrow h$ and we associate $\text{co-bcut}(B, h)$
 312 with $h \rightarrow B$. We use this correspondence as orientation. For example, we store the eccentricity of h in
 313 $\text{bcut}(B, h)$ with the arc $B \rightarrow h$ and the eccentricity of h in $\text{co-bcut}(B, h)$ with the arc $h \rightarrow B$.
 314

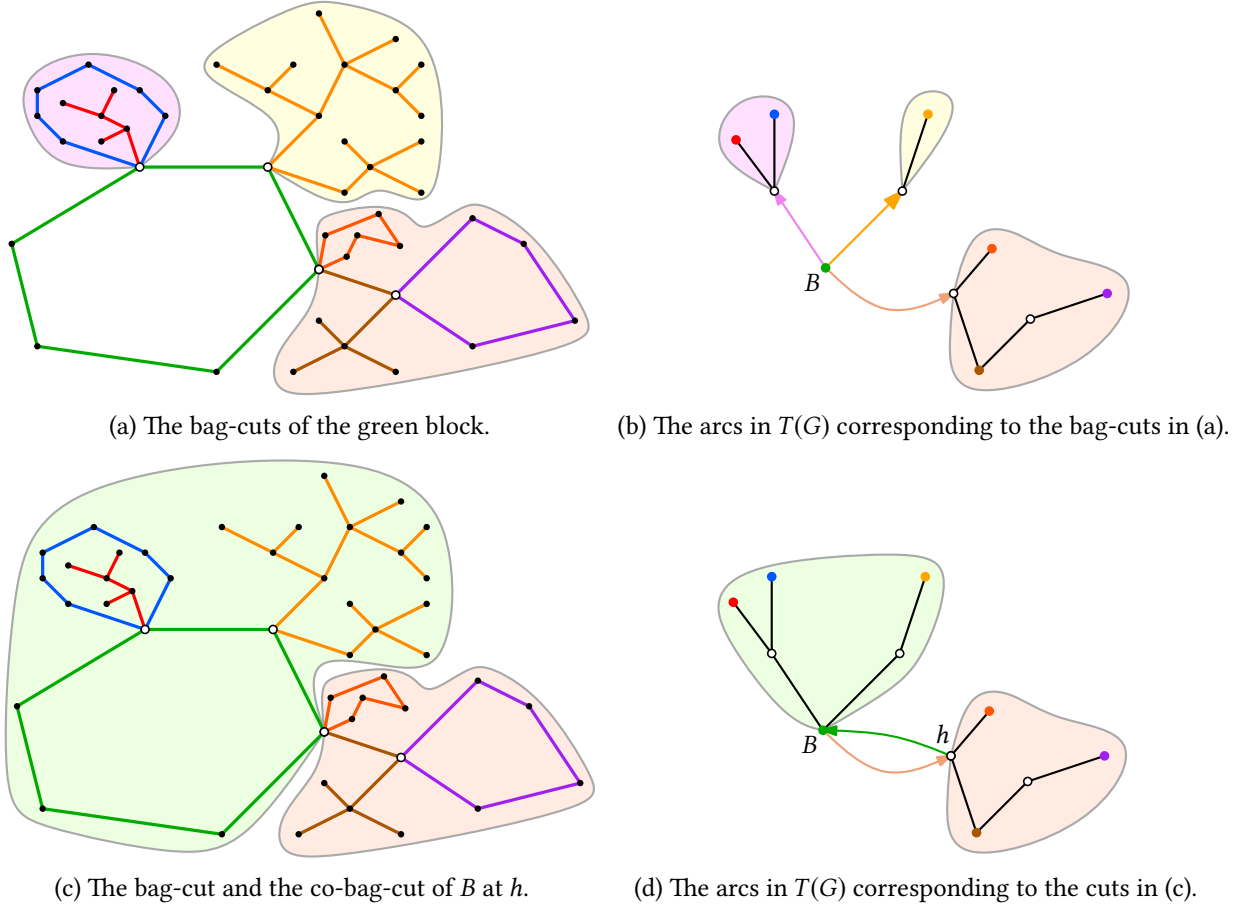


Figure 15: The bag-cuts of the green block (a) together with their corresponding arcs (of matching colour) in the tree structure (b). The bag-cut (salmon) and co-bag-cut (light green) corresponding to the two arcs of an edge from B to h in the tree structure (d) and in the network itself (c).

3.1 Eccentricity Queries

To support eccentricity queries on a bag B of a network G , we compress the bag-cuts of B like we compress the branches of uni-cyclic networks: for any hinge $h \in B$ we replace $\text{bcut}(B, h)$ with a vertex \hat{h} and an edge $h\hat{h}$ whose weight is the largest distance from h to any point in $\text{bcut}(B, h)$, i.e., $w_{h\hat{h}} = \text{ecc}_{\text{bcut}(B, h)}(h)$. We refer to the resulting network as the *perspective* of G from B , denoted by $\text{per}_G(B)$. The perspective of G from bag B preserves farthest distances of G , i.e., we have $\text{ecc}_{\text{per}_G(B)}(p) = \text{ecc}_G(p)$ for all p on B .

The perspective from a branch is a tree and the perspective from a block of a cactus network is a uni-cyclic network. Since, we already have efficient data structures for trees and uni-cyclic networks, the challenge lies in constructing these perspectives in linear time. As illustrated in Figure 16, we first construct the perspective $\text{per}(B^*)$ of an arbitrary bag B^* by computing the extended shortest path tree of each hinge h^* of B^* in $\text{bcut}(B^*, h^*)$. This takes linear time using a modified breadth-first search from each hinge h^* where we cut each simple cycle C at the antipodal \bar{v} of the vertex v where we first enter C . Using the data structure for queries in $\text{per}(B^*)$ and the extended shortest path trees from the hinges of B^* , we then construct the perspectives from the bags adjacent to B^* and continue in a breadth-first search fashion.




Figure 16: A schematic view of the construction of the perspective $\text{per}(B^*)$ from the initial bag B^* . For every hinge h^* of B^* , the bag-cut $\text{bcut}(B^*, h^*)$ is replaced with an edge $h^* \hat{h}^*$ weighted with the largest distance from h^* into $\text{bcut}(B^*, h^*)$. The extended shortest path trees from the hinges of B^* are outlined as arrows from the hinges of the network. An arrow pointing into the co-bag-cut of bag B at hinge h means that we obtain the extended shortest path tree from h into $\text{co-bcut}(B, h)$ as a by-product of the construction of the perspective from B^* .

Let B' be a bag neighboring B^* at hinge h^* , as shown in Figure 17. For all hinges h' of B' with $h' \neq h^*$, the extended shortest path tree of h' in $\text{bcut}(B', h')$ is a sub-tree of the extended shortest path tree of h^* in $\text{bcut}(B^*, h^*)$. So we know the weight of the edge $\hat{h}' h'$ in $\text{per}(B')$, where \hat{h}' represents $\text{bcut}(B', h')$. At hinge h^* , we have the extended shortest path trees of h^* in every co-bag-cut at h^* except for $\text{co-bcut}(B^*, h^*)$, the one leading back into B^* . The eccentricity of h^* in $\text{bcut}(B', h')$ is the largest distance from h^* into any co-bag-cut $\text{co-bcut}(B, h^*)$ for all bags B neighboring B' at h^* , i.e.,


$$\text{ecc}_{\text{bcut}(B', h^*)}(h^*) = \max\{\text{ecc}_{\text{co-bcut}(B, h^*)}(h^*) \mid B \text{ is a bag with } h^* \in B \neq B'\} . \quad (1)$$

329 We need $\text{ecc}_{\text{co-bcut}(B^*, h^*)}(h^*)$ to compute $\text{per}(B')$. The difference between the perspective $\text{per}_G(B^*)$ of G
 330 from B^* and the perspective $\text{per}_{\text{co-bcut}(B^*, h^*)}(B^*)$ of $\text{co-bcut}(B^*, h^*)$ from B^* is that the latter lacks the edge $\hat{h}^* h^*$
 331 where \hat{h}^* represents $\text{bcut}(B^*, h^*)$ in $\text{per}_G(B^*)$. We avoid constructing $\text{per}_{\text{co-bcut}(B^*, h^*)}(B^*)$ using the following
 332 observation. The farthest points from \hat{h}^* in $\text{per}_G(B^*, h^*)$ are also farthest from h^* in $\text{per}_{\text{co-bcut}(B^*, h^*)}(B^*) =$
 333 $\text{per}_G(B^*, h^*) - \hat{h}^* h^*$, but they are closer by $w_{h^* \hat{h}^*}$, i.e.,


$$\text{ecc}_{\text{co-bcut}(B^*, h^*)}(h^*) = \text{ecc}_{\text{per}(B^*) - \hat{h}^* h^*}(h^*) = \text{ecc}_{\text{per}(B^*)}(\hat{h}^*) - w_{h^* \hat{h}^*} = \text{ecc}_{\text{per}(B^*)}(\hat{h}^*) - \text{ecc}_{\text{bcut}(B^*, h^*)}(h^*) .$$

334 To see this, consider the extended shortest path tree from \hat{h}^* in $\text{per}_G(B^*)$. Removing the edge $\hat{h}^* h^*$ from
 335 this tree yields the extended shortest path tree from h^* in $\text{per}_G(B^*) - \hat{h}^* h^* = \text{co-bcut}(B^*, h^*)$. In this way,
 336 we obtain $\text{ecc}_{\text{co-bcut}(B^*, h^*)}(h^*)$ with a constant time query from \hat{h}^* in $\text{per}_G(B^*)$, as illustrated in Figure 17d.


337 Finally, we obtain the missing value $\text{ecc}_{\text{bcut}(B', h^*)}(h^*)$ by taking the maximum in (1). However, we need
 338 to avoid a dependence on the number of bags containing h^* , i.e., the degree of h^* in the tree structure.




(a) The bag-cuts of neighboring bag B' .



(b) The missing weight in the perspective from B' .



(c) The desired query in $\text{co-bcut}(B^*, h^*)$.



(d) The actual query in $\text{per}(B^*)$.

Figure 17: A schematic view of the construction of the perspective of a bag B' neighboring the initial bag B^* . From the construction of $\text{per}(B^*)$, we already have the distance information for the bag-cuts (blue) of B' at all hinges except for the one at h^* (green) and we have the distances for all co-bag-cuts of neighbors of B' (yellow) other than $\text{co-bcut}(B^*, h^*)$ (red). To obtain the missing value $\text{ecc}_{\text{co-bcut}(B^*, h^*)}(h^*)$ we would like to query from h^* in $\text{co-bcut}(B^*, h^*)$, but instead we perform a query from \hat{h}^* in $\text{per}(B^*)$ to avoid constructing a data structure for $\text{co-bcut}(B^*, h^*)$.

339 We augment the arcs of $T(G)$ with the following during the construction of B^* : When we compute the
340 extended shortest path trees from the hinges of B^* , we store $\text{ecc}_{\text{co-bcut}(B,h)}(h)$ with arc $h \rightarrow B$ in $T(G)$.
341 These values are, at first, unknown for arcs in $T(G)$ on paths towards B^* . With each hinge vertex we
342 store the largest and second largest known value among its adjacent arcs in $T(G)$ and two bags B_1 and
343 B_2 attaining these values. How does this help us compute the maximum in (1)? When we learn of the
344 missing value $\text{ecc}_{\text{co-bcut}(B^*,h^*)}(h^*)$ at h^* in $T(G)$, we have three cases depending on whether h^* has farthest
345 points in the direction of B^* or, if not, in the direction of B' . First, $\text{ecc}_{\text{co-bcut}(B^*,h^*)}(h^*)$ could be larger than
346 $\text{ecc}_{\text{co-bcut}(B_1,h^*)}(h^*)$ in which case $\text{ecc}_{\text{bcut}(B',h^*)}(h^*) = \text{ecc}_{\text{co-bcut}(B^*,h^*)}(h^*)$. Second, $\text{ecc}_{\text{co-bcut}(B^*,h^*)}(h^*)$ could be
347 strictly smaller than $\text{ecc}_{\text{co-bcut}(B_1,h^*)}(h^*)$ with $B_1 \neq B'$ in which case $\text{ecc}_{\text{bcut}(B',h^*)}(h^*) = \text{ecc}_{\text{co-bcut}(B_1,h^*)}(h^*)$.
348 Third, $\text{ecc}_{\text{co-bcut}(B^*,h^*)}(h^*)$ could be strictly smaller than $\text{ecc}_{\text{co-bcut}(B_1,h^*)}(h^*)$ with $B_1 = B'$ in which case
349 $\text{ecc}_{\text{bcut}(B',h^*)}(h^*) = \text{ecc}_{\text{co-bcut}(B_2,h^*)}(h^*)$. With the aforementioned bookkeeping, we can handle each of the
350 three cases in constant time and consequently construct $\text{per}(B')$ in time proportional to the size of B' .

351 With the above technique, we construct $\text{per}(B')$ re-using the preprocessing from the construction of
352 $\text{per}(B^*)$ and the augmented tree structure. In the same way, we construct the perspectives from all other
353 neighbors of B^* , then all perspectives from the neighbors of all neighbors of B^* and so forth for an overall
354 construction time of $O(n)$. We inherit the query times from the data structures of trees and uni-cyclic
355 networks; in summary this yields the following theorem for eccentricity queries in cactus networks.

356 **Theorem 10.** *Let G be a cactus network with n vertices. There is a data structure with $O(n)$ size and construction*
357 *time supporting eccentricity queries on G in $O(1)$ time from branches and in $O(\log l)$ time from blocks of size l .*

358 3.2 Farthest-Point Queries

359 Consider a farthest-point query from a point q in bag B , as illustrated in Figure 18. First, we perform a
360 farthest-point query from q in the farthest-point data structure of the perspective of G from B . This query
361 yields all farthest-points from q inside B and it returns the vertex \hat{h} representing the bag-cut $\text{bcut}(B, h)$
362 when q has farthest points in $\text{bcut}(B, h)$. Any path from q to a farthest point \bar{q} in $\text{bcut}(B, h)$ passes through h ,
363 hence \bar{q} is also farthest from h in $\text{bcut}(B, h)$. Recall that $\text{bcut}(B, h)$ consists of the co-bag-cuts $\text{co-bcut}(B', h)$
364 of the bags B' neighboring B at h . To find all farthest points from q in $\text{bcut}(B, h)$, we cascade the query into
365 the co-bag-cuts $\text{co-bcut}(B', h)$ of those neighbors B' of B that contain farthest points from h in $\text{bcut}(B, h)$,
366 i.e., $\text{ecc}_{\text{co-bcut}(B',h)}(h) = \text{ecc}_{\text{bcut}(B,h)}(h)$. For the decision into which co-bag-cuts to cascade, we consult the
367 eccentricity values stored with the arcs of the tree structure.

368 Now assume we cascade a query from q into the co-bag-cut $\text{co-bcut}(B', h)$, as in Figures 18c and 18d. To
369 query for the farthest points from h in $\text{co-bcut}(B', h)$, we perform a farthest-point query from \hat{h}' in $\text{per}(B')$,
370 where \hat{h}' is the vertex representing $\text{bcut}(B', h)$ in $\text{per}(B')$. The farthest points from \hat{h}' in $\text{per}(B')$ are also
371 farthest from h in $\text{per}(B') - \hat{h}'h$, which corresponds to $\text{co-bcut}(B', h)$. Since \hat{h}' is a pendant vertex in $\text{per}(B')$,
372 this query takes time proportional to the number of reported farthest points in $\text{per}(B')$. The farthest points
373 from \hat{h}' in $\text{per}(B')$ that lie in B' are farthest points from the original query point q , and those farthest points
374 from \hat{h}' in $\text{per}(B')$ that represent bag-cuts adjacent to B' indicate where we have to continue our search for
375 other farthest points from q . In this fashion, we propagate the query from q to the bags of G along paths to
376 farthest points from q . However, we might visit $\Omega(n)$ bags before we reach one that contains a farthest
377 point from q . We improve the query time by introducing shortcuts in the tree structure to bypass long
378 chains of bags without farthest points.




Figure 18: Answering a farthest-point query from q using the perspectives. Query points are indicated with orange circles, reported farthest points with orange squares. To process the query, we first query for q in the perspective from the bag containing the query point (b). Then we cascade the query into adjacent bags that lead to further farthest points (c and d). Here, we show only two cascaded queries, to report all farthest points, we need five more cascaded queries.

379 Consider a co-bag-cut $\text{co-bcut}(B, h)$ of bag B at hinge h . This co-bag-cut corresponds to the arc from h to
380 B in the tree structure and all bags in $\text{co-bcut}(B, h)$ are in the sub-tree $T_{h \rightarrow B}$ reachable from h through B . To
381 give a visual idea of our shortcuts, imagine we do the following: First, we color a bag B' in $T_{h \rightarrow B}$ red when
382 B' contains farthest points from h in $\text{co-bcut}(B, h)$. Second, we color an uncolored bag B' in $T_{h \rightarrow B}$ orange
383 when two paths from h to red bags split at B' . Finally, we color the remaining bags black. We seek to bypass
384 black bags, since these are irrelevant for our query. For now, we only consider arcs of $T_{h \rightarrow B}$ leading away
385 from h . The shortcut for arc $h' \rightarrow B'$ in $T_{h \rightarrow B}$ leads to the first arc $h'' \rightarrow B''$ where B'' is the closest orange
386 or red descendant of h' . The shortcuts reachable from h form a directed tree where all vertices representing
387 bags are either yellow or red; following all shortcuts in this tree takes time $O(r)$ where r is the number of
388 red bag vertices in $T_{h \rightarrow B}$. In other words, using these shortcuts, the number of bags visited when reporting
389 farthest points from h in $\text{co-bcut}(B, h)$ is linear in the number of bags containing these farthest points. An
390 example of a farthest-point query with and without using shortcuts is shown in Figure 19.

391 How can we determine these shortcuts efficiently? Recall that the propagation scheme starts at some bag
392 B^* where we first compute the extended shortest path tree of each hinge h in B^* in $\text{bcut}(B^*, h^*)$. We first
393 consider only the arcs of $T(G)$ leading away from B^* . The arcs leading to leaf bags of $T(G)$ need no shortcuts.
394 Consider an arc $h \rightarrow B$ and assume that all other arcs in the sub-tree $T_{h \rightarrow B}$ already have their shortcuts, as
395 illustrated in Figure 20 on page 23. We distinguish three cases using a query in $\text{per}(B)$: First, there could be
396 farthest points from \hat{h} in B (red case). Second, there could be farthest points from \hat{h} in two bag cuts of B

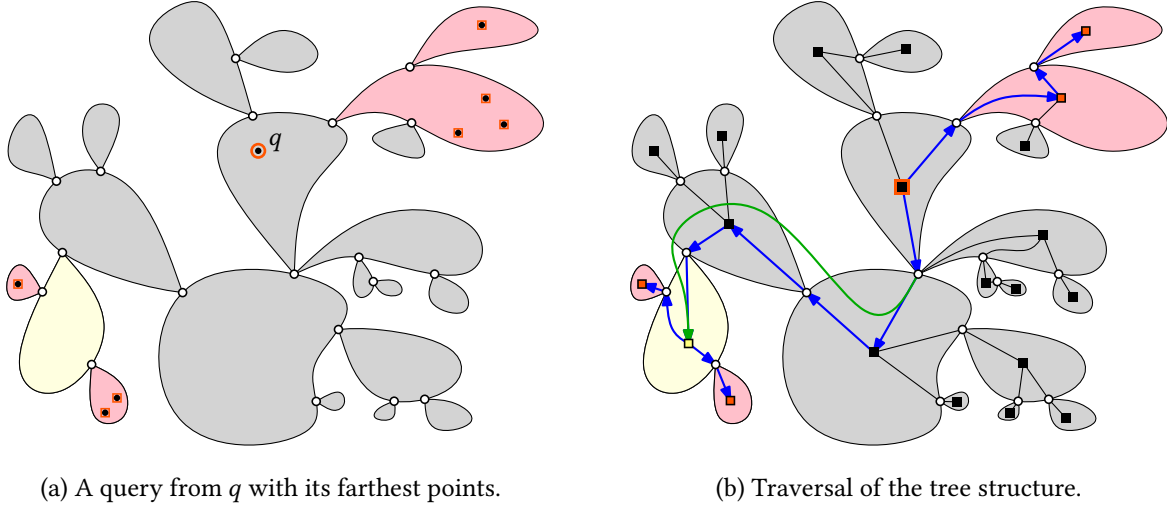


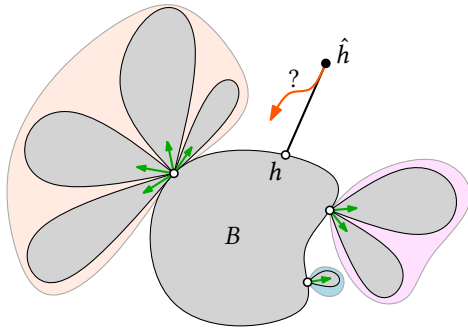
Figure 19: The traversal of the tree structure (b) during a farthest-point query (a). The bags containing farthest-points are colored red, the bags where two paths to farthest points split are colored yellow, and all other blocks are colored gray. In blue, we highlight the arcs visited during the query and, in green, we highlight a shortcut bypassing several gray bags.

397 (orange case). Third, all farthest points from \hat{h} could lie in a single bag cut $\text{bcut}(B, h')$ (black or orange case).
 398 In the first two cases, we introduce a trivial shortcut, i.e., the shortcut from $h \rightarrow B$ points to $h \rightarrow B$. In the
 399 third case, we inspect the information stored at the arcs incident to hinge h' to determine whether the
 400 farthest points from \hat{h} in $\text{bcut}(B, h')$ lie in multiple co-bag-cuts at h' (orange case) or in a single co-bag-cut
 401 co- $\text{bcut}(B', h')$ (black case). In the black case, we introduce a shortcut from $h \rightarrow B$ to the destination of
 402 the shortcut from $h' \rightarrow B'$. In this way, we obtain all shortcuts in the tree structure leading away from B^*
 403 without increasing our asymptotic bound on the size and construction time.


404 We employ our breadth-first-search propagation scheme to construct the shortcuts in the tree structure
 405 for arcs pointing towards B^* . At any hinge h of B^* , we are only missing the shortcut for the arc $h \rightarrow B^*$.
 406 With a query in $\text{per}(B^*)$, we can immediately determine whether the red, the orange or the black case from
 407 above applies; no shortcuts towards B^* are required for this step. With all shortcuts of the arcs from the
 408 hinges of B^* in place, we can compute the shortcuts towards B^* for the hinges of all blocks neighboring B^* ,
 409 then all hinges of the bags neighboring the neighbors of B^* and so forth.

410 Placing the shortcuts during the construction of our data structure for eccentricity queries takes only
 411 constant additional time and space per shortcut. How much time does a farthest point query take? First,
 412 we have to perform a farthest-point query in the perspective from the bag containing the query point.
 413 Then, we follow shortcuts towards bags containing more farthest points. The farthest-point queries in the
 414 perspective of subsequent bags B' take time linear in the number of reported farthest points in $\text{per}(B')$,
 415 since we query from pendant vertices of $\text{per}(B')$. Moreover, as all visited blocks are either orange or red,
 416 the number of visited blocks is linear in the number of red blocks, i.e., blocks containing farthest points.
 417 Altogether, this yields a query time of $O(k)$ from branches and $O(k + \log l)$ from blocks of size l .


418 **Theorem 11.** *Let G be a cactus network with n vertices. There is a data structure with $O(n)$ size and construction*
 419 *time supporting farthest-point queries on G in $O(k)$ time from branches and in $O(k + \log l)$ time from blocks of*
 420 *size l , where k is the number of reported farthest points.*




(a) The missing shortcut for $\text{co-bcut}(B, h)$.




(b) We do not place a shortcut when the perspective from B contains any farthest points within B itself.




(c) We do not introduce a shortcut when there are farthest points in several adjacent bag-cuts.



(d) When only one bag-cut contains farthest-points, we take a closer look at its co-bag-cuts.



(e) We do not introduce a shortcut when several co-bag-cuts contain farthest points.



(f) We only introduce a shortcut when only one co-bag-cut contains farthest points.

Figure 20: Determining a missing shortcut into a co-bag-cut of B at h , where the shortcuts for all co-bag cuts of the other hinges of B are known (a). We consider the farthest points from \hat{h} in $\text{per}(B)$: we place no shortcut when (b) there are farthest points in B , when (c) paths to farthest points split at B , and when (d,e) all farthest points lie in multiple co-bag-cuts of the same adjacent bag-cut. We only place a shortcut when (f) all farthest points lie in a single co-bag cut $\text{co-bcut}(B', h')$. In this case, the shortcut from $\text{co-bcut}(B, h)$ leads to the target of the shortcut into $\text{co-bcut}(B', h')$.

4 Conclusions and Future Work

In previous work [4], we obtain a data structure with construction time $O(m^2 \log n)$ for any network with n vertices and m edges, and with optimal query times for eccentricity queries and farthest-point queries. In this work, we improve the construction time to $O(n)$ for certain classes of networks without sacrificing optimal query time. In future work, we aim to achieve $o(m^2 \log n)$ construction time for more general classes of networks such as planar networks, k -almost-trees [10], and series parallel graphs [6].

References

- [1] Franz Aurenhammer, Rolf Klein, and Der-Tsai Lee. *Voronoi Diagrams and Delaunay Triangulations*. World Scientific, 2013. DOI: [10.1142/8685](https://doi.org/10.1142/8685).
- [2] Boaz Ben-Moshe, Binay Bhattacharya, Qiaosheng Shi, and Arie Tamir. “Efficient algorithms for center problems in cactus networks”. In: *Theoretical Computer Science* 378.3 (2007), pp. 237–252. DOI: [10.1016/j.tcs.2007.02.033](https://doi.org/10.1016/j.tcs.2007.02.033).
- [3] Prosenjit Bose, Jean-Lou De Carufel, Carsten Grimm, Anil Maheshwari, and Michiel Smid. “Optimal Data Structures for Farthest-Point Queries in Cactus Networks”. In: *Proceedings of the 25th Canadian Conference on Computational Geometry*. CCCG 2013. (Waterloo, Ontario, Canada, Aug. 8–10, 2013). 2013. URL: http://cccg.ca/proceedings/2013/papers/paper_23.pdf.
- [4] Prosenjit Bose, Kai Dannies, Jean-Lou De Carufel, Christoph Doell, Carsten Grimm, Anil Maheshwari, Stefan Schirra, and Michiel Smid. “Network Farthest-Point Diagrams”. In: *Journal of Computational Geometry* 4.1 (2013), pp. 182–211. URL: <http://jocg.org/v4n1p8/>.
- [5] Rainer E. Burkard and Jakob Krarup. “A linear algorithm for the pos/neg-weighted 1-median problem on a cactus”. In: *Computing* 60.3 (1998), pp. 193–215. DOI: [10.1007/BF02684332](https://doi.org/10.1007/BF02684332).
- [6] Richard J. Duffin. “Topology of Series-Parallel Networks”. In: *Journal of Mathematical Analysis and Applications* 10.2 (1965), pp. 303–318. DOI: [10.1016/0022-247X\(65\)90125-3](https://doi.org/10.1016/0022-247X(65)90125-3).
- [7] Martin Erwig. “The Graph Voronoi Diagram with Applications”. In: *Networks* 36.3 (2000), pp. 156–163. DOI: [10.1002/1097-0037\(200010\)36:3<156::AID-NET2>3.0.CO;2-L](https://doi.org/10.1002/1097-0037(200010)36:3<156::AID-NET2>3.0.CO;2-L).
- [8] Takehiro Furuta, Atsuo Suzuki, and Keisuke Inakawa. *The k -th Nearest Network Voronoi Diagram and its Application to Districting Problem of Ambulance Systems*. Technical Report 0501. Nanzan University, 2005. URL: <http://birdie.ic.nanzan-u.ac.jp/MCENTER/pdf/wp0501.pdf>.
- [9] Carsten Grimm. “Eccentricity Diagrams. On Charting Farthest-Point Information in Networks”. Diplomarbeit (Master’s Thesis). Magdeburg: Otto-von-Guericke Universität Magdeburg, Mar. 2012.
- [10] Yuri Gurevich, Larry J. Stockmeyer, and Uzi Vishkin. “Solving NP-Hard Problems on Graphs that are almost Trees and an Application to Facility Location Problems”. In: *Journal of the ACM* 31.3 (1984), pp. 459–473. DOI: [10.1145/828.322439](https://doi.org/10.1145/828.322439).
- [11] S. Louis Hakimi. “Optimum Locations of Switching Centers and the Absolute Centers and Medians of a Graph”. In: *Operations Research* 12.3 (1964), pp. 450–459. JSTOR: [168125](https://www.jstor.org/stable/168125).
- [12] S. Louis Hakimi, Martine Labbé, and Edward Schmeichel. “The Voronoi Partition of a Network and its Implications in Location Theory”. In: *ORSA Journal on Computing* 4.4 (1992), pp. 412–417. DOI: [10.1287/ijoc.4.4.412](https://doi.org/10.1287/ijoc.4.4.412).

- 459 [13] Pentti Hämäläinen. “The absolute Center of a unicyclic Network”. In: *Discrete Applied Mathematics*
460 25.3 (1989), pp. 311–315. DOI: [10.1016/0166-218X\(89\)90009-7](https://doi.org/10.1016/0166-218X(89)90009-7).
- 461 [14] Gabriel Y. Handler. “Minimax Location of a Facility in an undirected Tree Graph”. In: *Transportation*
462 *Science* 7.3 (1973), pp. 287–293. JSTOR: [25767706](https://www.jstor.org/stable/25767706).
- 463 [15] Pierre Hansen, Martine Labbé, and Brigitte Nicolas. “The Continuous Center Set of a Network”. In:
464 *Discrete Applied Mathematics* 30.2-3 (1991), pp. 181–195. DOI: [10.1016/0166-218X\(91\)90043-V](https://doi.org/10.1016/0166-218X(91)90043-V).
- 465 [16] Frank Harary and Geert Prins. “The Block-Cutpoint-Tree of a Graph”. In: *Publicationes Mathematicae*
466 *Debrecen* 13 (1966), pp. 103–107.
- 467 [17] John Hopcroft and Robert Tarjan. “Algorithm 447: Efficient Algorithms for Graph Manipulation”. In:
468 *Communications of the ACM* 16 (6 June 1973), pp. 372–378. DOI: [10.1145/362248.362272](https://doi.org/10.1145/362248.362272).
- 469 [18] Oded Kariv and S. Louis Hakimi. “An Algorithmic Approach to Network Location Problems. I: The
470 p -Centers”. In: *SIAM Journal on Applied Mathematics* 37.3 (1979), pp. 513–538. DOI: [10.1137/0137040](https://doi.org/10.1137/0137040).
- 471 [19] Rex K. Kincaid. “Exploiting Structure: Location Problems on Trees and Treelike Graphs”. In: *Founda-*
472 *tions of Location Analysis*. Springer US, 2011, pp. 315–334. DOI: [10.1007/978-1-4419-7572-0_14](https://doi.org/10.1007/978-1-4419-7572-0_14).
- 473 [20] Rex K. Kincaid and Timothy J. Lowe. “Locating an absolute center on graphs that are almost trees”.
474 In: *European Journal of Operational Research* 44.3 (1990), pp. 357–372. DOI: [10.1016/0377-2217\(90\)](https://doi.org/10.1016/0377-2217(90)90247-9)
475 [90247-9](https://doi.org/10.1016/0377-2217(90)90247-9).
- 476 [21] Mohammad Kolahdouzan and Cyrus Shahabi. “Voronoi-based K Nearest Neighbor Search for Spatial
477 Network Databases”. In: *Proceedings of the Thirtieth International Conference on Very Large Data Bases*.
478 VLDB ’04. (Toronto, Ontario, Canada, Aug. 31–Sept. 3, 2004). VLDB Endowment, 2004, pp. 840–851.
479 URL: <http://www.vldb.org/conf/2004/RS21P6.PDF>.
- 480 [22] Martine Labbé, Dominique Peeters, and Jacques-François Thisse. “Location on Networks”. In: *Network*
481 *Routing*. Vol. 8. Handbooks in Operations Research and Management Science. Elsevier, 1995. Chap. 7,
482 pp. 551–624. DOI: [10.1016/S0927-0507\(05\)80111-2](https://doi.org/10.1016/S0927-0507(05)80111-2).
- 483 [23] Yu-Feng Lan, Yue-Li Wang, and Hitoshi Suzuki. “A linear-time Algorithm for solving the Center
484 Problem on weighted Cactus Graphs”. In: *Information Processing Letters* 71.5–6 (1999), pp. 205–212.
485 DOI: [10.1016/S0020-0190\(99\)00111-8](https://doi.org/10.1016/S0020-0190(99)00111-8).
- 486 [24] Atsuyuki Okabe, Barry Boots, Kokichi Sugihara, and Sung Nok Chiu. *Spatial Tessellations: Concepts*
487 *and Applications of Voronoi Diagrams*. 2nd Edition. Probability and Statistics. John Wiley & Sons, Ltd.,
488 2000. DOI: [10.1002/9780470317013](https://doi.org/10.1002/9780470317013).
- 489 [25] Atsuyuki Okabe, Toshiaki Satoh, Takehiro Furuta, Atsuo Suzuki, and K. Okano. “Generalized Network
490 Voronoi Diagrams: Concepts, Computational Methods, and Applications”. In: *International Journal of*
491 *Geographical Information Science* 22.9 (2008), pp. 965–994. DOI: [10.1080/13658810701587891](https://doi.org/10.1080/13658810701587891).
- 492 [26] Atsuyuki Okabe and Kokichi Sugihara. *Spatial Analysis Along Networks: Statistical and Computational*
493 *Methods*. Statistics in Practice. John Wiley & Sons, Ltd., 2012. DOI: [10.1002/9781119967101](https://doi.org/10.1002/9781119967101).
- 494 [27] Atsuyuki Okabe and Atsuo Suzuki. “Locational Optimization Problems solved through Voronoi
495 Diagrams”. In: *European Journal of Operational Research* 98.3 (1997), pp. 445–456. DOI: [10.1016/S0377-](https://doi.org/10.1016/S0377-2217(97)80001-X)
496 [2217\(97\)80001-X](https://doi.org/10.1016/S0377-2217(97)80001-X).

- 497 [28] Dimitris Papadias, Jun Zhang, Nikos Mamoulis, and Yufei Tao. “Query Processing in Spatial Network
498 Databases”. In: *Proceedings of the 29th International Conference on Very Large Data Bases*. VLDB ’03.
499 (Berlin, Germany, Sept. 9–12, 2003). VLDB Endowment, 2003, pp. 802–813. URL: <http://www.vldb.org/conf/2003/papers/S24P02.pdf>.
500
- 501 [29] Charles S. ReVelle and Horst A. Eiselt. “Location Analysis: A Synthesis and Survey”. In: *European
502 Journal of Operational Research* 165.1 (2005), pp. 1–19. DOI: [10.1016/j.ejor.2003.11.032](https://doi.org/10.1016/j.ejor.2003.11.032).
- 503 [30] Qiaosheng Shi. “Efficient Algorithms for Network Center/Covering Location Optimization Prob-
504 lems”. PhD thesis. Burnaby, British Columbia, Canada: School of Computing Science, Simon Fraser
505 University, 2008. URL: <http://summit.sfu.ca/item/8893>.
- 506 [31] David Taniar and Wenny Rahayu. “A Taxonomy for Nearest Neighbour Queries in Spatial Databases”.
507 In: *Journal of Computer and System Sciences* 79.7 (2013), pp. 1017–1039. DOI: [10.1016/j.jcss.2013.
508 01.017](https://doi.org/10.1016/j.jcss.2013.01.017).
- 509 [32] Barbaros Ç. Tansel. “Discrete Center Problems”. In: *Foundations of Location Analysis*. International
510 Series in Operations Research & Management Science. Springer US, 2011, pp. 79–106. DOI: [10.1007/
511 978-1-4419-7572-0_5](https://doi.org/10.1007/978-1-4419-7572-0_5).
- 512 [33] Barbaros Ç. Tansel, Richard L. Francis, and Timothy J. Lowe. “Location on Networks: A Survey.
513 Part I: The p -Center and p -Median Problems”. In: *Management Science* 29.4 (1983), pp. 482–497. DOI:
514 [10.1287/mnsc.29.4.482](https://doi.org/10.1287/mnsc.29.4.482).
- 515 [34] Barbaros Ç. Tansel, Richard L. Francis, and Timothy J. Lowe. “Location on Networks: A Survey.
516 Part II: Exploiting Tree Network Structure”. In: *Management Science* 29.4 (1983), pp. 498–511. DOI:
517 [10.1287/mnsc.29.4.498](https://doi.org/10.1287/mnsc.29.4.498).