Theta-3 is connected

Oswin Aichholzer* Sang Won Bae † Luis Barba ‡ § Prosenjit Bose ‡ Matias Korman ¶
André van Renssen ‡ Perouz Taslakian ‖ Sander Verdonschot ‡

Abstract

In this paper, we show that the θ-graph with three cones is connected. We also provide an alternative proof of the connectivity of the Yao-graph with three cones.

1 Introduction

Introduced independently by Clarkson [6] in 1987 and Keil [9] in 1988, the θ-graph of a set P of points in the plane is constructed as follows. We consider each point $p \in P$ and partition the plane into $m \geq 2$ cones (regions in the plane between two rays originating from the same point) with apex p, each defined by two rays at consecutive multiples of $2\pi/m$ radians from the negative y-axis. We label the cones C_0 through C_{m-1}, in clockwise order around p, starting from the cone containing the positive y-axis from p if m is odd, or having this axis as its left boundary if m is even; see Figure 1. If the apex is not clear from the context, we use C_i^p to denote the cone C_i with apex p. We sometimes referred to C_i^p as the *i-cone* of p. To build the θ-graph, we consider each point p and connect it by an edge with the *closest* point in each of its cones. We measure distance by projecting each point onto the bisector of that cone instead of using the Euclidean distance. We use this definition of distance in the remainder of the paper, except for Section 4, which deals with Yao graphs. For simplicity, we assume that no two points of P lie on a line parallel to either the boundary or the angle bisector of a cone, guaranteeing that each point connects to at most one point in each cone. We call the θ-graph with m cones the θ_m-graph.

For θ-graphs with an even number of cones, proving connectedness is easy. As the first $m/2$ cones cover exactly the right half-plane, each point will have an edge to a point to its right, if such a point exists. Thus, we can find a path from any point to the rightmost point and, by concatenating these, a path between any pair of points. Unfortunately, if m is odd this property does not hold, as no set of cones covers *exactly* the right half-plane. Therefore, a point is not guaranteed to have an edge to a point to its right, even if such point exists.

The fact that θ-graphs with more than 6 cones are connected has been known for a long time. In fact, they even guarantee the existence of a *short* path between every pair of points. The length of this path is bounded by a constant times the straight-line Euclidean distance between the two points [3, 5, 6, 9, 11]. Graphs that have this property are called *geometric t-spanners* for some constant $t > 0$. For more information on geometric t-spanners, see the book by Narasimhan and Smid [10].

For a long time, very little was known about θ-graphs with fewer than 7 cones. Bonichon *et al.* [2] broke ground in this area in 2010, by showing that the θ_6-graph is a geometric spanner. Subsequently, both the θ_4- and θ_5-graphs have been shown to be constant spanners [1, 4]. El Molla [8] already showed that the θ_2- and θ_3-graphs are not constant spanners. The θ_3-graph is the last θ-graph for which connectedness has not been proven. In this paper, we settle this question by proving that the θ_3-graph is always connected.

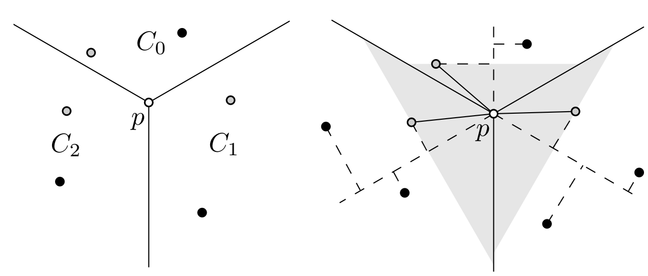

Figure 1: Left: A point p and its three cones in the θ_3-graph. Right: Point p adds an edge to the closest point in each of its cones, where distance is measured by projecting points onto the bisector of the cone.

*Institute for Software Technology, Graz University of Technology. Research of OA partially supported by the ESF EUROCORES programme EuroGIGA - CRP 'ComPoSe', Austrian Science Fund (FWF): I648-N18.

†Department of Computer Science, Kyonggi University. Work by S.W. Bae was supported by the Contents Convergence Software Research Center funded by the GRRC Program of Gyeonggi Province, South Korea.

‡School of Computer Science, Carleton University. Research supported in part by NSERC.

§Boursier FRIA du FNRS, Département d'Informatique, Université Libre de Bruxelles

¶Departament de Matemática Aplicada II, Universitat Politécnica de Catalunya. Received support of the Secretary for Universities and Research of the Ministry of Economy and Knowledge of the Government of Catalonia, the European Union, and projects MINECO MTM2012-30951, Gen. Cat. DGR2009SGR1040, ESF EUROCORES programme EuroGIGA – CRP 'ComPoSe': MICINN Project EUI-EURC-2011-4306.

‖College of Science and Engineering, American University of Armenia

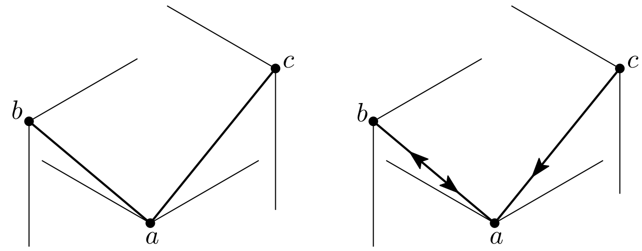

Figure 2: Left: A point set for which θ-routing does not find a path from a to c, as it keeps cycling between a and b. Right: The directed version of the graph is not strongly connected, as there is no path from either a or b to c.

The question of connectedness about the θ_3 graph is interesting because the θ_3-graph has some unique properties that cause standard proof techniques for θ-graphs to fail. As such, we hope that the techniques we develop here will lead to more insight into the structure of other θ-graphs. As an example, most proofs for a larger number of cones show that the θ-routing algorithm (always follow the edge in the cone that contains the destination) returns a short path between any two points. But in the θ_3-graph, θ-routing is not guaranteed to ever reach the destination. The smallest point set that exhibits this behavior has three points, such that for each point, both other points lie in the same cone; see Figure 2. In fact, this example shows not only that this exact routing strategy fails; it shows that if we consider the edges to be directed (from the point that added them, to the closest point in its cone), the graph is not strongly connected. Our proof requires more global methods than previous proofs on θ-graphs.

Most proofs for a larger number of cones use induction on the distance between points or on the size of the empty triangle between a point and its closest point. In the θ_3-graph however, both of these measures can increase when we follow an edge. Thus, applying induction on these distances seems a difficult task. An induction on the number of points similarly fails, as inserting a new point may remove edges that were present before, and it is not obvious that the endpoints of those edges are still connected in the new graph.

The θ_3-graph is strongly related to the Yao-3-graph, where each point also connects to the closest point in each cone, but the distance measure is the standard Euclidean distance. This graph was shown to be connected by Damian and Kumbhar [7]. Their proof uses induction on a rhomboid distance-measure that was tailored specifically for the Yao-3-graph. Since the 'closest' point for the θ_3-graph can be much further away than in the Yao-3-graph, this method of induction does not translate to the θ_3-graph either. Conversely, we show that our proof extends to the Yao-3-graph, providing an alternative proof for its connectivity.

2 Properties of the θ_3-graph

For $i \in \{0, 1, 2\}$, the edge connecting a point with its closest point in cone C_i is called an *i-edge*. Note that an edge can have one or two roles depending on the position of its endpoints. An example is depicted in Figure 2, where edge ab is both the 0-edge of a and the 1-edge of b.

Lemma 1 *For all $i \in \{0, 1, 2\}$, no two i-edges of the θ_3-graph can cross.*

Proof. We consider only 0-edges of P; the proof is analogous for 1- and 2-edges. For a contradiction, assume that there are two 0-edges that cross at a point s. Call these edges u_1v_1 and u_2v_2, such that v_1 is in the 0-cone of u_1 and v_2 in the 0-cone of u_2. Assume without loss of generality that the y-coordinate of v_1 is smaller than that of v_2; see Figure 3 for an illustration. Because s lies on segments u_1v_1 and u_2v_2, s lies in the 0-cones of both u_1 and u_2. Therefore, the 0-cone of s is contained in the intersection of the 0-cones of u_1 and u_2. As v_1 lies in cone C_0 of s, point v_1 lies in cone C_0 of u_2 as well. Because we assumed that the y-coordinate of v_1 is less than that of v_2, we conclude that v_1 is closer to u_2 than v_2. Thus, the edge u_2v_2 is not a 0-edge yielding a contradiction. □

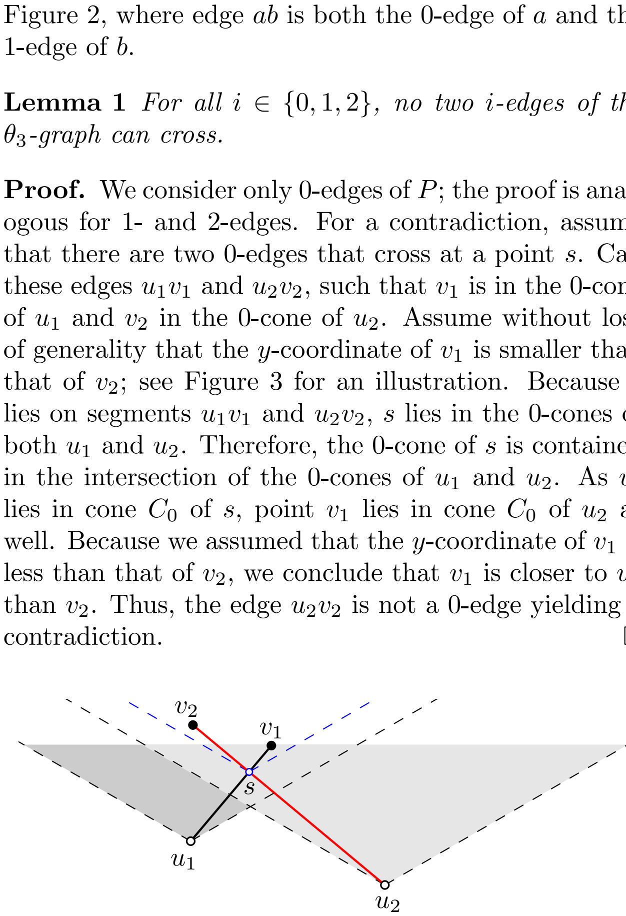

Figure 3: Two 0-edges u_1v_1 and u_2v_2 such that $v_1 \in C_0^{u_1}$ and $v_2 \in v_1 \in C_0^{u_2}$ cannot cross because the lowest point among v_1 and v_2 will be adjacent to both u_1 and u_2.

We say that a cone is *empty* if it contains no point of P in its interior. A point having an empty i-cone is called an *i-sink*.

Given a point p of P, the *i-path from p* is defined recursively as follows: If the i-cone of p is empty, the i-path from p consists of the single point p. Otherwise, let q be the closest point to p in its i-cone. The i-path from p is defined as the union of edge pq with the i-path from q.

Lemma 2 *Every i-path of the θ_3-graph is well-defined and has an i-sink as one of its endpoints.*

Proof. We consider only 0-paths; the proof is analogous for the other paths. A 0-path from a point p is well defined because the closest point in the 0-cone of p always lies above p. Therefore, the sequence of points

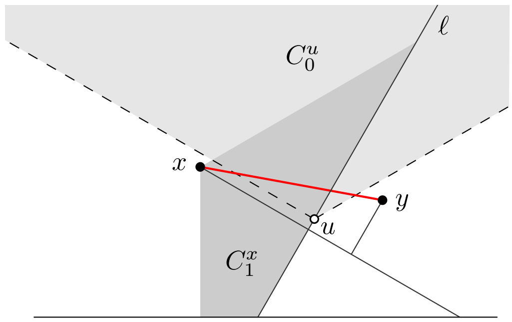

Figure 4: The last case in the proof of Lemma 3 where it is shown that empty i-cones cannot be crossed by edges of the θ_3-graph.

in the 0-path from p is monotonically increasing in the y-coordinate. Because P is a finite set, the depth of the recursion is finite and must end at a point having an empty 0-cone. □

Lemma 3 *If a cone of a point is empty, then no edge of the θ_3-graph can cross this cone.*

Proof. We consider only 0-cones for this proof; analogous arguments hold for the other cones. Let u be a point of P with an empty 0-cone. We prove the lemma by contradiction, so assume that there exists an edge xy that crosses C_0^u. Since no edge between two points in the same cone can cross another cone, assume without loss of generality that $x \in C_2^u$ and $y \in C_1^u$.

Note that y cannot lie in C_0^x, since either C_0^x does not intersect C_1^u (if $u \notin C_0^x$) or the line segment between x and y does not intersect C_0^u (if $u \in C_0^x$). Therefore, y lies in C_1^x.

If $u \in C_0^x$, then C_1^x does not intersect C_0^u and hence, the line segment between x and y cannot intersect C_0^u either. Therefore, both u and y lie in C_1^x. Let ℓ be the perpendicular to the bisector of C_1^x that passes through u. For the edge xy to exist, the projection of y on the bisector of C_1^x must be closer to x than that of u, i.e., y must lie to the left of ℓ. However, all points lying to the left of ℓ are contained in $C_0^u \cup C_2^u$ yielding a contradiction as $y \in C_1^u$; see Figure 4 for an illustration of this case. □

As a consequence of Lemmas 1 and 3, two sinks connected by an i-path partition the remaining points into two sets such that no i-path can connect a point in one set to a point in the other set, as any such path would cross either the i-path between the sinks, or the empty cone of one of the sinks. Such a construction is called an *i-barrier*; see Figure 5 for an illustration.

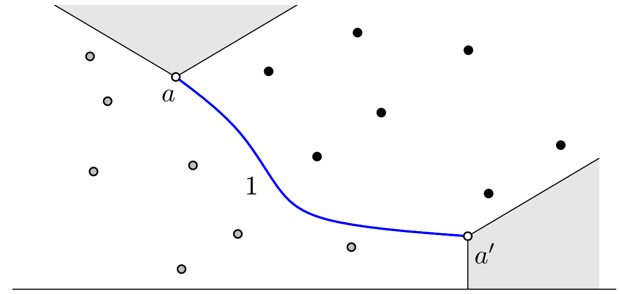

Figure 5: A 1-barrier, defined by the 1-path joining a with a', splits the remaining points into two sets such that no two points in different sets can be joined by a 1-path.

3 Proving connectedness

In this section we prove that the θ_3-graph of any given point set is connected. We start by proving that three given 0-sinks in a specific configuration are always connected. We then prove that if the θ_3-graph has at least two disjoint connected components, then there exist three 0-sinks that are in this configuration and are not all in the same component.

Although the edges of the θ_3-graph are not directed, by Lemma 2 we can think of an i-path as *oriented* towards the i-sink it reaches. An i-path oriented from point a to point b is denoted by $a \to b$. The following lemma is depicted in Figure 6.

Lemma 4 *Given three 0-sinks a, b, and c, such that (i) a lies to the left of b and b lies to the left of c, and (ii) the 1-path from a ends at a 1-sink a', whose 0-path ends at c, then a, b, and c belong to the same connected component.*

Proof. Since there is a path from a to c via a', a and c must be in the same component. We show that b belongs to this same connected component. The proof proceeds by induction on the number of 0-sinks to the right of c.

In the base case, there are no 0-sinks to the right of c. Consider the 1-sink b' at the end of the 1-path from b; see Figure 6 (right). If b' and a' are the same point,

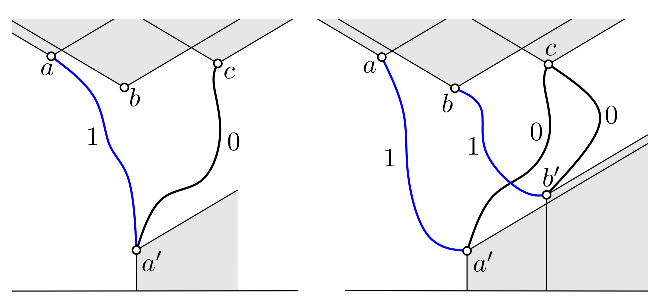

Figure 6: Left: The configuration of three 0-sinks described in Lemma 4. Right: The configuration in the base case of the induction where no 0-sink lies to the right of c.

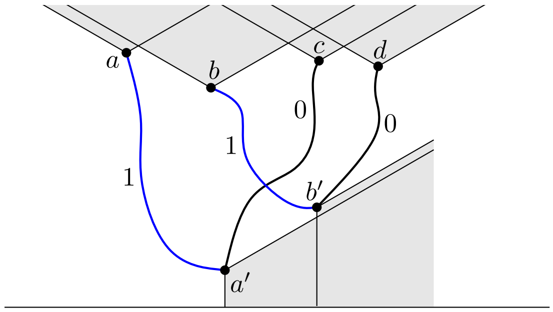

Figure 7: The configuration of the inductive step where the induction hypothesis can be applied on 0-sinks b, c and d.

then b is in the same connected component as a and we are done, so assume that this is not the case. Since the 1-path $a \to a'$ forms a 1-barrier, b' must lie to the right of a'. The 1-path $b \to b'$ also has to cross the 0-path $a' \to c$, as otherwise $a' \to c$ would cross the empty cone of b', which is impossible by Lemma 3. Because the 0-path $a' \to c$ forms a 0-barrier, the 0-path from b' cannot end up to the left of c. Moreover, since there are no 0-sinks to the right of c, the 0-path from b' must end at c. Thus, there is a path connecting b and c, which proves the lemma in the base case.

For the inductive step, let k be the number of 0-sinks to the right of c and assume that the lemma holds for any triple of 0-sinks with fewer than k 0-sinks to their right. By the same argument as in the base case, we have a 1-path from b to a 1-sink b' that lies to the right of a'. Now consider the 0-sink d at the end of the 0-path from b'; see Figure 7. Since the 0-path $a' \to c$ forms a 0-barrier, d cannot lie to the left of c. If d and c are the same point, we have a path connecting b and c as in the base case, so assume that this is not the case. Thus d lies to the right of c. Now b, c, and d form a triple of 0-sinks that satisfy criteria (i) and (ii). And since d is a 0-sink to the right of c, there are fewer than k 0-sinks to the right of d. Thus, by induction, we have that b is in the same connected component as c, which proves the lemma. □

Theorem 5 *The θ_3-graph is connected.*

Proof. Assume for sake of a contradiction that there exists a point set P whose θ_3-graph G is not connected. From each point, we can follow its 0-path to end up at a 0-sink. Therefore, G must contain at least one 0-sink for each connected component. Let a be the leftmost 0-sink, and let A be the connected component of G that contains a. Now let b be the leftmost 0-sink that does not belong to A.

We use Lemma 4 to show that, in fact, b must belong to A as well. Before we can do this, we need to define two barriers. The first barrier is formed by the 2-path from b, ending at a 2-sink b'. Because a lies in C_2^b, point b does not have an empty 2-cone and hence, b' differs from b. The second barrier is formed by the 0-path from b', which ends at a 0-sink c; see Figure 8. Since b is the leftmost 0-sink that does not belong to A, either c and b are the same point, or c lies to the right of b.

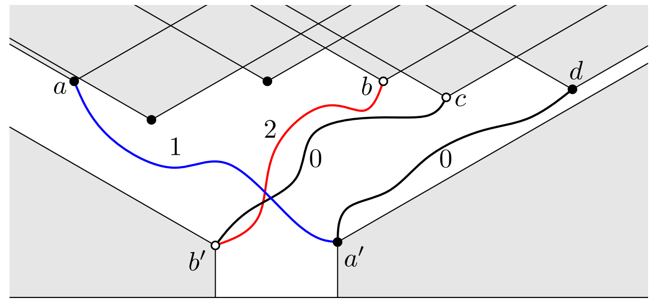

Figure 8: Two 0-sinks a and b are assumed to lie in different components such that both a and b are the leftmost 0-sinks in their component. The 1-path from a ends at a 1-sink a' whose 0-path ends at a 0-sink d lying to the right of b. The 0-sinks a, b and d jointly satisfy the criteria of Lemma 4.

Now consider the 1-sink a' at the end of the 1-path from a. This point has to lie to the right of both barriers $b \to b'$ and $b' \to c$, as otherwise these paths would cross the empty cone C_1 of a', which is not allowed by Lemma 3. Because the path $a \to a'$ is a 1-path and the barriers in question consist of 0- and 2-edges, these crossings are possible. Now let d be the 0-sink at the end of the 0-path from a'. Since this path cannot cross the 0-barrier $b' \to c$, d cannot lie to the left of c.

Because d belongs to component A, if c and d are the same point, c belongs to component A. Otherwise, if c and d are distinct points, then a, b, and d jointly satisfy the criteria of Lemma 4, which gives us that b belongs to component A as well—a contradiction since b is the leftmost 0-sink that does not belong to A. This contradiction comes from our assumption that G is not connected. Therefore, the θ_3-graph of any point set is connected. □

4 The Yao-3-graph

The construction of the Yao-3-graph is very similar to that of the θ_3-graph. The only difference is in the way distance is measured: the θ-graph uses the length of the projection onto the bisector, whereas the Yao-graph uses the Euclidean distance. Therefore, in every cone a point is connected to its closest Euclidean neighbor. We denote by $|pq|$ the Euclidean distance between two points p and q.

We show that like the θ_3-graph, the Yao-3-graph is also connected. To this end, we re-introduce the three basic lemmas we had for the θ_3-graph and show that the same properties hold for the Yao-3-graph.

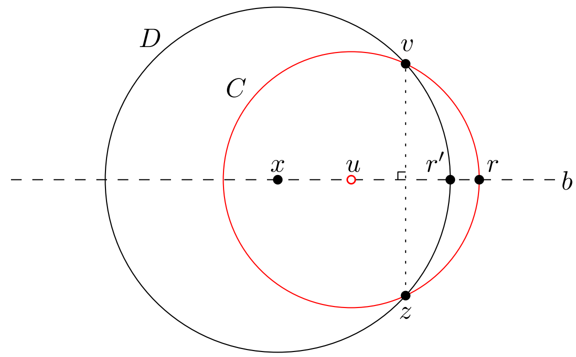

Figure 9: Point x lies to the left of point u and the arcs vr' and $r'z$ are enclosed by circle C centered at u, having radius $|uv|$.

We first prove a geometric auxiliary lemma depicted in Figure 9.

Lemma 6 *Given a non-vertical line b and a circle C centered at a point u on b, let v and z be two points on C such b bisects the segment vz. Let x be a point on b and let D be the circle centered at x having radius $|xv|$. If x lies to the left of u, then the right-side arc of D between v and z is enclosed by C; otherwise, the left-side arc of D between v and z is enclosed by C.*

Proof. Assume that x lies to the left of u; the proof of the other case is analogous. Let r and r' respectively be the right intersections of C and D with line b; see Figure 9. Hence, arcs vr' and $r'z$ lie either entirely inside C or entirely outside C. Therefore, it suffices to show that r' is enclosed by C, i.e., $|ur'| \leq |ur|$. Since x lies to the left of u, we can rewrite $|ur'|$ as $|xr'| - |xu|$. Since $|xr'| = |xv|$ and $|ur| = |uv|$, we thus need to show that $|xv| \leq |xu| + |uv|$. This follows from the triangle inequality. □

The proof of the following lemma is similar to that of Lemma 1.

Lemma 7 *For all $i \in \{0, 1, 2\}$, no two i-edges of the* Yao-3-*graph can cross.*

Proof. We look at the 0-edges. The cases for the other edges are analogous. Let uv be a 0-edge such that $v \in C_0^u$ and assume without loss of generality that v lies to the right of u. We prove the lemma by contradiction, so assume that some 0-edge xy crosses uv and let $y \in C_0^x$. Note that for xy to cross uv, C_0^x must contain some part of uv. Hence v lies in C_0^x.

Let k be the line through the right boundary of C_0^u and let l be the line through u such that the angle between l and the vertical line through u is $\pi/6$. We consider four cases, depending on the location of x with respect to u; see Figure 11 (left): (a) $x \in C_0^u$ to the left of the line uv, (b) $x \in C_2^u$ above k, (c) $x \in C_2^u$ below k or $x \in C_1^u$ below l, (d) $x \in C_1^u$ above l or $x \in C_0^u$ to the right of the line uv.

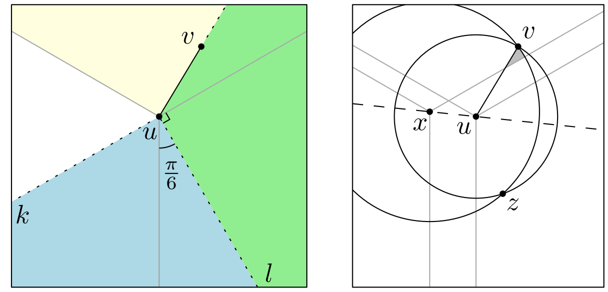

Figure 10: Left: The four cases. Right: The case when x lies in C_2^u and above k.

Case (a): $x \in C_0^u$ to the left of the line uv. Since v lies inside C_0^x and v lies to the right of u, x lies in the circle centered at u having radius $|uv|$. Thus, x lies closer to u than v, contradicting the existence of edge uv.

Case (b): $x \in C_2^u$ above k. We apply Lemma 6 as follows, see Figure 11 (right): Let C be the circle centered at u having radius $|uv|$. Let the line through u and x be bisector b, the bisector of v and z. Note that this implies that z lies outside C_0^u. Let D be the circle centered at x having radius $|xv|$. Since x lies to the left of u, Lemma 6 gives us that the right arc vz of circle D is enclosed by circle C. Since the area in which y must lie for xy to cross uv is bounded by the right boundary of C_0^x, edge uv, and the right arc vz of circle D, it is enclosed by C. Therefore, there does not exist a point $y \in C_0^x$ such that xy intersects uv.

Case (c): $x \in C_2^u$ below k or $x \in C_1^u$ below l. Since u lies in C_0^x, y needs to be closer to x than u for edge xy to exist. Hence it must lie inside the circle C centered at x having radius $|xu|$. Look at the lower half-plane defined by the line through u perpendicular to C and note that C is contained in this half-plane. However, the half-plane does not intersect C_0^u to the right of u and hence no point y inside the half-plane can be used to form an edge xy that crosses uv.

Case (d): $x \in C_1^u$ above l or $x \in C_0^u$ to the right of the line uv. We apply Lemma 6 as follows, see Figure 11 (right): Let C be the circle centered at u having radius $|uv|$. Let the line through u and x be bisector b. Note that z lies outside C_0^x. Let D be the circle centered at x having radius $|xv|$. Since x lies to the right of u, Lemma 6 gives us that the left arc vz of circle D is enclosed by circle C. Since the area in which y must lie for xy to cross uv is bounded by edge uv, the left arc vz of circle D, and either the left boundary of C_0^x (if $u \notin C_0^x$) or the line ux (if $u \in C_0^x$), it is enclosed by C. Therefore, there does not exist a point $y \in C_0^x$ such that xy intersects uv. □

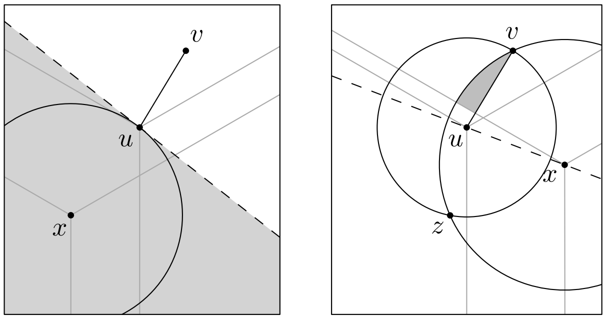

Figure 11: Left: The case when $x \in C_2^u$ below k or $x \in C_1^u$ below l. Right: The case when $x \in C_1^u$ above l or $x \in C_0^u$ to the right of the line uv.

Lemma 8 *Every i-path of the* Yao-3-*graph is well-defined and has an i-sink as one of its endpoints.*

Proof. The proof of this lemma is analogous to Lemma 2 for the θ_3-graph. □

Lemma 9 *If a cone of a point is empty, then no edge in the* Yao-3-*graph can cross this cone.*

Proof. We assume without loss of generality that C_0^u does not contain any points. We prove the lemma by contradiction, so assume that there exists an edge xy that crosses C_0^u. Since no edge between two points in the same cone can cross another cone, let $x \in C_2^u$ and $y \in C_1^u$.

Point y cannot lie in C_0^x, since either C_0^x does not intersect C_1^u (if $u \notin C_0^x$) or the line segment between x and y does not intersect C_0^u (if $u \in C_0^x$). Hence y must lie in C_1^x.

If $u \in C_0^x$, C_1^x does not intersect C_0^u and thus the line segment between x and y cannot intersect C_0^u either. Therefore both u and y lie in C_1^x. For the edge xy to exist, y must be closer to x than u, i.e., y must lie in the circle centered at x having radius $|xu|$. This circle is contained in the half-plane to the left of the line through u perpendicular to the circle.

If x lies on or above the horizontal line through u, the half-plane does not intersect C_1^u. If x lies below the horizontal line through u, the half-plane does not intersect C_1^u above u and thus xy would not cross C_0^u. Since y is enclosed by the circle, the circle is contained in the half-plane, and there is no point p in the half-plane such that $p \in C_1^u$ and px crosses C_0^u, xy cannot cross C_0^u either. □

Using Lemmas 7, 8 and 9, the proof of Theorem 5 translates directly to the Yao-3-graph yielding the following result.

Theorem 10 *The* Yao-3-*graph is connected.*

Acknowledgments. This problem was introduced during the Fields Workshop on Discrete and Computational Geometry held at Carleton University in Ottawa, Canada.

References

[1] L. Barba, P. Bose, J.-L. De Carufel, A. van Renssen, and S. Verdonschot. On the stretch factor of the Theta-4 graph. In *To appear in WADS*, 2013.

[2] N. Bonichon, C. Gavoille, N. Hanusse, and D. Ilcinkas. Connections between theta-graphs, Delaunay triangulations, and orthogonal surfaces. In *Graph-theoretic concepts in computer science*, volume 6410 of *Lecture Notes in Computer Science*, pages 266–278, 2010.

[3] P. Bose, J.-L. De Carufel, P. Morin, A. van Renssen, and S. Verdonschot. Optimal bounds on theta-graphs: More is not always better. In *Proceedings of CCCG*, pages 305–310, 2012.

[4] P. Bose, P. Morin, A. van Renssen, and S. Verdonschot. The θ_5-graph is a spanner. *To appear in the proceedings of WG'13*, 2013.

[5] P. Bose, A. van Renssen, and S. Verdonschot. On the spanning ratio of theta-graphs. In *Proceedings of WADS*, 2013.

[6] K. L. Clarkson. Approximation algorithms for shortest path motion planning. In *STOC*, pages 56–65, 1987.

[7] M. Damian and A. Kumbhar. Undirected connectivity of sparse yao graphs. In *FOMC*, pages 25–32, 2011.

[8] N. M. El Molla. *Yao spanners for wireless ad hoc networks*. PhD thesis, Villanova University, 2009.

[9] J. M. Keil. Approximating the complete Euclidean graph. In *SWAT*, volume 318 of *Lecture Notes in Computer Science*, pages 208–213, 1988.

[10] G. Narasimhan and M. Smid. *Geometric Spanner Networks*. Cambridge University Press, 2007.

[11] J. Ruppert and R. Seidel. Approximating the d-dimensional complete Euclidean graph. In *CCCG*, pages 207–210, 1991.