On Plane Constrained Bounded-Degree Spanners[*][†]

Prosenjit Bose[‡] Rolf Fagerberg[§] André van Renssen[¶]
Sander Verdonschot[‡]

Abstract

Let P be a finite set of points in the plane and S a set of non-crossing line segments with endpoints in P. The visibility graph of P with respect to S, denoted $Vis(P, S)$, has vertex set P and an edge for each pair of vertices u, v in P for which no line segment of S properly intersects uv. We show that the constrained half-θ_6-graph (which is identical to the constrained Delaunay graph whose empty visible region is an equilateral triangle) is a plane 2-spanner of $Vis(P, S)$. We then show how to construct a plane 6-spanner of $Vis(P, S)$ with maximum degree $6 + c$, where c is the maximum number of segments of S incident to a vertex.

1 Introduction

A geometric graph G is a graph whose vertices are points in the plane and whose edges are line segments between pairs of vertices. A graph G is called plane if no two edges intersect properly. Every edge is weighted by the Euclidean distance between its endpoints. The distance between two vertices u and v in G, denoted by $d_G(u, v)$ or simply $d(u, v)$ when G is clear from the context, is defined as the sum of the weights of the edges along the shortest path between u and v in G. A subgraph H of G is a t-spanner of G (for $t \geq 1$) if for each pair of vertices u and v, $d_H(u, v) \leq t \cdot d_G(u, v)$. The smallest value t for which H is a t-spanner is the *spanning ratio* or *stretch factor* of H. The graph G is referred to as the *underlying graph* of H. The spanning properties of various geometric graphs have been studied extensively in the literature (see [8, 13] for a comprehensive overview of the topic). However, most of the research has focused on constructing spanners where the underlying graph is the complete Euclidean geometric graph. We study this problem in a more general setting with the introduction of line segment *constraints*.

Specifically, let P be a set of vertices in the plane and let S be a set of line segments with endpoints in P, with no two line segments intersecting properly. The line segments of S are called *constraints*. Two vertices u and v can *see each other* if and only if either the line segment uv does not properly intersect any constraint or uv is itself a constraint. If two vertices u and v can see each other, the line segment uv is a *visibility edge*. The *visibility graph* of P with respect to a set of constraints S, denoted $Vis(P, S)$, has P as vertex set and all visibility edges as edge set. In other words, it is the complete graph on P minus all edges that properly intersect one or more constraints in S.

This setting has been studied extensively within the context of motion planning amid obstacles. Clarkson [10] was one of the first to study this problem and showed how to construct a linear-sized $(1+\epsilon)$-spanner of $Vis(P, S)$. Subsequently, Das [11] showed how to construct a spanner of $Vis(P, S)$

[*]An extended abstract of this paper appeared in the proceedings of the 10th Latin American Symposium on Theoretical Informatics (LATIN 2012) [5].

[†]This work is supported in part by the Natural Science and Engineering Research Council of Canada, Carleton University's President's 2010 Doctoral Fellowship, the Ontario Ministry of Research and Innovation, and the Danish Council for Independent Research, Natural Sciences, grant DFF-1323-00247, and JST ERATO Grant Number JPMJER1201, Japan.

[‡]Carleton University, Ottawa, Canada. jit@scs.carleton.ca, sander@cg.scs.carleton.ca

[§]University of Southern Denmark, Odense, Demark. rolf@imada.sdu.dk

[¶]University of Sydney, Sydney, Australia. andre.vanrenssen@sydney.edu.au

with constant spanning ratio and constant degree. Bose and Keil [7] showed that the Constrained Delaunay Triangulation is a 2.42-spanner of $Vis(P, S)$. In this article, we show that the constrained half-θ_6-graph (which is identical to the constrained Delaunay graph whose empty visible region is an equilateral triangle) is a plane 2-spanner of $Vis(P, S)$ by generalizing the approach used by Bose *et al.* [6]. This improves the upper bound on the spanning ratio of 36 implied by Bose *et al.* [4]. A key difficulty in proving this result stems from the fact that the constrained Delaunay graph is **not** necessarily a triangulation (see Figure 1). We then generalize the elegant construction of Bonichon *et al.* [2] to show how to construct a plane 6-spanner of $Vis(P, S)$ with maximum degree $6 + c$, where $c = \max\{c(v)|v \in P\}$ and $c(v)$ is the number of constraints incident to a vertex v.

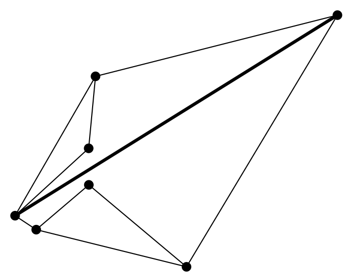

Figure 1: The constrained half-θ_6-graph is not necessarily a triangulation. The thick line segment represents a constraint

2 Preliminaries

We define a *cone* C to be the region in the plane between two rays originating from a vertex referred to as the apex of the cone. We let six rays originate from each vertex, with angles to the positive x-axis being multiples of $\pi/3$ (see Figure 2). Each pair of consecutive rays defines a cone. For ease of exposition, we only consider point sets in general position: no two vertices define a line parallel to one of the rays that define the cones and no three vertices are collinear. These assumptions imply that we can consider the cones to be open. If a point set is not in general position, one can easily find a suitable rotation of the point set to put it in general position.

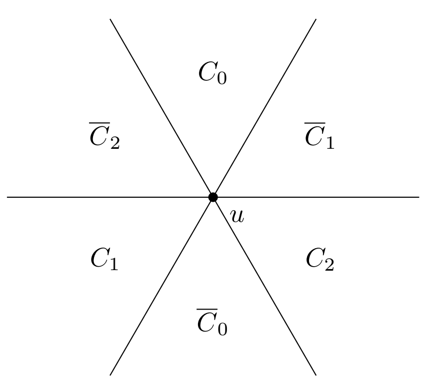

Figure 2: The cones having apex u

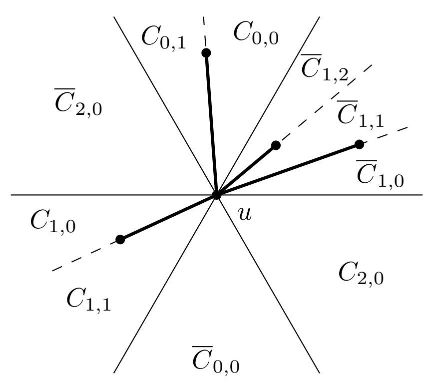

Figure 3: The subcones having apex u. Constraints are shown as thick line segments

Let $(\overline{C}_1, C_0, \overline{C}_2, C_1, \overline{C}_0, C_2)$ be the sequence of cones in counterclockwise order starting from the positive x-axis. The cones C_0, C_1, and C_2 are called *positive* cones and $\overline{C}_0$, $\overline{C}_1$, and $\overline{C}_2$ are called *negative* cones. By using addition and subtraction modulo 3 on the indices, positive cone C_i

has negative cone $\overline{C}_{i+1}$ as clockwise next cone and negative cone $\overline{C}_{i-1}$ as counterclockwise next cone. A similar statement holds for negative cones. We use C_i^u and $\overline{C}_j^u$ to denote cones C_i and $\overline{C}_j$ with apex u. Note that for any two vertices u and v, $v \in C_i^u$ if and only if $u \in \overline{C}_i^v$.

Let vertex u be an endpoint of a constraint c and let the other endpoint v lie in cone C_i^u. The lines through all such constraints c split C_i^u into several parts. We call these parts *subcones* and denote the j-th subcone of C_i^u by $C_{i,j}^u$, numbered in counterclockwise order (see Figure 3). When a constraint $c = (u, v)$ splits a cone of u into two subcones, we define v to lie in both of these subcones. We call a subcone of a positive cone a positive subcone and a subcone of a negative cone a negative subcone. We consider a cone that is not split to be a single subcone.

We now introduce the constrained half-θ_6-graph, a generalized version of the half-θ_6-graph as described by Bonichon *et al.* [1]: for each positive subcone of each vertex u, add an edge from u to the closest vertex in or on the boundary of that subcone that can see u, where distance is measured along the bisector of the original cone (not the subcone) (see Figure 4). More formally, we add an edge between two vertices u and v if v can see u, $v \in C_{i,j}^u$, and for all vertices $w \in C_{i,j}^u$ that can see u, $|uv'| \leq |uw'|$, where v' and w' denote the projection of v and w on the bisector of C_i^u and $|xy|$ denotes the length of the line segment between two vertices x and y. Note that our assumption of general position implies that each vertex adds at most one edge to the graph for each of its positive subcones.

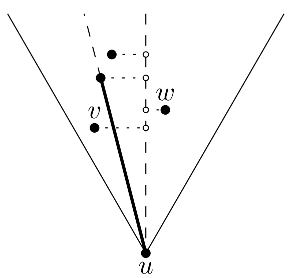

Figure 4: Three vertices are projected onto the bisector of a cone of u. Vertex v is the closest vertex in the left subcone and w is the closest vertex in the right subcone

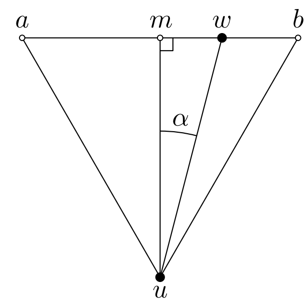

Figure 5: Canonical triangle T_{uw}

Given a vertex w in a positive cone C_i of vertex u, we define the *canonical triangle* T_{uw} to be the triangle defined by the borders of C_i^u and the line through w perpendicular to the bisector of C_i^u (see Figure 5). Note that for each pair of vertices there exists a unique canonical triangle. We say that a region is *empty* if it does not contain any vertices.

3 Spanning Ratio of the Constrained Half-θ_6-Graph

In this section we show that the constrained half-θ_6-graph is a plane 2-spanner of the visibility graph $Vis(P, S)$. To do this, we first prove a property of visibility graphs. Recall that a region is *empty* if it does not contain any vertices.

Lemma 1 *Let u, v, and w be three arbitrary points in the plane such that uw and vw are visibility edges and w is not the endpoint of a constraint intersecting the interior of triangle uvw. Then there exists a convex chain of visibility edges from u to v in triangle uvw, such that the polygon defined by uw, wv and the convex chain is empty and does not contain any constraints.*

Proof. Let Q be the set of vertices of $Vis(P, S)$ inside triangle uvw. If Q is empty, no constraint can cross uv, since one of its endpoints would have to be inside uvw, so our convex chain is simply uv. Otherwise, we build the convex hull of $Q \cup \{u, v\}$. Note that uv is part of the convex hull since

Q lies inside uvw to one side of the line through uv. When we remove this edge, we get a convex chain from u to v in triangle uvw. By the definition of a convex hull, the polygon defined by uw, wv and the convex chain is empty.

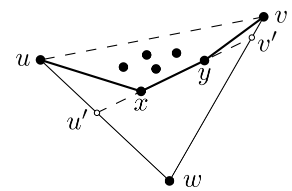

Figure 6: A convex chain from u to v and intersections u' and v' of the triangle and the line through x and y

Next, we show that two consecutive vertices x and y along the convex chain can see each other. Let u' be the intersection of uw and the line through x and y and let v' be the intersection of vw and the line through x and y (see Figure 6). Since w is not the endpoint of a constraint intersecting the interior of triangle uvw and, by construction, both u' and v' can see w, any constraint crossing xy would need to have an endpoint inside $u'wv'$. But the polygon defined by uw, wv and the convex chain is empty, so this is not possible. Therefore x can see y.

Finally, since the polygon defined by uw, wv and the convex chain is empty and consists of visibility edges, any constraint intersecting its interior needs to have w as an endpoint, which is not allowed. Hence, the polygon does not contain any constraints. $\square$

In the proof of Lemma 1, note that u, v, and w actually need not be part of the point set P. The lemma holds for any three points in the plane satisfying the requirements, if one considers the visibility edge as a line segment between any two points in the plane which is not intersected by a constraint. Lemma 1 will sometimes be used with this interpretation in mind later in the paper. Using this lemma, we proceed to improve the upper bound on the spanning ratio of the constrained half-θ_6-graph implied by Bose *et al.* [4].

Theorem 1 *Let u and w be vertices, with w in a positive cone of u, such that uw is a visibility edge. Let m be the midpoint of the side of T_{uw} opposing u, and let α be the unsigned angle between the lines uw and um. There exists a path connecting u and w in the constrained half-θ_6-graph of length at most $(\sqrt{3} \cdot \cos\alpha + \sin\alpha) \cdot |uw|$ that lies inside T_{uw}.*

Proof. We assume without loss of generality that $w \in C^u_{0,j}$. We prove the theorem by induction on the area of T_{uw}. Formally, we perform induction on the rank, when ordered by area, of the triangles T_{xy} for all pairs of vertices x and y that can see each other. Let $\delta(x, y)$ denote the length of the shortest path from x to y in the constrained half-θ_6-graph that lies inside T_{xy}. Let a and b be the upper left and right corner of T_{uw}, and let A and B be the triangles uaw and ubw (see Figure 7). Our inductive hypothesis is the following:

- If A is empty, then $\delta(u, w) \leq |ub| + |bw|$.
- If B is empty, then $\delta(u, w) \leq |ua| + |aw|$.
- If neither A nor B is empty, then $\delta(u, w) \leq \max\{|ua| + |aw|, |ub| + |bw|\}$.

We first note that this induction hypothesis implies the theorem: using the side of T_{uw} as the unit of length, we have that $\delta(u, w) \leq (\sqrt{3} \cdot \cos\alpha + \sin\alpha) \cdot |uw|$ (see Figure 8).

Base case: Triangle T_{uw} has minimal area. Since the triangle is a smallest canonical triangle, w is the closest vertex to u in its positive subcone. Hence the edge uw is in the constrained half-θ_6-graph, and $\delta(u, w) = |uw|$. From the triangle inequality, we have that $|uw| \leq \min\{|ua| + |aw|, |ub| + |bw|\}$, so the induction hypothesis holds.

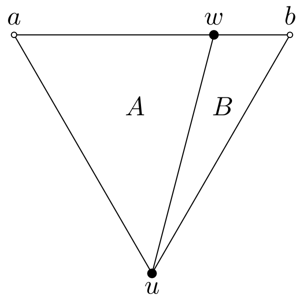

Figure 7: Triangles A and B

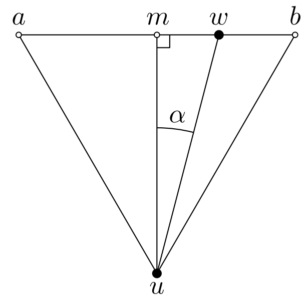

Figure 8: Canonical triangle T_{uw}

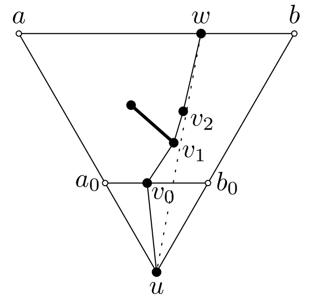

Figure 9: Convex chain from v_0 to w

Induction step: We assume that the induction hypothesis holds for all pairs of vertices that can see each other and have a canonical triangle whose area is smaller than the area of T_{uw}. If uw is an edge in the constrained half-θ_6-graph, the induction hypothesis follows by the same argument as in the base case. If there is no edge between u and w, let v_0 be the visible vertex closest to u in the positive subcone containing w, and let a_0 and b_0 be the upper left and right corner of T_{uv_0} (see Figure 9). By definition, $\delta(u, w) \leq |uv_0| + \delta(v_0, w)$, and by the triangle inequality, $|uv_0| \leq \min\{|ua_0| + |a_0v_0|, |ub_0| + |b_0v_0|\}$. We assume without loss of generality that v_0 lies to the left of uw, which means that A is not empty.

Since uw and uv_0 are visibility edges, by applying Lemma 1 to triangle v_0uw, a convex chain $v_0, ..., v_k = w$ of visibility edges connecting v_0 and w exists (see Figure 9). Note that, since v_0 is the closest visible vertex to u, every vertex along the convex chain lies above the horizontal line through v_0.

When looking at two consecutive vertices v_{i-1} and v_i along the convex chain, there are three types of configurations: (i) $v_{i-1} \in C_1^{v_i}$, (ii) $v_i \in C_0^{v_{i-1}}$ and v_i lies to the right of or has the same x-coordinate as v_{i-1}, (iii) $v_i \in C_0^{v_{i-1}}$ and v_i lies to the left of v_{i-1}. Let $A_i = v_{i-1}a_iv_i$ and $B_i = v_{i-1}b_iv_i$, the vertices a_i and b_i will be defined for each case. By convexity, the direction of $\overrightarrow{v_iv_{i+1}}$ is rotating counterclockwise for increasing i. Thus, these configurations occur in the order Type (i), Type (ii), and Type (iii) along the convex chain from v_0 to w. We bound $\delta(v_{i-1}, v_i)$ as follows (see Figure 10):

Type (i): If $v_{i-1} \in C_1^{v_i}$, let a_i and b_i be the upper left and lower corner of $T_{v_iv_{i-1}}$. Triangle B_i lies between the convex chain and uw, so it must be empty by Lemma 1. Since v_i can see v_{i-1} and $T_{v_iv_{i-1}}$ has smaller area than T_{uw}, the induction hypothesis gives that $\delta(v_{i-1}, v_i)$ is at most $|v_{i-1}a_i| + |a_iv_i|$.

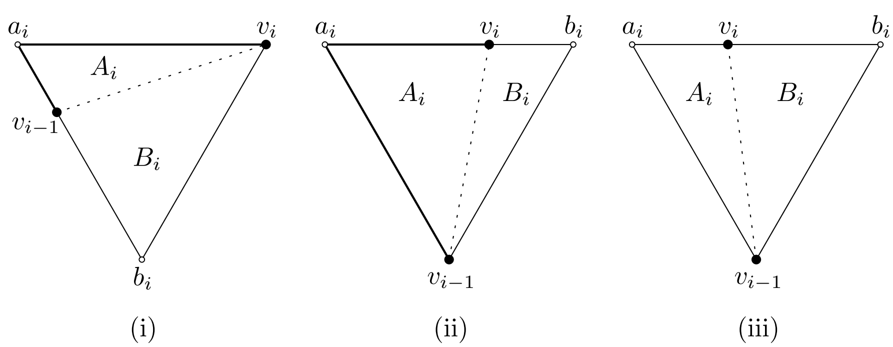

Figure 10: Charging the three types of configurations

Type (ii): If $v_i \in C_0^{v_{i-1}}$, let a_i and b_i be the left and right corner of $T_{v_{i-1}v_i}$. Since v_i can see

v_{i-1} and $T_{v_{i-1}v_i}$ has smaller area than T_{uw}, the induction hypothesis applies. Whether A_i and B_i are empty or not, $\delta(v_{i-1}, v_i)$ is at most $\max\{|v_{i-1}a_i| + |a_iv_i|, |v_{i-1}b_i| + |b_iv_i|\}$. Since v_i lies to the right of or has the same x-coordinate as v_{i-1}, we know $|v_{i-1}a_i| + |a_iv_i| \geq |v_{i-1}b_i| + |b_iv_i|$, so $\delta(v_{i-1}, v_i)$ is at most $|v_{i-1}a_i| + |a_iv_i|$.

Type (iii): If $v_i \in C_0^{v_{i-1}}$ and v_i lies to the left of v_{i-1}, let a_i and b_i be the left and right corner of $T_{v_{i-1}v_i}$. Since v_i can see v_{i-1} and $T_{v_{i-1}v_i}$ has smaller area than T_{uw}, we can apply the induction hypothesis. Thus, if B_i is empty, $\delta(v_{i-1}, v_i)$ is at most $|v_{i-1}a_i| + |a_iv_i|$ and if B_i is not empty, $\delta(v_{i-1}, v_i)$ is at most $|v_{i-1}b_i| + |b_iv_i|$.

Recall that a and b are the upper left and right corner of T_{uw} and that B is the triangle ubw (see Figure 7). To complete the proof, we consider three cases: (a) $\angle awu \leq \pi/2$, (b) $\angle awu > \pi/2$ and B is empty, (c) $\angle awu > \pi/2$ and B is not empty.

Case (a): If $\angle awu \leq \pi/2$, the convex chain cannot contain any Type (iii) configurations: for Type (iii) configurations to occur, v_i needs to lie to the left of v_{i-1}. However, by construction, v_i lies to the right of the line through v_{i-1} and w. Hence, since $\angle awv_{i-1} < \angle awu \leq \pi/2$, v_i lies to the right of v_{i-1}. We can now bound $\delta(u, w)$ as follows using the bounds on Type (i) and Type (ii) configurations outlined above (see Figure 11):

$$\begin{aligned}
\delta(u, w) &\leq |uv_0| + \sum_{i=1}^{k} \delta(v_{i-1}, v_i) \\
&\leq |ua_0| + |a_0v_0| + \sum_{i=1}^{k} (|v_{i-1}a_i| + |a_iv_i|) \\
&= |ua| + |aw|
\end{aligned}$$

We see that the latter is equal to $|ua| + |aw|$ as required.

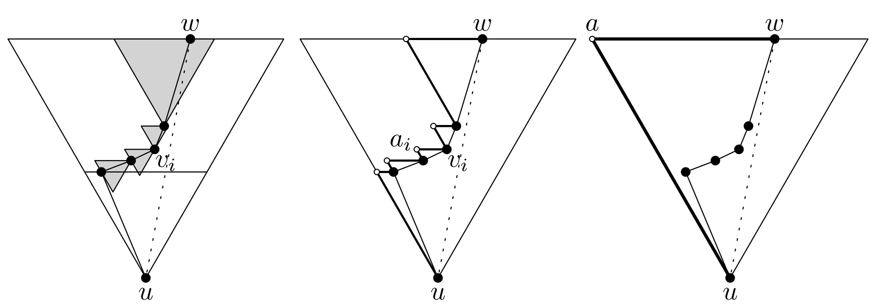

Figure 11: Visualization of the paths (thick lines) in the inequalities of case (a)

Case (b): If $\angle awu > \pi/2$ and B is empty, the convex chain can contain Type (iii) configurations. However, since B is empty and the area between the convex chain and uw is empty (by Lemma 1), all triangles B_i are also empty. Hence using the induction hypothesis, $\delta(v_{i-1}, v_i)$ is at most $|v_{i-1}a_i| + |a_iv_i|$ for all i. Using these bounds on the lengths of the paths between the vertices along the convex chain, we can bound $\delta(u, w)$ as in the previous case. Therefore, $\delta(u, w) \leq |ua| + |aw|$ as required.

Case (c): If $\angle awu > \pi/2$ and B is not empty, the convex chain can contain Type (iii) configurations. Since B is not empty, the triangles B_i need not be empty. Recall that v_0 lies in A, hence neither A nor B are empty. Therefore, it suffices to prove that $\delta(u, w) \leq \max\{|ua| + |aw|, |ub| + |bw|\} = |ub| + |bw|$. Let $T_{v_jv_{j+1}}$ be the first Type (iii) configuration along the convex chain (if it has any), let a' and b' be the upper left and right corner of T_{uv_j}, and let b'' be the upper right corner of T_{v_jw} (see Figure 12). Note that since $\angle awu > \pi/2$ and v_j lies to the left of uw, $|a'v_j|$ is smaller than $|b'v_j|$.

$$\begin{aligned}
\delta(u, w) &\leq |uv_0| + \sum_{i=1}^{k} \delta(v_{i-1}, v_i) \\
&\leq |ua_0| + |a_0v_0| + \sum_{i=1}^{j}(|v_{i-1}a_i| + |a_iv_i|) + \sum_{i=j+1}^{k}(|v_{i-1}b_i| + |b_iv_i|) \\
&= |ua'| + |a'v_j| + |v_jb''| + |b''w| \\
&\leq |ub'| + |b'v_j| + |v_jb''| + |b''w| \\
&= |ub| + |bw|
\end{aligned}$$

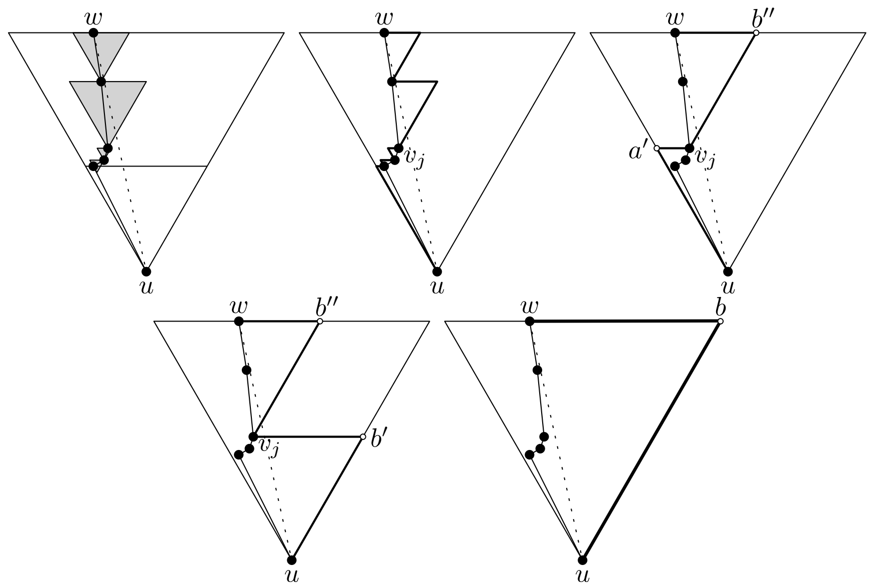

Figure 12: Visualization of the paths (thick lines) in the inequalities of case (c) □

Since the expression $\sqrt{3} \cdot \cos\alpha + \sin\alpha$ is increasing for $\alpha \in [0, \pi/6]$, the maximum value is attained by inserting the extreme value $\pi/6$. This leads to the following corollary.

Corollary 2 *The constrained half-θ_6-graph is a 2-spanner of the visibility graph.*

Next, we prove that the constrained half-θ_6-graph is plane.

Lemma 2 *Let u, v, x, and y be four distinct vertices such that the two canonical triangles T_{uv} and T_{xy} intersect. Then at least one of the corners of one canonical triangle is contained in the other canonical triangle.*

Proof. If one triangle contains the other triangle, it contains all of its corners. Therefore we focus on the case where neither triangle contains the other.

By definition, the upper boundaries of T_{uv} and T_{xy} are parallel, the left boundaries of T_{uv} and T_{xy} are parallel, and the right boundaries of T_{uv} and T_{xy} are parallel. Because we assume that no two vertices define a line parallel to one of the rays that define the cones, we assume, without loss of generality, that the upper boundary of T_{uv} lies below the upper boundary of T_{xy}. The upper boundary of T_{uv} must lie above the lower corner of T_{xy}, since otherwise the triangles do not intersect. If the upper left (right) corner of T_{uv} lies to the right (left) of the right (left) boundary of T_{xy}, the triangles cannot intersect. Hence, either one of the upper corners of T_{uv} is contained in T_{xy} or the upper boundary of T_{uv} intersects both the left and right boundary of T_{xy}. In the latter case, the fact that the left boundaries of T_{uv} and T_{xy} are parallel and the right boundaries of T_{uv} and T_{xy} are parallel, implies that the lower corner of T_{xy} is contained in T_{uv}. $\square$

Lemma 3 *The constrained half-θ_6-graph is plane.*

Proof. We prove the lemma by contradiction. Assume that two edges uv and xy cross at a point p. Since the two edges are contained in their canonical triangles, these triangles must intersect. By Lemma 2 we know that at least one of the corners of one triangle lies inside the other. We focus on the case where the upper right corner of T_{xy} lies inside T_{uv}. The other cases are analogous. Since uv and xy cross, this also means that either x or y must lie in T_{uv}. In the remainder, we assume that $y \in T_{uv}$. The arguments used for the case where $x \in T_{uv}$ are analogous.

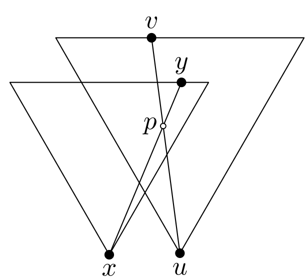

Figure 13: Edges uv and xy intersect at point p

Assume without loss of generality that $v \in C^u_{0,j}$ (see Figure 13). If $y \in C^u_{0,j}$, we look at triangle upy. Since both u and y can see p, we get by Lemma 1 that either u can see y or upy contains a vertex. In both cases, u can see a vertex in this subcone that is closer than v, contradicting the existence of the edge uv.

If $y \notin C^u_{0,j}$, there exists a constraint uz such that v lies to one side of the line through uz and y lies on the other side. Since this constraint cannot cross yp, z lies inside upy and is therefore closer to u than v. Since by definition z can see u, this also contradicts the existence of uv. $\square$

4 Bounding the Maximum Degree

In this section, we show how to construct a bounded degree subgraph G_9 of the constrained half-θ_6-graph that is a 6-spanner of the visibility graph. Given a vertex u and one of its negative subcones, we define the *canonical sequence* of this subcone as the vertices in this subcone that are neighbors of u in the constrained half-θ_6-graph, in counterclockwise order (see Figure 14). These vertices all have u as their closest visible vertex in a positive subcone. The *canonical path* is defined by connecting consecutive vertices in the canonical sequence. This definition differs slightly from the one used by Bonichon *et al.* [2].

To construct G_9, we start with a graph with vertex set P and no edges. Then for each negative subcone of each vertex $u \in P$, we add the canonical path and an edge between u and the closest vertex along this path, where distance is measured using the projections of the vertices onto the bisector of the cone containing the subcone. A given edge may be added by several vertices, but it appears only once in G_9. This construction is similar to the construction of the unconstrained degree-9 half-θ_6-graph described by Bonichon *et al.* [2]. We proceed to prove that G_9 is a spanning subgraph of the constrained half-θ_6-graph with spanning ratio 3.

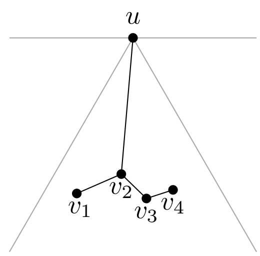

Figure 14: The edges that are added to G_9 for a negative subcone of a vertex u with canonical sequence v_1, v_2, v_3 and v_4

Lemma 4 *G_9 is a subgraph of the constrained half-θ_6-graph.*

Proof. Given a vertex u, we look at one of its negative subcones, say $\overline{C}^u_{0,j}$. The edges added to G_9 for this subcone can be divided into two types: edges of the canonical path, and the edge between u and the closest vertex along the canonical path. Since every vertex along the canonical path is by definition connected to u in the constrained half-θ_6-graph, it remains to show that the edges of the canonical path are part of the constrained half-θ_6-graph.

Let v and w be two consecutive vertices in the canonical path of $\overline{C}^u_{0,j}$, with v before w in counterclockwise order. By applying Lemma 1 on the visibility edges vu and wu, we get a convex chain $v = x_0, x_1, \ldots, x_{k-1}, x_k = w$ of $k \geq 1$ visibility edges, which together with vu and wu form a polygon Q empty of vertices and constraints.

Since Q is empty, v is not the endpoint of a constraint lying between vu and vx_1. Hence, x_1 cannot be in cone C^v_0, otherwise x_1 would be closer to v than u in the subcone of v that contains u. Similarly, x_{k-1} cannot lie in cone C^w_0. By convexity of the chain, this implies that no vertex on the chain can lie in cone C_0 of another vertex on the chain. Hence, since Q is empty, all vertices x_i can see u.

We first show that $k = 1$, i.e. that the chain is just the line vw. We prove this by contradiction, so assume that $k > 1$. Hence, there is at least one vertex x_i with $0 < i < k$. As such a vertex is not part of the canonical path in $\overline{C}^u_{0,j}$, it must see a closest vertex y different from u in the subcone of $C^{x_i}_0$ that contains u. As vertices on the chain cannot lie in C_0 of each other, y cannot be a vertex on the chain. As Q is empty, y must therefore lie strictly outside of Q, and yx_i must properly intersect either vu or wu. But this contradicts the planarity of the constrained half-θ_6-graph, as yx_i, vu, and wu would all be edges of this graph. Hence, $k = 1$ and the chain is a single visibility edge vw.

It remains to show that vw is an edge of the constrained half-θ_6-graph. Assume without loss of generality that w lies in C^v_2 (the case that v lies in C^w_1 is similar). We need to show that w is the closest visible vertex in subcone $C^v_{2,j}$. We prove this by contradiction, so assume another vertex x in $C^v_{2,j}$ is the closest. Vertex x lies in T_{vw}, which is partitioned into a part inside Q, a part to the right of wu, and a part below vw (see Figure 15). If x lies to the right of wu, we would have intersecting edges vx and wu, contradicting planarity of the constrained half-θ_6-graph. As Q is empty, x must lie below vw (see Figure 15).

Applying Lemma 1 on the visibility edges vx and vw, we get a convex chain $x = x_0, x_1, \ldots, x_{k-1}$, $x_k = w$ of visibility edges and an empty polygon R. Vertex x_1 cannot lie in C^x_0, as this would contradict that x is the closest visible vertex to v in $C^v_{2,j}$. Hence, since Q and R are empty, x can see u. Since v and w are two consecutive vertices in the canonical sequence of $\overline{C}^u_{0,j}$, x is not part of this canonical sequence. So it must see a closest vertex y different from u in the subcone of C^x_0 that contains u. Neither v nor the convex chain from x to w lie in C^x_0. As Q and R are empty, xy must properly intersect either vu or wu, contradicting the planarity of the constrained half-θ_6-graph. □

For future reference, we note that during the proof of Lemma 4 the following two properties were shown.

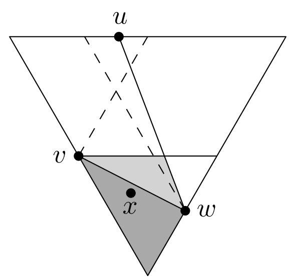

Figure 15: T_{vw} is partitioned into a part inside Q (light gray), a part to the right of wu (white), and a part below vw (dark gray)

Corollary 3 *Let u, v, and w be three vertices such that v and w are neighbors along a canonical path of u in $\overline{C}_i^u$. Vertex w cannot lie in C_i^v or $\overline{C}_i^v$.*

Corollary 4 *Let u, v, and w be three vertices such that v and w are neighbors along a canonical path of u in $\overline{C}_i^u$. Triangle uvw is empty and does not contain any constraints.*

Theorem 5 *G_9 is a 3-spanner of the constrained half-θ_6-graph.*

Proof. We prove the theorem by showing that for every edge uw in the constrained half-θ_6-graph, where w lies in a negative cone of u, G_9 contains a spanning path between u and w of length at most $3 \cdot |uw|$. This path will consist of a part of the canonical path in the subcone of u that contains w plus the edge between u and the closest canonical vertex in that subcone.

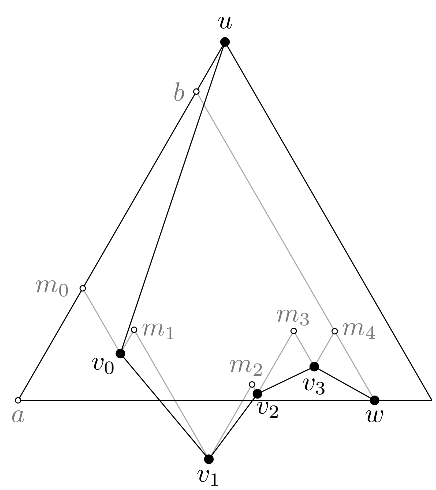

Figure 16: Bounding the length of the canonical path

We assume without loss of generality that $w \in \overline{C}_0^u$. Let v_0 be the vertex closest to u on the canonical path in the subcone $\overline{C}_{0,j}^u$ that contains w and let $v_0, v_1, ..., v_k = w$ be the vertices along the canonical path from v_0 to w (see Figure 16). Let l_j and r_j denote the rays defining the left and right boundaries of $C_0^{v_j}$ for $0 \leq j \leq k$ and let r denote the ray defining the right boundary of $\overline{C}_0^u$ (as seen from u). Let m_j be the intersection of l_j and r_{j-1}, for $1 \leq j \leq k$, and let m_0 be the intersection of l_0 and r. Let a be the intersection of r and the horizontal line through w and let b be the intersection of l_k and r. The length of the path between u and w in G_9 can now be bounded as follows:

$$\begin{aligned}
d_{G_9}(u,w) &\leq |uv_0| + \sum_{j=1}^{k} |v_{j-1}v_j| \\
&\leq |um_0| + |m_0v_0| + \sum_{j=1}^{k} |m_jv_j| + \sum_{j=1}^{k} |v_{j-1}m_j| \\
&= |um_0| + \sum_{j=0}^{k} |m_jv_j| + \sum_{j=1}^{k} |v_{j-1}m_j|
\end{aligned}$$

Since u lies in C_0 of each of the vertices along the canonical path, all m_jv_j project onto wb and all $v_{j-1}m_j$ project onto m_0b, when projecting along lines parallel to the boundaries of $\overline{C}_0^u$ instead of using orthogonal projections. By Corollary 3 no edge on the canonical path can lie in C_0 of one of its endpoints, hence the projections of m_jv_j onto wb do not overlap. For the same reason, the projections of $v_{j-1}m_j$ onto m_0b do not overlap. Hence, we have that $\sum_{j=0}^{k} |m_jv_j| = |wb|$ and $\sum_{j=1}^{k} |v_{j-1}m_j| = |m_0b|$.

$$\begin{aligned}
d_{G_9}(u,w) &= |um_0| + \sum_{j=0}^{k} |m_jv_j| + \sum_{j=1}^{k} |v_{j-1}m_j| \\
&= |um_0| + |wb| + |m_0b| \\
&\leq |ua| + 2 \cdot |wa|
\end{aligned}$$

Let α be $\angle auw$. Using some basic trigonometry, we get $|ua| = |uw| \cdot \cos\alpha + |uw| \cdot \sin\alpha/\sqrt{3}$ and $|wa| = 2 \cdot |uw| \cdot \sin\alpha/\sqrt{3}$. Thus the spanning ratio can be expressed as:

$$\frac{d_{G_9}(u,w)}{|uw|} \leq \cos\alpha + 5 \cdot \frac{\sin\alpha}{\sqrt{3}}$$

Since this is a non-decreasing function in α for $0 < \alpha \leq \pi/3$, its maximum value is obtained when $\alpha = \pi/3$, where the spanning ratio is 3. $\square$

It follows from Theorems 1 and 5 that G_9 is a 6-spanner of the visibility graph.

Corollary 6 *G_9 is a 6-spanner of the visibility graph.*

To bound the degrees of the vertices, we use a charging scheme that charges the edges of a vertex to its cones. Summing the charge for all cones of a vertex then bounds its degree.

Recalling that the edges of G_9 are generated by canonical paths, consider a canonical path in $\overline{C}_{i,j}^u$, created by a vertex u. We use v to indicate an arbitrary vertex along the canonical path, and we let v' be the closest vertex to u along the canonical path. The edges of G_9 generated by this canonical path are charged to cones as follows:

- The edge uv' is charged to $\overline{C}_i^u$ and to $C_i^{v'}$.
- An edge of the canonical path that lies in $\overline{C}_{i+1}^v$ is charged to C_i^v.
- An edge of the canonical path that lies in $\overline{C}_{i-1}^v$ is charged to C_i^v.
- An edge of the canonical path that lies in C_{i+1}^v is charged to $\overline{C}_{i-1}^v$.
- An edge of the canonical path that lies in C_{i-1}^v is charged to $\overline{C}_{i+1}^v$.

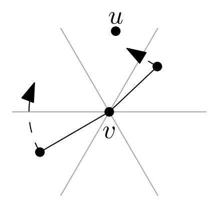

Figure 17: Two edges of a canonical path and the associated charges

Essentially, the edge between u and v' is charged to the cones that contain it and edges along the canonical path are charged to the adjacent cone that is closer to the cone of v that contains u. In other words, all charges are shifted one cone towards the positive cone containing u (see Figure 17).

By Corollary 3, no edge on the canonical path can lie in C_i^v or $\overline{C}_i^v$, so the charging scheme above is exhaustive. Note that each edge is charged to both of its endpoints and therefore the charge on a vertex is an upper bound on its degree (only an upper bound, since an edge can be generated and charged by several canonical paths).

Lemma 5 *Let v be a vertex that is incident to at least two constraints in the same positive cone C_i^v. Let $C_{i,j}^v$ be a subcone between two constraints and let u be the closest visible vertex in this subcone. Let $\overline{C}_{i,k}^u$ be the subcone of u that contains v and (when uv is a constraint) intersects $C_{i,j}^v$. Then v is the only vertex on the canonical path of $\overline{C}_{i,k}^u$.*

Proof. Let vw_1 and vw_2 be the two constraints between which subcone $C_{i,j}^v$ lies. By applying Lemma 1 on these visibility edges, we get a convex chain $w_1 = x_0, x_1, \ldots, x_k = w_2$ which together with vw_1 and vw_2 form a polygon $Q \subset C_{i,j}^v$ empty of vertices and constraints. Since u is the closest vertex visible to v inside $C_{i,j}^v$, u must be the vertex on this chain closest to v. In particular, it is at least as close to v as w_1 and w_2. Since vw_1 and vw_2 are constraints and Q is empty, there can be no vertex other than v in $\overline{C}_{i,k}^u$ from which u is visible. Hence, v is the only vertex on the canonical path of $\overline{C}_{i,k}^u$. □

Lemma 6 *Each positive cone C_i of a vertex v has a charge of at most* $\max\{2, c_i(v)+1\}$*, where $c_i(v)$ is the number of incident constraints in C_i^v.*

Proof. Let u be a vertex such that v is part of the canonical path of u. We first show that if this canonical path charges C_i^v, then u must lie in C_i^v. Assume u lies in C_j^v, $j \neq i$. Since all charges of this canonical path are shifted one cone towards C_j^v, a charge to C_i^v would have to come from $\overline{C}_j^v$. However, by Corollary 3, no edge on the canonical path of a vertex in C_j^v can lie in $\overline{C}_j^v$.

Next, we observe that there can be only one such vertex u for each subcone of C_i^v. This follows because v is only part of canonical paths of vertices u of which uv is an edge in the constrained half-θ_6-graph, and there is at most one edge for each positive subcone.

If C_i^v is a single subcone and v is not the closest vertex to u on its canonical path, C_i^v is charged for at most two edges along a single canonical path. Hence, its charge is at most 2. If v is the closest vertex to u, the negative cones adjacent to this positive cone cannot contain any vertices of the canonical path. If they did, these vertices would be closer to u than v is, as distance is measured using the projection onto the bisector of the cone of u. Hence, if v is the closest vertex to u, the positive cone containing u is charged 1. Thus, when the positive cone is a single subcone, the cone is charged 2 if it has an edge of the canonical path in each adjacent negative cone, and at most 1 otherwise.

Next, we look at the case where C_i^v is not a single subcone. For each subcone, except the first and last, the canonical path of the vertex u from that subcone consists only of v, by Lemma 5. Hence, we get a charge of 1 per subcone and a charge of at most $c_i(v) - 1$ in total for all subcones except the first and last subcone. We complete the proof by showing that the vertices u of the first and the last subcone can add a charge of at most 1 each.

Consider the first subcone $C^v_{i,0}$. The argument for the last subcone is symmetric. If v is the closest vertex to u on its canonical path, the negative cones adjacent to this positive cone cannot contain any vertices of the canonical path, since these would be closer to u than v is. Hence, the vertex u of this subcone adds a charge of 1. If v is not the closest vertex to u, we argue that v is the end of the canonical path of the vertex u of the subcone, implying that u can add a charge of at most 1: Let x be the other endpoint of the constraint that defines the subcone. Since u is the closest visible vertex in this subcone of v, it cannot lie further from v than x. If u is x, constraint uv splits $\overline{C}^u_i$ and only one of these two parts intersects the first subcone of v. Hence v is the end of the canonical path of u. If u is not x, u lies closer to v than x. Any vertex y before v (in counterclockwise order) on the canonical path would have to lie in C^v_{i+1} or $\overline{C}^v_{i-1}$, since by Corollary 3, y cannot lie in C^v_i or $\overline{C}^v_i$. Since y must also lie in $\overline{C}^u_i$ to be on this canonical path, vertex u is not visible from y due to the constraint xv. Hence, no such vertex y can exist on the canonical path, implying that v is the end of the canonical path.

Summing up all charges, each positive cone is charged at most $c_i(v) + 1$ if $c_i(v) \geq 1$, and at most 2 otherwise. Hence, a positive cone is charged at most $\max\{2, c_i(v) + 1\}$. $\square$

Corollary 7 *If the i-th positive cone of a vertex v has a charge of $c_i(v) + 2$, then $c_i(v) = 0$, i.e. it does not contain any constraints having v as an endpoint in C_i and is charged for two edges in the adjacent negative cones.*

Lemma 7 *Each negative cone $\overline{C}_i$ of a vertex v has a charge of at most $c_{\bar{i}}(v) + 1$, where $c_{\bar{i}}(v)$ is the number of incident constraints in $\overline{C}^v_i$.*

Proof. A negative cone of a vertex v is charged by the edge to the closest vertex in each of its subcones and it is charged by the two adjacent positive cones if edges of canonical paths lie in those cones (see Figure 18). We first show that vertices that do not lie in the positive subcones directly adjacent to $\overline{C}^v_i$ cannot have an edge involving v along their canonical paths. Let u be a vertex that does not lie in a positive subcone directly adjacent to $\overline{C}^v_i$ and let vx be the constraint closest to $\overline{C}^v_i$ that defines the boundary of the subcone of v that contains u. For u to have an edge along its canonical path that is charged to $\overline{C}^v_i$, it needs to lie further from u than x, since otherwise no vertex creating such an edge is visible to u. However, this implies that v would not connect to u, thus it would not part of the canonical path of u.

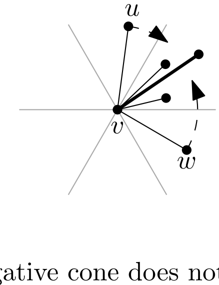

Figure 18: If vw is present, the negative cone does not contain edges having v as endpoint

As v can only be part of the canonical path of a single vertex in each of its positive subcones, we need to consider only the charges to $\overline{C}^v_i$ from the canonical path created by the closest visible vertices in the two positive subcones directly adjacent to $\overline{C}^v_i$. Let these vertices be u and w.

Next, we show that every negative cone can be charged by at most one edge in total from its adjacent positive cones. Suppose that w lies in a positive cone of v and vw is part of the canonical path of u. Then w lies in a negative cone of u, which means that u lies in a positive cone of w and cannot be part of a canonical path for w. It remains to show that this negative cone of v cannot be charged by an edge vu' from a canonical path of a different vertex w'. Since uvw forms a triangle in constrained half-θ_6-graph and this graph is planar, no edge of $u'vw'$ can cross any of the edges of uvw. This implies that either u' and w' lie inside uvw or u and w lie inside $u'vw'$. However, by Corollary 4, triangles xyz formed by a vertex x and two vertices y and z that are neighbors along the canonical path of x are empty. Therefore, u' and w' cannot lie inside uvw and u and w cannot lie inside $u'vw'$. Thus every negative cone charged by at most one edge in total from its adjacent positive cones.

Finally, we show that if one of uv or vw is present, the negative cone does not have an edge to the closest vertex in that cone and it contains no constraint that has v as an endpoint. We first show that if one of uv or vw is present, the negative cone does not have an edge to the closest vertex in that cone. We assume without loss of generality that vw is present, $u \in C_i^v \cap C_i^w$, and $w \in C_{i-1}^v$. Since v and w are neighbors on the canonical path of u, we know that the triangle uvw is part of the constrained half-θ_6-graph and, by Corollary 4, this triangle is empty. Furthermore, since uw is an edge of the constrained half-θ_6-graph and, by Lemma 3, the constrained half-θ_6-graph is plane, v cannot have an edge to the closest vertex beyond uw. Hence the negative cone does not have an edge to the closest vertex in that cone. By the same argument, the negative cone cannot contain a constraint that has v as an endpoint.

It follows that if this negative cone contains no constraint that has v as an endpoint, it is charged at most 1, by one of uv, vw, or the edge to the closest. Also, if this negative cone does contain constraints that have v as an endpoint, it is not charged by edges in the adjacent positive cones and hence its charge is at most $c_{\bar{i}}(v)+1$, one for the closest in each of its subcones. □

Theorem 8 *Every vertex v in G_9 has degree at most $c(v) + 9$.*

Proof. From Lemmas 6 and 7, each positive cone has charge at most $c_i(v) + 2$ and each negative cone has charge at most $c_{\bar{i}}(v) + 1$, where $c_i(v)$ and $c_{\bar{i}}(v)$ are the number of constraints in the i-th positive and negative cone. Since a vertex has three positive and three negative cones and the $c_i(v)$ and $c_{\bar{i}}(v)$ sum up to $c(v)$, this implies that the total degree of a vertex is at most $c(v)+9$. □

5 Bounding the Maximum Degree Further

In this section, we show how to reduce the maximum degree further, resulting in a plane 6-spanner G_6 of the visibility graph in which the degree of any node v is bounded by $c(v) + 6$.

By Lemmas 6 and 7 we see that if we can avoid the case where a positive cone gets a charge of $c_i(v)+2$, then every cone is charged at most $c_i(v)+1$, for a total charge of $c(v)+6$. By Corollary 7, this case only happens when a positive cone has $c_i(v) = 0$ and is charged for two edges in the adjacent negative cones. This situation is depicted in Figure 19, where x, v, and y are all on the canonical path of u. We construct G_6 by performing a transformation on G_9 for all positive cones in this situation.

We now describe the transformation. We assume without loss of generality that the positive cone in question is C_0^v. We call a vertex v the *closest canonical vertex* in a negative subcone of u when, among the vertices of the canonical path of u in that subcone, v is closest to u.

We first note that if x is the closest canonical vertex in one of the at most two subcones of $\overline{C}_2^v$ that contain it, the edge vx is charged to C_0^v, since vx is an edge of the canonical path induced by u, and it is also charged to cone $\overline{C}_2^v$, since it is the closest canonical vertex in one of its subcones. Since we need to charge it only once to account for the degree of v, we can remove the charge to C_0^v, reducing its charge by 1 as desired. Similarly, if y is the closest canonical vertex in one of the at most two subcones of $\overline{C}_1^v$ that contain it, it is charged to both C_0^v and $\overline{C}_1^v$, so we can reduce the charge to C_0^v by 1. Therefore, we only perform a transformation if neither x nor y is the closest canonical vertex in the subcones of v that contain them.

In that case, the transformation proceeds as follows. First, we add an edge between x and y. Next, we look at the sequence of vertices between v and the closest canonical vertex on the

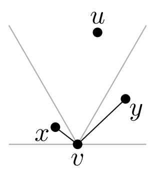

Figure 19: A positive cone having charge 2

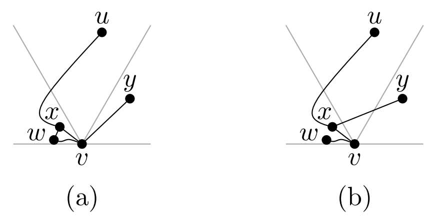

Figure 20: Transforming G_9 (a) into G_6 (b)

canonical path induced by u. If this sequence includes x, we remove vy. Otherwise we remove vx. Note that by Corollary 4, triangles uxv and uvy are empty and do not contain any constraints and therefore the edge xy does not intersect any constraints.

We assume without loss of generality that vy is removed. By removing vy and adding xy, we reduce the degree of v at the cost of increasing the degree of x. Hence, we need to find a way to balance the degree of x. Since x lies in $\overline{C}_2^v$ and the edge xv is part of the constrained half-θ_6-graph, x lies on a canonical path of v in $\overline{C}_2^v$ and, since x is not the closest canonical vertex to v on this canonical path, x has a neighbor w along this canonical path. We claim that x is the last vertex along the canonical path of v in $\overline{C}_2^v$ and thus w is uniquely defined. This follows because for any vertex z later than x along that canonical path, either z must lie in triangle uvx, contradicting its emptiness by Corollary 4, or the edges zv and xu of the constrained half-θ_6-graph must intersect, contradicting its planarity by Lemma 3. To balance the degree of x, we remove edge xw, if w lies in $\overline{C}_0^x$ and w is not the closest canonical vertex in a subcone of $\overline{C}_0^x$ that contains it. Otherwise xw is not removed. The situation before the transformation is shown in Figure 20 (a) and the situation after the transformation is shown in Figure 20 (b). A curved line segment denotes a part of a canonical path plus the edge from its closest canonical vertex.

To construct G_6, we apply this transformation on each positive cone matching the situation above. Note that since edge uv is part of the constrained half-θ_6-graph, which is plane, and G_9 is a subgraph of the constrained half-θ_6-graph, the edges added by this transformation cannot be part of G_9 as they cross uv. Hence, since only edges of G_9 are removed, there are no conflicts among the transformations of different cones, i.e. no cone will add an edge that was removed by another cone and vice versa. Before we prove that this construction yields a graph of maximum degree $6 + c$, we first show that the resulting graph is still a 3-spanner of the constrained half-θ_6-graph.

Lemma 8 *Let vx be an edge of G_9 and let x lie in a negative cone $\overline{C}_i$ of v. If x is not the closest canonical vertex in either of the at most two subcones of $\overline{C}_i^v$ that contain it, then the edge vx is used by at most one canonical path.*

Proof. We prove the lemma by contraposition. Assume that edge vx is part of two canonical paths of two vertices u and w. For v and x to be neighbors on a canonical path of u and w, these vertices need to lie in $C_{i+1}^v \cap C_{i+1}^x$ or $C_{i-1}^v \cap C_{i-1}^x$, by Corollary 3. By Corollary 4 and planarity of the constrained half-θ_6-graph, u and w cannot lie in the same region, hence one lies in $C_{i+1}^v \cap C_{i+1}^x$ and one lies in $C_{i-1}^v \cap C_{i-1}^x$. We assume without loss of generality that $u \in C_{i+1}^v \cap C_{i+1}^x$ and $w \in C_{i-1}^v \cap C_{i-1}^x$ (see Figure 21). Thus uvx and wvx form two disjoint triangles in the constrained half-θ_6-graph and, by Corollary 4, both triangles are empty. Furthermore, since the constrained half-θ_6-graph is plane, no edge from v can cross ux or wx, making vx the only edge of v in $\overline{C}_i$. Therefore, x is the closest canonical vertex in any subcone of $\overline{C}_i^v$ that contains it.

□

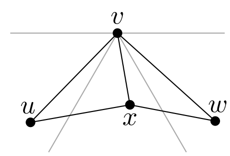

Figure 21: If edge vx is part of two canonical paths, x is the only neighbor of v in the negative cone of v

Lemma 9 *For every edge uw in the constrained half-θ_6-graph, there exists a path in G_6 of length at most $3 \cdot |uw|$.*

Proof. In the proof of Theorem 5 we showed that for every edge uw in the constrained half-θ_6-graph, where w lies in a negative cone of u, G_9 contains a spanning path between u and w of length at most $3 \cdot |uw|$, consisting of a part of the canonical path in the subcone of u that contains w plus the edge between u and the closest canonical vertex in that subcone. We now show that G_6 also contains a spanning path between u and w of length at most $3 \cdot |uw|$.

Note that in the construction, we never remove an edge vx with x being the closest canonical vertex in a negative subcone of v. This means two things: 1) For any spanning path in G_9, its last edge still exists in G_6. 2) By Lemma 8, any removed edge is part of a single canonical path, so we need to argue only about this single canonical path and the spanning paths using it.

During the construction of G_6, two types of edges are removed: Type 1, represented by vy in Figure 20, and Type 2, represented by xw in Figure 20. We first show that no edge removal of either of these types removes edge vx in Figure 20. A Type 1 removal that has v as the middle vertex in the configuration, as shown in Figure 20, is called *centered at v*. A Type 1 removal of vy affects the single canonical path containing xv and vy (see Figure 20). We note that no Type 1 removal involving v can be centered at x or y, since v lies in a positive cone of both x and y and a Type 1 removal requires both neighbors of the center vertex to lie in negative cones. This implies that Type 1 removals are non-overlapping (i.e. their configurations do not share edges) and, in particular, it implies that edge vx is not removed by this type of removal.

A Type 2 removal of xw affects the canonical path that contains w and x (see Figure 20). As argued during the construction of G_6, x is the last vertex along a canonical path of v and the edge xw is removed if w lies in a negative cone of x and w is not a closest canonical vertex to x. We now show that edge vx cannot be removed by a Type 2 removal: For it to be removed, we need that either x lies in a negative cone of v and v is the last vertex along this canonical path, or v lies in a negative cone of x and x is the last vertex along this canonical path. However, since v is not the last vertex along the canonical path that contains v and x (it is followed by y) and v does not lie in a negative cone of x, neither condition is satisfied.

Now that we know that for every Type 1 removal, edge vx is still present in the final G_6, we look at the spanning paths in G_6. Every spanning path present in G_9 can be affected by several non-overlapping Type 1 removals, as well as by a Type 2 removal at either end. By applying the triangle inequality to Figure 20, it follows that $|xy| \leq |xv| + |vy|$. Combined with the fact that for every Type 1 removal, vx is present in G_6, it follows that there still exists a spanning path between u and any vertex w along its canonical path, except possibly the last vertex x on either end, as the edge connecting x to its neighbor along the canonical path could be removed by a Type 2 removal. However, we perform a Type 2 removal only when u and x are part of a Type 1 configuration centered at u and ux is the edge of this configuration that was not removed (see Figure 20, where v acts as the node called u in the Type 2 argument above). Furthermore, we showed that in this case ux is still present in G_6. Hence, there exists a spanning path of length at most $3 \cdot |uw|$ between u and any vertex w along its canonical path.

Thus, we have proven that for every edge uw in the constrained half-θ_6-graph, where w lies in a negative cone of u, also G_6 contains a spanning path between u and w of length at most $3 \cdot |uw|$. $\square$

Lemma 10 *Every vertex v in G_6 has degree at most $c(v) + 6$.*

Proof. To bound the degree, we look at the charges of the vertices. We prove that after the transformation each positive cone has charge at most $c_i(v) + 1$ and each negative cone has charge at most $c_{\bar{i}}(v) + 1$. This implies that the total degree of a vertex is at most $c(v) + 6$. Since the charge of the negative cones is already at most $c_{\bar{i}}(v) + 1$, we focus on positive cones having charge $c_i(v) + 2$. By Corollary 7, this means that these cones have charge 2 and $c_i(v) = 0$.

Let v be a vertex such that one of its positive cones C_i^v has charge 2, let u be the vertex whose canonical path charged 2 to C_i^v, and let $x \in \overline{C}_{i-1}^v$ and $y \in \overline{C}_{i+1}^v$ be the neighbors of v on this canonical path (see Figure 19). If x or y is the closest canonical vertex in a subcone of $\overline{C}_{i-1}^v$ or $\overline{C}_{i+1}^v$, this edge has been charged to both that negative cone and C_i^v. Hence we can remove the charge to C_i^v while maintaining that the charge is an upper bound on the degree of v.

If neither x nor y is the closest canonical vertex in a subcone of $\overline{C}_{i-1}^v$ or $\overline{C}_{i+1}^v$, edge xy is added. We assume without loss of generality that edge vy is removed. Thus vy need not be charged, decreasing the charge of C_i^v to 1. Since vy was charged to $\overline{C}_{i-1}^y$ and this charge is removed, we charge edge xy to $\overline{C}_{i-1}^y$. Thus the charge of y does not change.

It remains to show that we can charge xy to x. Recall that x lies on the canonical path of v in $\overline{C}_{i-1}^v$, is the last vertex on this canonical path, and has w as neighbor on this canonical path (see Figure 20). Since vertices uvx and vwx each form a triangle of neighboring vertices on a canonical path in the constrained half-θ_6-graph, by Corollary 4 they are empty and do not contain any constraints. This implies that x is not the endpoint of any constraint in C_{i-1}^x. Hence, since x is the last vertex along the canonical path of v, C_{i-1}^x has charge at most 1 by Lemma 6 and Corollary 7. Now, consider the neighbor w of x. Vertex w can be in one of two cones with respect to x: C_{i+1}^x and $\overline{C}_i^x$. If $w \in C_{i+1}^x$, xw is charged to $\overline{C}_i^x$. Thus the charge of C_{i-1}^x is 0 and we can charge xy to it.

If $w \in \overline{C}_i^x$ and w is the closest canonical vertex to x in a subcone of $\overline{C}_i^x$, xw has been charged to both C_{i-1}^x and $\overline{C}_i^x$. We can remove that charge from C_{i-1}^x and instead charge xy to it, while keeping the charge of C_{i-1}^x at 1. If $w \in \overline{C}_i^x$ and w is not the closest canonical vertex to x in a subcone of $\overline{C}_i^x$ that contains it, xw was removed during the transformation. Since this edge was charged to C_{i-1}^x, we can now charge xy to C_{i-1}^x, while keeping the charge of C_{i-1}^x at 1. $\square$

Lemma 11 *G_6 is a plane subgraph of the visibility graph.*

Proof. Since G_9 is a plane subgraph of the visibility graph by Lemmas 3 and 4, only the edges added in the transformation from G_9 to G_6 can violate the lemma. An added edge xy can potentially intersect edges of G_6 that are in the constrained half-θ_6-graph, other edges that were added (recall that added edges are not in the constrained half-θ_6-graph, so these two cases are disjoint), and constraints.

First, we consider intersections of xy with edges of G_6 that are in the constrained half-θ_6-graph. Since xy was added in the transformation, x, y, and v are part of a canonical path of some vertex u (see Figure 20). Thus, in the constrained half-θ_6-graph uvx and uvy form two triangles, each containing a pair of neighboring vertices along the canonical path, which are empty by Corollary 4. Since the constrained half-θ_6-graph is plane and xy lies inside $uxvy$, the only edge of the constrained half-θ_6-graph that can intersect xy is uv. We now argue that uv is not in G_6. By construction, uv can only be part of G_9 if v is the closest vertex to u on this canonical path, or if uv are neighboring vertices along another canonical path of some vertex t. The former cannot be the case, by the conditions for adding xy in the transformation (see Figure 20). Assume the latter is the case. If $u \in C_i^v$, then either $t \in C_{i+1}^u \cap C_{i+1}^v$ or $t \in C_{i-1}^u \cap C_{i-1}^v$, by Corollary 3. If $t \in C_{i-1}^u \cap C_{i-1}^v$, the triangle uvt contains all of $\overline{C}_i^u \cap \overline{C}_{i+1}^v$, which contains y, as shown in Figure 22.

As uvt is empty by Corollary 4, this is a contradiction. If $t \in C_{i+1}^u \cap C_{i+1}^v$, a similar contradiction based on x arises. This shows that uv is not in G_9, and hence not in G_6 either, as edges added in the transformation from G_9 to G_6 are not in the constrained half-θ_6-graph.

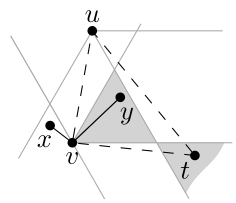

Figure 22: If $t \in C^u_{i-1} \cap C^v_{i-1}$, the triangle uvt contains all of $\overline{C}^u_i \cap \overline{C}^v_{i+1}$, which contains y

Next, we consider intersections of xy with other added edges. By Corollary 4 the quadrilateral $uxvy$ does not contain any vertices. Its sides ux, xv, vy, and yu are edges of the constrained half-θ_6-graph, which we showed above cannot be intersected by added edges. Hence, the only possibility for an added edge to intersect xy is the edge uv. However, uv cannot be an added edge, as we argued above. Thus, xy cannot intersect an added edge.

Finally, we consider intersection of xy with constraints. By Corollary 4, triangles uxv and uvy are empty and do not contain any constraints. Hence, since edge xy is contained in $uxvy$, it does not intersect any constraints. □

6 Conclusion

We showed that the constrained half-θ_6-graph is a plane 2-spanner of $Vis(P, S)$. We then generalized the construction of Bonichon *et al.* [2] to show how to construct a plane 6-spanner of $Vis(P, S)$ with maximum degree $6 + c$, where $c = \max\{c(v)|v \in P\}$ and $c(v)$ is the number of constraints incident to a vertex v.

A number of open problems still remain. For example, since constrained θ-graphs with at least 6 cones were recently shown to be spanners [9], a logical next question is to see if the method shown in this paper can be generalized to work for any constrained θ-graph. It would also be interesting to see if the degree can be reduced further still, while remaining a spanner of $Vis(P, S)$.

Furthermore, it would be interesting to see if it is possible to reduce the maximum degree of the vertices further. This was recently shown to be possible in the unconstrained setting [3, 12], which raises the question whether the approaches used in the unconstrained setting work in the constrained setting as well. Since these two approaches use different graphs as a starting point and thus require different edge removal rules and shortcutting techniques, it could very well be the case that only one of them results in a plane graph that respects the constraints.

References

[1] Nicolas Bonichon, Cyril Gavoille, Nicolas Hanusse, and David Ilcinkas. Connections between theta-graphs, Delaunay triangulations, and orthogonal surfaces. In *Proceedings of the 36th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2010)*, pages 266–278, 2010.

[2] Nicolas Bonichon, Cyril Gavoille, Nicolas Hanusse, and Ljubomir Perkovic. Plane spanners of maximum degree six. In *Proceedings of the 37th International Colloquium on Automata, Languages, and Programming (ICALP 2010)*, pages 19–30, 2010.

[3] Nicolas Bonichon, Iyad Kanj, Ljubomir Perković, and Ge Xia. There are plane spanners of degree 4 and moderate stretch factor. *Discrete & Computational Geometry*, 53(3):514–546, 2015.

[4] Prosenjit Bose, Jean-Lou De Carufel, and André van Renssen. Constrained generalized Delaunay graphs are plane spanners. In *Proceedings of the Computational Intelligence in Information Systems (CIIS 2016)*, volume 532 of *Advances in Intelligent Systems and Computing*, pages 281–293, 2016.

[5] Prosenjit Bose, Rolf Fagerberg, André van Renssen, and Sander Verdonschot. On plane constrained bounded-degree spanners. In *Proceedings of the 10th Latin American Symposium on Theoretical Informatics (LATIN 2012)*, volume 7256 of *Lecture Notes in Computer Science*, pages 85–96, 2012.

[6] Prosenjit Bose, Rolf Fagerberg, André van Renssen, and Sander Verdonschot. Optimal local routing on Delaunay triangulations defined by empty equilateral triangles. *SIAM Journal on Computing*, 44(6):1626–1649, 2015.

[7] Prosenjit Bose and J. Mark Keil. On the stretch factor of the constrained Delaunay triangulation. In *Proceedings of the 3rd International Symposium onVoronoi Diagrams in Science and Engineering (ISVD 2006)*, pages 25–31, 2006.

[8] Prosenjit Bose and Michiel Smid. On plane geometric spanners: A survey and open problems. *Computational Geometry: Theory and Applications*, 46(7):818–830, 2013.

[9] Prosenjit Bose and André van Renssen. Upper bounds on the spanning ratio of constrained theta-graphs. In *Proceedings of the 11th Latin American Symposium on Theoretical Informatics (LATIN 2014)*, volume 8392 of *Lecture Notes in Computer Science*, pages 108–119, 2014.

[10] Ken Clarkson. Approximation algorithms for shortest path motion planning. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing*, pages 56–65, 1987.

[11] Gautam Das. The visibility graph contains a bounded-degree spanner. In *Proceedings of the 9th Canadian Conference on Computational Geometry (CCCG 1997)*, pages 70–75, 1997.

[12] Iyad Kanj, Ljubomir Perkovic, and Duru Turkoglu. Degree four plane spanners: Simpler and better. *Journal of Computational Geometry*, 8(2):3–31, 2017.

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